

Piecewise Linear Morse Theory

Dissertation zur Erlangung des Doktorgrades

vorgelegt am

Fachbereich Mathematik und Informatik
der Freien Universität Berlin

2016

von

Romain Grunert

Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
rgrunert@inf.fu-berlin.de

Betreuer:

Prof. Dr. Günter Rote
Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
Germany
rote@inf.fu-berlin.de

Gutachter:

Prof. Dr. Günter Rote
Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
Germany
rote@inf.fu-berlin.de

Prof. Dr. Wolfgang Kühnel
Institut für Geometrie und Topologie
Universität Stuttgart
Pfaffenwaldring 57
70550 Stuttgart
Germany
kuehnel@mathematik.uni-stuttgart.de

Datum der Disputation: 18. November 2016

This research was partially supported by the DFG Collaborative Research
Center TRR 109, 'Discretization in Geometry and Dynamics.'

Contents

Preface	vii
1 Introduction	1
1.1 Smooth Morse Theory	2
1.2 Piecewise Linear Topology	5
1.2.1 Polyhedral Sets and Polyhedral Complexes	5
1.2.2 Interval Complexes	9
1.2.3 Operations on Complexes	10
1.2.4 Piecewise Linear Maps	15
1.2.5 Subcomplexes	18
1.2.6 Subdivisions	20
1.2.7 Combinatorial Equivalences and Piecewise Linear Homeomorphisms	21
1.2.8 Pseudo-Radial Projection	29
1.2.9 Combinatorial Manifolds	34
1.2.10 Homotopy and Homology	37
1.2.11 Pachner Moves and Recognition of 3-Spheres	41
1.2.12 Schoenflies Conjecture	44
1.3 Discrete Morse Theory	45
2 Level Set Complexes	49
2.1 Definitions and Basic Observations	50
2.1.1 Notation	50
2.1.2 Level Set Cells	52
2.1.3 Definition of Level Set Complexes	57
2.2 Canonical Representations	60
2.2.1 Definition of Canonical Representations and Importance for Determining the Combinatorial Structure of Level Set Complexes	60
2.2.2 Criteria Characterising Canonical Representations	63
2.3 Combinatorial Equivalences and Collapses	80
2.3.1 Natural Combinatorial Equivalences between Level Set Complexes	80
2.3.2 More Combinatorially Equivalent Level Set Complexes and First Conclusions	86
2.3.3 Collapsing Lower Level Set Complexes	91

3	Piecewise Linear Morse Functions	97
3.1	Local Equivalence	98
3.1.1	Definition and First Observations	99
3.1.2	Characterisations Based on Signs Instead of Values	104
3.2	Regular and Critical Points	110
3.2.1	Reference Triples	111
3.2.2	Constructing Link Complexes	116
3.2.3	Restricting Criticality to Vertices	120
3.2.4	Existence of PL Morse Functions and Their Construction from Discrete Morse Functions	122
3.2.5	Comparison with Previous Notions	132
4	More Characterisations for Regular Points	147
4.1	Unknotted Sphere Pairs	148
4.2	Level Preserving Isotopies	154
4.2.1	Definition and Basic Properties	154
4.2.2	Characterising Regular Points by Local Level Preserving Isotopies	156
4.2.3	Constructing Level Preserving Isotopies across Vertices	158
5	Piecewise Linear Analogues of Classical Results from Smooth Morse Theory	167
5.1	The Fundamental Results on the Behaviour of Level Sets	167
5.2	Consequences: Homotopy Type, Morse Inequalities, and Reeb Theorem	173
6	Computational Aspects	177
6.1	Recognising Regular Vertices	177
6.1.1	Algorithmic Solutions in Low Dimensions	178
6.1.2	Undecidability in Higher Dimensions	181
6.2	Computing Isotopies	190
6.2.1	General Remarks on the Computation of Level Preserving Isotopies and the Combinatorially Equivalent Subdi- visions Representing Them	191
6.2.2	Dimension 1	193
6.2.3	Dimension 2	193
6.2.4	Merging Isotopies	198
6.2.5	Dimension 3	200
6.2.6	Dimension 4	205

Preface

Morse theory was originally designed for the study of smooth manifolds by means of smooth maps defined on them. The results of Morse theory allow to infer global topological properties of the considered manifold from the local behaviour of a real-valued smooth map at critical points, provided that the map does not exhibit certain degeneracies at the critical points. Maps without such degeneracies are called Morse functions. With respect to such a map, each point of the manifold is either regular or non-degenerate critical. Points of the latter kind are classified according to their index. The fundamental results of Morse theory are obtained by examining level sets, in particular so called sublevel sets formed by all those points whose value is at most a certain threshold h . When the level h varies continuously, a change in topology of the corresponding level sets occurs only when the level crosses a critical value, that is, a value attained at some critical point. If only a single critical point is involved, the change in topology can be related to the index of that critical point.

The main objective of this thesis is a detailed development of the fundamental results of an analogous theory for piecewise linear manifolds with piecewise linear maps, with a focus on the question how regular points can be characterised in this setting and on the treatment of level sets. The special interest in the characterisation of regular points originates in some motivating open problems explained in the following.

Motivation. The initial motivation for the studies were some questions raised by Rote in an extended abstract [42] on piecewise linear Morse theory summarising a talk given at a workshop in Oberwolfach. The abstract resumes a discussion started in an earlier paper by Chiang, Lenz, Lu, and Rote on the construction of contour trees [13]. The paper contains two theorems concerning the characterisation of regular and critical points for piecewise linear maps on three-dimensional combinatorial manifolds and the absence of topological changes in the level sets of such maps when no critical values are crossed. It is remarked that although one would expect such statements to be well-known facts, no readily applicable results were found in the literature. The question arising in [13] and [42] from the first theorem characterising critical points for three-dimensional manifolds is the following: Which conditions characterise critical points of piecewise linear maps on combinatorial d -manifolds where d is greater than three? For the 4-dimensional case a characterisation is suggested and it is conjectured that the algorithmic question whether a given point is critical or not is undecidable for sufficiently large dimension. The second theorem expresses the absence of topological changes when the level h varies in inter-

vals $[a, b]$ without critical values by the existence of a level-preserving isotopy $\phi: M_{=b} \times [a, b] \rightarrow M$, where M is the considered manifold and $M_{=b}$ the level set consisting of all points with value b . It is noted that one part of the proof does not carry over to higher dimensions and that the constructed isotopy ϕ is only piecewise linear when restricted to a fixed level, that is to some subset $M_{=b} \times \{h\}$, but not necessarily piecewise linear on its whole domain. This raises the problem of modifying the construction in such a way that it proves the existence of a level-preserving isotopy for arbitrary dimensions. Moreover the constructed map defining the isotopy should be piecewise linear on its whole domain.

The results of the study of piecewise linear Morse theory in this thesis allow to address the aforementioned questions and problems from [13] and [42] as follows: In Chapters 3 and 4, several characterisations for regular and critical points, including some previously suggested notions, are proposed and most of them are shown to be equivalent. We also prove that it is undecidable for $d \geq 6$ whether a given point in a d -dimensional simplicial complex is regular or critical. This result is derived in Corollary 6.6 from Novikov's undecidability result for the $(d - 1)$ -sphere recognition problem with $(d - 1) \geq 5$. It remains open whether the problem of recognising regular points is still undecidable when the input is restricted to complexes that are combinatorial manifolds. For intervals without critical values, a modified construction is presented that defines a piecewise linear level-preserving isotopy for the level sets in that interval. In principle, the construction outlined in Subsection 4.2.3 works in any dimension, but it relies on homeomorphisms whose existence is ensured in theory by the defining properties of a regular point. In higher dimensions however, there are no algorithms available that would actually compute the theoretically ensured homeomorphisms for arbitrary inputs.

Previous Work. Apart from the motivating papers by Chiang, Lenz, Lu, and Rote [13] and by Rote [42], there are several other publications that deal with notions and results that can be attributed to piecewise linear Morse theory. Among the first papers that study piecewise linear maps on combinatorial manifolds in terms of notions familiar from Morse theory are a paper of Eells and Kuiper [17] and a paper of Kosinski [29]. Both papers appeared in 1962 and define piecewise linear analogues for the notions of regular points and non-degenerate critical points of a certain index. Kosinski states that classical Morse theory results like the Morse inequalities or the Reeb theorem can be derived based on the suggested notions, but detailed proofs are not provided. Eells and Kuiper prove that any combinatorial manifold has a map linear on cells of its barycentric subdivision such that all critical points of the map are non-degenerate. It is mentioned as well that the Morse inequalities hold for maps without degenerate critical points. However, in their approach, which is based on deformations, the study of level sets plays no significant role. An approach based on deformations is also used for an analogue of Morse theory in the category of topological manifolds in a paper of Morse [38], which is referred to by Kosinski and by Eells and Kuiper. It defines analogues for regular and non-degenerate critical points for continuous maps on arbitrary topological manifolds and shows amongst other results that the Morse inequalities hold for continuous maps without degenerate critical points as well. Some of the state-

ments in [38] deal with level sets, but direct analogues to the theorems typically stated in textbooks on Morse theory, namely the diffeomorphism between level sets when no critical levels are passed [36, Theorem 3.1, p. 12] and the homotopy equivalence to the attachment of a λ -cell when passing a single critical point of index λ [36, Theorem 3.2, p. 14], are not obtained.

The definitions for regular and non-degenerate critical points of piecewise linear maps in [17] and [29] are both based on the idea of using piecewise linear analogues of the properties stated by the Morse Lemma [36, Lemma 2.2, p. 6] as characterising properties for the different kinds of points. The details of the definitions however differ slightly. Since the characterisation of regular and critical points is one of the motivating problems of this thesis, several more or less different characterisations are eventually shown to be equivalent, including the definitions of Eells and Kuiper [17] and the definitions of Kosinski [29]. Particularly relevant for this thesis is a notion of equivalence introduced by Kosinski for the definitions of regular and critical points, because several characterisations of such points considered here are obtained from different characterisations of a similarly defined notion of equivalence termed local equivalence. We verify that our notion of local equivalence agrees indeed with Kosinski's notion of equivalence. The objects that are classified by these equivalence relations are triples consisting of a domain, a piecewise linear map defined on that domain, and a point in that domain. Among the various characterisations we derive for the equivalence relations, the following criterion is particularly convenient for the comparison with other notions of regularity and criticality outlined in the next paragraph: Two points with given maps on their domains are considered as equivalent if their links are piecewise linearly homeomorphic in such a way that the homeomorphism respects the separation in an upper and a lower part by the respective level sets through the points.

The approach taken by Eells and Kuiper and by Kosinski, which is also chosen here, is probably the most obvious way to transfer the notions of regular and non-degenerate critical points, and consequently the notion of a Morse function, from the smooth to the piecewise linear category. But later on, different notions of regularity and criticality for piecewise linear maps on combinatorial manifolds have been studied. For two such approaches a brief comparison with the standard one is included in the thesis. The first approach is due to Banchoff [5, 4], the second one to Brehm and Kühnel [9, 30]. It turns out that not only the standard notions of regularity and criticality, but also these two notions can be described in terms of a classification of the points with respect to a suitable equivalence relation, where the three relations have different degrees of coarseness. The finest equivalence relations are Kosinski's relation and the local equivalence relation mentioned above, which define the standard notions of regularity and criticality and were characterised by homeomorphic links with matching upper and lower parts. The equivalence that induces the approach of Brehm and Kühnel is coarser and requires only that the lower parts of the links of the two points have the same homology. Even coarser is the equivalence generating Banchoff's classification. Here it suffices that the lower parts of the links have the same Euler characteristic. The comparison of the different notions of regularity and criticality is presented in Subsection 3.2.5.

A more recent and substantially different approach to develop Morse theory for combinatorial manifolds received much attention and has found a lot of applications in subsequent research. It was suggested by Forman [18, 19] and is

called discrete Morse theory. Instead of piecewise linear maps on the domain of a complex, it considers maps that assign values to the cells of the complex. If the map fulfils certain properties, it is called a discrete Morse function. Alternatively, the theory can be described in terms of certain partial matchings in the Hasse diagram of the face poset of the complex, so that the notion of a Morse function can be avoided completely in this variant of discrete Morse theory. The complex is not required to be a combinatorial manifold; the theory is applicable even to arbitrary finite regular CW-complexes. In spite of the differences between discrete Morse functions and piecewise linear Morse functions in the sense used here, we show in Lemma 3.26 that any discrete Morse function on a combinatorial manifold induces by a simple construction a piecewise linear Morse function such that each critical point of the piecewise linear Morse function corresponds to a critical cell of the discrete Morse function or its restriction to the boundary. This result is the main reason why the thesis includes a suggestion for an extended definition of regularity and non-degenerate criticality for piecewise linear maps that also treats boundary points, although we do not attempt to develop a piecewise linear Morse theory for manifolds with boundary.

Main Results. The most important results of the thesis are related to the characterisation of regular points. The definition suggested here and previously in [17] and [29] can be stated in a nutshell as follows: A point p is regular with respect to a piecewise linear map f if there is a piecewise linear local coordinate system for a neighbourhood of p such that expressing f in terms of these coordinates yields a linear map in some neighbourhood in Euclidean space. We express this requirement in Definition 3.12 using the notion of local equivalence introduced and studied in Section 3.1. In this study we identify several equivalent characterisations for the notion of local equivalence. A list of these characterisations is given in Theorem 3.11. Each of them induces a corresponding characterisation for regular points. These characterisations are listed at the beginning of Section 4.1.

Two of these characterisations induced by different characterisations of local equivalence are particularly convenient. For stating them, it is necessary to observe that the link $\text{lk}(p)$ of a point p is separated into an upper level link $\text{lk}(p)_{\geq f(p)}$ consisting of those points in the link with f -value at least $f(p)$ and a symmetrically defined lower level link $\text{lk}(p)_{\leq f(p)}$. The first characterisation requires that $\text{lk}(p)$ is piecewise linearly homeomorphic to a standard sphere in such a way that the separation into upper and lower level link corresponds to the separation of the standard sphere into two half-spheres by an equatorial plane. The second rephrases the condition by using combinatorially equivalent subdivisions instead of piecewise linear homeomorphisms: A point is regular if and only if some subdivision of $\text{lk}(p)$ is combinatorially equivalent to some subdivision of the standard sphere such that vertices in the upper level link correspond to vertices in the upper half-sphere and vertices in the lower level link correspond to vertices in the lower half-sphere. The underlying facts for the local equivalence relation that imply these characterisations are stated and proven in Lemmata 3.7, 3.8, 3.9 and Corollary 3.10.

Based on these observations, it is not hard to verify that the characterisations for low dimensions observed in [13] match with the definition used here. More-

over, the characterisation in terms of the separation in upper and lower level link can be rephrased as in Lemma 4.1 by requiring that the level set $\text{lk}(p)_{=f(p)}$ is an unknotted co-dimension 1 sphere in the sphere $\text{lk}(p)$. On the other hand, the original definition of regular points in terms of a local coordinate system can easily be transformed into the characterisation stated in Theorem 4.10, which requires the existence of a piecewise linear level-preserving isotopy in some neighbourhood of the point.

Under a suitable general position assumption on the map f , criticality can only occur at vertices. Theorem 4.19 characterises a regular vertex v of a combinatorial manifold under this assumption by the existence of a global piecewise linear level-preserving isotopy in a sufficiently small interval $[f(v) - \epsilon, f(v) + \epsilon]$. Still assuming general position for the map f defined on the combinatorial manifold M , we also prove in Theorem 4.20 that for intervals $[a, b]$ without critical values a piecewise linear level-preserving isotopy $\phi: M_{=b} \times [a, b] \rightarrow M$ exists. This generalises the result in [13] for the 3-dimensional case to arbitrary dimensions and strengthens it in the sense that the isotopy is piecewise linear on its whole domain $M_{=b} \times [a, b]$. Furthermore, this statement is a counterpart of a corresponding statement of smooth Morse theory (see Fact 1.3 in the introductory summary of smooth Morse theory presented in Section 1.1).

For the algorithmic problem of recognising regular vertices raised in [13], we obtain by a reduction from the sphere recognition problem the result stated in Corollary 6.6 that it is undecidable whether a certain vertex of a given simplicial complex of dimension $d \geq 6$ with a map linear on cells is regular or not. However, since the reduction produces from negative instances for the sphere recognition problem instances for the regular vertex recognition problem, where the simplicial complex is not a manifold, the question remains open whether the problem remains undecidable when restricted to combinatorial manifolds.

From the existence of a level-preserving isotopy for intervals without critical values, we can derive in Theorem 5.1 a piecewise linear analogue of the fundamental result from smooth Morse theory on the behaviour of level sets stated in Fact 1.2: If no critical point of a combinatorial manifold M has a value in the interval $[a, b]$, then the level sets $M_{\leq a}$ and $M_{\leq b}$ consisting of all the points with value at most a or b respectively are piecewise linearly homeomorphic and the latter deformation retracts to the former. For the other fundamental result on level sets, Fact 1.4, which describes the change in homotopy type of the level sets when passing a critical point as the attachment of a cell, the corresponding statement for the piecewise linear category is shown in Theorem 5.2.

From these two fundamental results, the equivalence of the homotopy type of the manifold M with a CW-complex (Theorem 5.3), the Morse inequalities (Theorem 5.4), and the Reeb theorem (Theorem 5.5) follow in much the same way as their smooth counterparts (Facts 1.5, 1.6, and 1.7). At least the Morse inequalities and the Reeb theorem for piecewise linear maps have been observed previously in the literature [9, 17, 29, 30].

Another interesting result presented in Lemma 3.26 shows that any discrete Morse function g on a combinatorial manifold M can be transformed into a piecewise linear Morse function f such that critical points of f and critical cells of g correspond to each other and the pieces on which f is linear are the cells of a derived subdivision of M . An analogous result for Banchoff's notion of regular and critical points was found previously by Bloch [8]. Moreover, the result can be regarded as a generalisation of the existence proof for piecewise linear Morse

functions given by Eells and Kuiper [17], which appeared long before the notion of a discrete Morse function was introduced: Applying the transformation to the trivial example of a discrete Morse function that assigns to each cell its dimension yields exactly the piecewise linear Morse function used in that proof.

Outline of the Thesis. The first chapter serves as an introduction that reviews the basic facts relevant for our study and fixes some notation used throughout the thesis. It is divided into three sections. The first section briefly reviews the results from classical smooth Morse theory for which we develop piecewise linear analogues. The second section covers various topics from piecewise linear topology providing the basic tools and results for our work. The last section summarises the required definitions and results from discrete Morse theory.

The level sets with respect to a map linear on cells of a simplicial complex are domains of polytopal complexes whose structure is induced in a natural way from the original simplicial complex. The complexes obtained in that way are called level set complexes and their properties are studied in Chapter 2. The main concern is to describe the rules that govern the influence of the combinatorial structure of the original complex on the combinatorial structure of the level set complex. The results serve as a toolbox for the proofs later on, which is applied whenever simple observations regarding level sets and level set complexes are needed.

Chapter 3 defines the kind of maps that we consider as piecewise linear Morse functions. For that aim, several characterisations for a notion of local equivalence of maps at points are shown to be equivalent, and regular and non-degenerate critical points are defined based on that notion. The obtained notions are compared to previously suggested definitions by Eells and Kuiper, by Kosinski, by Banchoff, and by Brehm and Kühnel. We also show that any discrete Morse function on a combinatorial manifold can be transformed into a piecewise linear Morse function linear on cells of the derived of the combinatorial manifold such that each critical point of the piecewise linear map is induced by a corresponding critical cell of the discrete Morse function.

Additional properties and characterisations of regular points not solely based on the different characterisations of local equivalence are studied in Chapter 4. A short first part deals with a characterisation of regular points in terms of unknotted sphere or ball pairs, which relates the problem of characterising regular points to the Schoenflies conjecture. The longer second part is concerned with level-preserving piecewise linear isotopies.

In Chapter 5, we state and prove, applying the results of the previous chapters, piecewise linear analogues of classical Morse theory statements. The selection of the analogues presented is based on Sections 3 to 5 of Milnor's textbook on Morse theory [36].

The last chapter considers two aspects of piecewise linear Morse theory from a computational point of view. The first section studies the problem of recognising regular points algorithmically. The problem is connected to the problem of sphere recognition, which yields recognition algorithms in low dimensions, but undecidability results in higher dimensions, when arbitrary simplicial complexes are allowed as input. The second section derives an upper bound for the number of cells in combinatorially equivalent complexes that represent a level-preserving piecewise linear isotopy. For that aim, the construction outlined in Chapter 4

for obtaining such an isotopy is analysed with regard to the number of cells produced by the various subdivision steps performed during the construction.

Acknowledgements. The work on this thesis within an inspiring research environment was an outstanding experience for me. I am very grateful to all people that helped creating that friendly environment and for the support I received from them. First of all I want to thank my advisor Günter Rote for introducing me to the topic of the thesis and giving me the opportunity to work on it under his mentoring. I also owe a lot of thanks to Bruno Benedetti for insightful discussions and valuable suggestions. I am grateful to Professor Wolfgang Kühnel for agreeing to review this thesis. Moreover, I would particularly like to thank all colleagues from the work group Theoretical Computer Science at Freie Universität Berlin for the pleasant atmosphere I enjoyed during my studies. I also thank all other people that supported me in my research, gave me advice, or encouraged me.

Chapter 1

Introduction

The results of classical smooth Morse theory provide methods to infer global topological information on a smooth manifold from examining how some suitable smooth function defined on the manifold behaves locally at certain so-called critical points. These points are usually defined and classified into different types according to the derivatives of the function at the respective point. But there is an equivalent characterisation of such points that can be transferred from the the context of differential topology and smooth manifolds to analogous characterisations in other contexts, namely topological manifolds in general topology or combinatorial manifolds in piecewise linear topology. This work addresses implications of such a definition in the latter case of piecewise linear topology. We study the relation to other characterisations of critical points in this setting and how results similar to those of classical smooth Morse theory can be proven.

For establishing a piecewise linear analogue of Morse theory, the considered objects are combinatorial manifolds without boundary, represented by a finite simplicial complex, with real-valued piecewise linear functions defined on their domains. Furthermore we assume that the functions are in general position in the following sense: Any simplicial complex endowed with a piecewise linear function on its domain can be subdivided in such a way that each restriction of the function to some simplex of the subdivision is a linear map. Therefore we can assume without loss of generality that the function is linear on the cells of the given complex. Under this assumption, the function is uniquely determined by its values at the vertices of the simplicial complex. We say that the function is in general position if the values at the vertices are all distinct.

For a lot of basic results and technical lemmata, the assumptions can be relaxed to various degrees without substantial changes to the proofs. So many results will be stated under less restrictive assumptions, with the following generalisations being the most frequently occurring ones in this work: The general position assumption can be dropped in some situations; in particular for results regarding the combinatorial structure of level sets, combinatorial manifolds can be replaced by arbitrary simplicial complexes, or complexes whose cells may be arbitrary convex polytopes, sometimes even unbounded polyhedral sets; occasionally we consider combinatorial manifolds with boundary.

The introduction gives a review of basic notions and facts which serve as tools throughout this work. It starts with a short survey on smooth Morse theory, continues with selected topics from piecewise linear topology that are

relevant for this work, and concludes with a short summary of discrete Morse theory.

1.1 Smooth Morse Theory

This section provides a collection of some basic results from smooth Morse theory, which serve as reference for their piecewise linear counterparts discussed in the main part. Detailed presentations of Morse theory are offered for example by the textbooks of Milnor [36] and Matsumoto [33]; Rote and Vegter give a short survey of basic results in [43, Section 7.4]. Readers not familiar with the notions used in the following may consult these references.

Definitions. Morse theory offers methods to obtain topological information on a smooth manifold M by studying some smooth real-valued function f defined on the manifold. Crucial for this study are the points of M where the derivative of f vanishes, called the *critical points* of f . A critical point p of f is *non-degenerate* if the Hessian of f at p is non-degenerate; in this case the *index* of p is the number of negative eigenvalues of the Hessian. A smooth function $f: M \rightarrow \mathbb{R}$ is called a *Morse function*, if all its critical points are non-degenerate. The terms *regular point* and *ordinary point* are used synonymously to denote a non-critical point. A real number h is called a *critical value* of f if its preimage under f contains a critical point, otherwise it is a *regular value* of f .

An alternative characterisation of critical points is suggested by the Morse lemma [33, Theorem 2.16, p. 44], [36, Lemma 2.2, p. 6]. It states that near a non-degenerate critical point p , using a suitable coordinate system centred at p , the function f can be represented by a shifted normalised non-degenerate quadratic form. For regular points, a similar local representation of f by a shifted non-constant linear form exists. Eells and Kuiper have included this additional observation regarding regular points in their statement of the Morse lemma, labelled as “Proposition” in [17, p. 8]. One easily checks that the existence of a local representation of f near p by one of these forms is also sufficient for the point p being regular or non-degenerate critical. The following summary restates the Morse lemma with the mentioned supplements.

Fact 1.1. Let M be a smooth d -manifold with a smooth function $f: M \rightarrow \mathbb{R}$ defined on it. Then a point p in M is a regular point if and only if there is a p -centred local coordinate system (X_1, \dots, X_d) such that f has the following representation in terms of these coordinates:

$$f = f(p) + X_1$$

A point p in M is a non-degenerate critical point of index λ if and only if there is a p -centred local coordinate system (X_1, \dots, X_d) such that the representation of f in terms of these coordinates is as follows:

$$f = f(p) - X_1^2 - \dots - X_\lambda^2 + X_{\lambda+1}^2 + \dots + X_d^2$$

An important tool in Morse theory is the study of level sets. In smooth Morse theory, the type of level sets typically considered is the set of points in M whose value under some Morse function f is at most a given value h in \mathbb{R} , which we denote by $M_{\leq h}$. The essential results of Morse theory relate the topological

changes of level sets, at least up to homotopy equivalence, when going from a level set $M_{\leq a}$ to a level set $M_{\leq b}$, to the critical points in the preimage under f of the interval $[a, b]$. For compact manifolds M , this yields a description of topological changes when going from the empty set, equal to $M_{\leq a}$ for some a less than the global minimum value of f on M , to the whole manifold M , equal to $M_{\leq b}$ for some b greater than the global maximum value of f on M . For this purpose, the whole interval, say $[\min_{p \in M} f(p) - \epsilon, \max_{p \in M} f(p) + \epsilon]$, is subdivided into subintervals such that each interval either contains no critical value or at most one critical value in its interior. We can assume without loss of generality that intervals of the latter kind are of the form $[h - \epsilon, h + \epsilon]$ for some critical value h and some real number $\epsilon > 0$. This way, taking the topological changes necessary for going from the start of some subinterval to the start of the next, we obtain a sequence of topological changes for going from the empty set to M .

Before we state the results, we briefly recall the notion of CW-complexes used in one of the statements. More detailed summaries are given for example by Hatcher [25, pp. 5–8, 519–529] and Forman [18, Section 1]. CW-complexes are constructed starting from the empty set by iteratively attaching cells to the previously obtained complex. For CW-complexes, the cells are topological balls; hence the boundary of a d -cell is a $(d-1)$ -dimensional topological sphere. A space Y is said to be obtained from X by the attachment of a d -cell, if there is a continuous map g , called the attaching map, from the boundary S^{d-1} of a d -ball B^d to the space X such that Y is the quotient space obtained from the disjoint union of X and B^d by identifying s and $g(s)$ for each $s \in S^{d-1}$. The construction of the quotient space is colloquially described as gluing B^d at its boundary to X along g . A finite CW-complex is any element of a finite sequence of spaces such that the first space of the sequence is the empty space and each subsequent space is obtained from its predecessor by the attachment of a cell. Hatcher [25] requires that the attachments occur in order of increasing dimension; Forman [18, Theorem 1.5] notes that any finite CW-complex is homotopy equivalent to such a one fulfilling this requirement.

Results. The first observation is that for intervals without critical values, the topology of the level sets remains unchanged.

Fact 1.2 ([33, Theorem 3.1, p. 73], [36, Theorem 3.1, p. 12]). Let M be a smooth manifold with a real-valued smooth function f defined on it. Assume that the preimage of the interval $[a, b]$ under f is compact and contains no critical points. Then the level set $M_{\leq a} = \{p \in M \mid f(p) \leq a\}$ is diffeomorphic to the level set $M_{\leq b} = \{p \in M \mid f(p) \leq b\}$. Furthermore, the latter set $M_{\leq b}$ deformation retracts to the former set $M_{\leq a}$.

Another result illustrating the absence of topological changes in intervals without critical values is the existence of an isotopy between the level sets in the interval. Matsumoto [33, Theorem 2.31, p. 66] expresses this isotopy as a diffeomorphism between the preimage of the interval, $f^{-1}([a, b])$, and the product $f^{-1}(a) \times [0, 1]$ of the preimage of a single value with an interval. The proof provided there implies that the constructed diffeomorphism has an additional property, which has a particularly neat formulation when the interval $[0, 1]$ is reparametrised linearly to $[a, b]$: The obtained diffeomorphism $\Phi: f^{-1}(a) \times [a, b] \rightarrow f^{-1}([a, b])$ is level preserving in the sense that for all lev-

els h in $[a, b]$ and points p in $f^{-1}(a)$ the equality $f(\Phi(p, h)) = h$ holds. This yields the following slightly strengthened version of the theorem in Matsumoto's textbook.

Fact 1.3. Let M be a smooth manifold with a smooth function $f: M \rightarrow \mathbb{R}$ such that $f^{-1}([a, b])$ contains no critical points. Then there is a level preserving isotopy from $f^{-1}(a)$ to $f^{-1}(b)$, in other words a diffeomorphism

$$\Phi: f^{-1}(a) \times [a, b] \rightarrow f^{-1}([a, b])$$

such that the equation $f(\Phi(p, h)) = h$ holds for all h in $[a, b]$ and all p in $f^{-1}(a)$.

It remains to describe the topological changes in the level sets, when the interval $[a, b]$ contains a critical value. We restrict our attention to the case that the preimage $f^{-1}(h)$ of the critical value h contains only a single critical point. For the case of several critical points with the same value, see the generalised result stated in [36, Remark 3.3, p. 19]. Alternatively, this case can be avoided by slightly modifying the considered function: For any Morse function, there is an arbitrarily close Morse function with the same critical points but distinct values at each of them [33, Theorem 2.34, p. 69]. The topology change when crossing a critical level of a single non-degenerate critical point is equivalent to the attachment of a cell:

Fact 1.4 ([36, Theorem 3.2, p. 14]). Let p be a non-degenerate critical point of index λ of a Morse function f on a smooth manifold M with value $f(p) = h$. Assume that for some $\epsilon > 0$ the preimage $f^{-1}([h - \epsilon, h + \epsilon])$ is compact and p is the only critical point contained in it. Then $M_{\leq h + \epsilon}$ is homotopy equivalent to $M_{\leq h - \epsilon}$ with a λ -cell attached.

It is possible to achieve a diffeomorphism instead of a homotopy equivalence by thickening the attached cell to a suitably attached handle and smoothing out the obtained space [33, Theorem 3.2, p. 77].

Applying the cell attachments for each critical point, starting with the empty set, we obtain a CW-complex that is homotopy equivalent to the manifold M .

Fact 1.5 ([36, Theorem 3.5, p. 20]). Let f be a Morse function on a smooth manifold M such that $M_{\leq h}$ is compact for every real number h . Then M is homotopy equivalent to a CW-complex whose cells are in bijection with the critical points of f : one λ -cell for each critical point of index λ .

Using handle attachment instead of cell attachment, we obtain a handle decomposition of M [33, Theorem 3.4, p. 81].

We conclude with two implications, the Morse inequalities, which relate Betti numbers and indices of critical points, and the Reeb theorem, stating that a compact manifold with exactly two critical points is homeomorphic to a sphere.

Fact 1.6 (Morse inequalities [36, Theorem 5.2, p. 29; Inequality (4 $_{\lambda}$), p. 30]). Let M be a compact smooth manifold with a Morse function f . Denote by c_{λ} the number of critical points of index λ of f , and by β_{λ} the λ -th Betti number of M . Then the following inequalities and equalities hold:

1. The Weak Morse Inequalities. For each λ :

$$\beta_{\lambda} \leq c_{\lambda}$$

2. An equality representing the Euler characteristic $\chi(M)$ of M in two ways:

$$\chi(M) = \sum_{\lambda} (-1)^{\lambda} c_{\lambda} = \sum_{\lambda} (-1)^{\lambda} \beta_{\lambda}$$

3. The Strong Morse Inequalities. For each λ :

$$\sum_{i=0}^{\lambda} (-1)^i \beta_{\lambda-i} \leq \sum_{i=0}^{\lambda} (-1)^i c_{\lambda-i}$$

Fact 1.7 (Reeb theorem [36, Theorem 4.1, p. 25]). If a smooth manifold has a Morse function with exactly two critical points, then it is homeomorphic to a sphere.

1.2 Piecewise Linear Topology

This section serves as a survey on definitions and results from piecewise linear topology that are the basic tools for our study. The treatment is restricted to finite complexes embedded geometrically in some Euclidean space. The combinatorial structure of the complexes given by the partial ordering induced by the face relation receives special attention. The results are mainly taken from the textbooks of Rourke and Sanderson [44], Glaser [21], and Hudson [26]. For basic properties of convex polyhedral sets, we also refer to Grünbaum [23].

1.2.1 Polyhedral Sets and Polyhedral Complexes

The combinatorial manifolds that we want to investigate by means of piecewise linear Morse theory are given as simplicial complexes. As in the case of smooth Morse theory, an important tool for this investigation is the study of level sets. A level set is the preimage under a given function of a certain level range. We are primarily interested in level sets that are preimages of closed intervals. In Chapter 2, we show that a level set of this kind can be regarded as the domain of a complex, whose combinatorial structure is naturally induced by that of the considered manifold. However, the cells of this naturally induced level set complex are not necessarily simplices; in general, we obtain a complex whose cells are convex polytopes. It turns out that the basic results on level set complexes in Chapter 2 have more concise formulations when we generalise the considered domains and level ranges to a class of complexes that contains as subclasses simplicial complexes, complexes whose cells are convex polytopes, and complexes formed by closed intervals. In this subsection, we introduce this class of complexes, which we call polyhedral complexes. Out of habit, we stick to the term “cell” for the members of the complexes, even if the members sometimes do not fulfil all conditions that are commonly required from cells. We also mention the notion of maps linear on cells, including cellular and simplicial maps as special cases.

Polyhedral Sets. For our polyhedral complexes, the cells are taken from the class of polyhedral sets. We recall the definition of polyhedral sets, following the terminology used by Grünbaum [23, Section 2.6, p. 26].

Definition 1.8 (polyhedral set). A subset of the Euclidean space \mathbb{R}^d is called a *polyhedral set*, if it is the intersection of finitely many closed half-spaces of \mathbb{R}^d . A subset F of a polyhedral set P is a *face* of P , denoted by $F \sqsubseteq P$, if it is either empty, or the whole set P , or the intersection of P with a supporting hyperplane of P [23, Section 2.4, p. 17]. When $(H_i^+)_{i \in I}$ is a family of half-spaces whose intersection is the polyhedral set P and $(H_i)_{i \in I}$ is the corresponding family of hyperplanes such that H_i bounds H_i^+ for each i , then for any subset $J \subseteq I$ and the corresponding subfamily of hyperplanes $(H_i)_{i \in J}$, the intersection $P \cap \bigcap_{i \in J} H_i$ of P with the hyperplanes is a face of P , and any non-empty face of P is representable as such an intersection [23, Section 2.6, pp. 26–27]. For a polyhedral set P , we denote by $\mathcal{F}(P)$ the set of faces of P . A *vertex* of P is a point v in P such that the singleton set $\{v\}$ is a face of P . The set of vertices of P is denoted by $\text{Vrt}(P)$. The union of the proper faces of P is the *boundary* of P , the remaining points form the *interior* of P . These definitions agree with the topological notions of boundary and interior, when P is considered as a subset of its affine hull.

The term “convex polyhedron” seems to be more commonly used for a polyhedral set, but switching to Grünbaum’s notion avoids confusion with the more general concept of a polyhedron defined in the book on piecewise linear topology by Rourke and Sanderson, our main reference for facts from this field [44, p. 2].

A *convex polytope* can be characterised as a bounded polyhedral set [23, Theorem 3.1.3, p. 32], or equivalently as the convex hull of a finite set of points [23, p. 31]. It is the convex hull of its vertices, and each face is the convex hull of some subset of vertices. Note that the polyhedral sets contained in \mathbb{R}^1 are exactly the closed intervals.

Polyhedral Complexes. Our definition of polyhedral complexes agrees with the one given by De Loera, Rambau, and Santos in [16, pp. 45–46]: We adopt the closure properties from the usual definition of simplicial complexes, see for example [43, p. 280], and extend the allowed class of cells from simplices to polyhedral sets, as suggested in [23, p. 41]. Therefore, the simplicial complexes form a subclass of the polyhedral complexes. Allowing only convex polytopes as cells yields another important subclass. In addition to these well known subclasses, we introduce the notion of interval complexes: Since Morse theory deals with real-valued functions, complexes embedded in \mathbb{R} and hence formed by closed intervals occur naturally in our study. The notions of subcomplexes and subdivisions, as defined for example in [44, p. 15], extend to polyhedral complexes in the obvious way.

Definition 1.9 (polyhedral complex). A *polyhedral complex* is a finite family of polyhedral sets in some \mathbb{R}^n such that the following closure properties hold:

1. Each face of a family member is again a family member.
2. The intersection of two family members is a face of both members, possibly the empty set, and thus a family member by the first property.

The family members are called the *cells* of the complex. A *polytopal complex* is a polyhedral complex such that each cell is a convex polytope. A *simplicial complex* is a polyhedral complex whose cells are simplices; in particular, it is a

polytopal complex. An *interval complex* is a polyhedral complex embedded in the one-dimensional euclidean space \mathbb{R} . In other words, the cells of an interval complex are closed intervals which intersect, if at all, only at their endpoints.

The *domain* or *underlying set* $|K|$ of a polyhedral complex K is the union of its cells. For each point p in the domain $|K|$, there is a unique cell of K containing p in its interior, namely the intersection of all cells in K containing p . In other words, $|K|$ is the disjoint union of the interiors of cells in K . A *subcomplex* of a polyhedral complex is a subfamily that is a polyhedral complex, and a *subdivision* of a polyhedral complex L is a polyhedral complex K such that K and L have the same domain and each member of K is contained in a member of L .

The face relation \sqsubseteq among the cells of a polyhedral complex induces a partial ordering on the complex. The notion of vertices extends from cells to complexes: A *vertex* of K is a vertex of one of its members, or equivalently, a point v in $|K|$ such that $\{v\}$ is a cell of K ; the set of vertices of K is denoted by $\text{Vrt}(K)$. Sometimes the cell $\{v\} \in K$ is also called a vertex. In this sense a vertex of K corresponds to a 0-dimensional cell of K .

Simplicial and polytopal complexes are standard notions in combinatorial topology; the latter kind however appears under various names in the literature. Polytopal complexes are called “convex linear cell complexes” by Glaser [21, p. 9] and Hudson [26, p. 5], or simply “cell complexes” by Rourke and Sanderson [44, p. 14]. Grünbaum [23, p. 39] mentions the terms “convex complex” and “geometric cell complex”, and also the term “polyhedral complex”, which we adopted for the more general class. The term “polytopal complex” chosen here is used amongst others by Adiprasito and Benedetti [1, 2]. Our terminology of polyhedral and polytopal complexes intends to reflect the popular convention that polyhedra are allowed to be unbounded polyhedral sets whereas polytopes are restricted to be bounded.

As noted in the definition, the face relation defines a poset structure on the complex. Since each cell of a complex is a face of some cell, at least of itself, the terms face and cell are often used synonymously. The term cell is preferred when considering the polyhedral set as a member of a complex, the term face emphasises its embedding in the poset structure. If some cell F is a face of some cell S , we also call S a *co-face* of F .

For polytopal and simplicial complexes, the standard basic example is the complex formed by the set of faces of a single convex polytope or simplex. These examples generalise to the basic example of a polyhedral complex, namely the complex associated with a polyhedral set P . The set of cells of this complex coincides with the set $\mathcal{F}(P)$ of faces of P . If the set P is a convex polytope, a simplex, or a closed interval, then the associated complex is polytopal, simplicial, or an interval complex respectively. The poset structure on $\mathcal{F}(P)$ induced by the face relation is often referred to as the face lattice of P .

The proper faces $\mathcal{F}(P) \setminus \{P\}$ of a polyhedral set P form a subcomplex of $\mathcal{F}(P)$. The domain of this subcomplex is the topological boundary of the polyhedral set in its affine hull. It is a common abuse of notation to denote the associated complex $\mathcal{F}(P)$ by P , and the boundary complex $\mathcal{F}(P) \setminus \{P\}$ by $\text{bd}(P)$ [44, p. 14]. The combination with another frequent convention for simplified notation, namely dropping the braces enclosing the element of a singleton set, leads to further ambiguity: A vertex v may denote a point p in space,

the 0-simplex $\{p\}$ containing the point, or the complex $\{\emptyset, \{p\}\}$ associated with the 0-simplex. The exact meaning is usually clear from the context, and the identification of the three objects often seems so natural that the inaccuracy is hardly ever noticed.

Maps Linear on Cells. Our primary object of study are complexes that have a map defined on their domain such that the restriction to each cell is a linear map. By a linear map, we mean any map obtained by restricting an affine linear map on some affine subspace to some domain contained in the subspace. The following definitions concerning such maps are taken from Rourke and Sanderson's textbook [44, pp. 16–17].

Definition 1.10. A map $f: |K| \rightarrow X \subseteq \mathbb{R}^m$ defined on the domain of a polyhedral complex K is called *linear on cells* of K , if for each cell $S \in K$, the restriction of f to S is a linear map. For two polytopal complexes K and L , a map $f: |K| \rightarrow |L|$ is called *cellular*, if f is linear on cells of K and for each cell S of K , the image $f(S)$ is a cell in L . A cellular map $f: |K| \rightarrow |L|$ is called *simplicial*, if K and L are simplicial complexes.

A map that is linear on cells of a polytopal complex K is uniquely determined by its values on vertices. All other values can be inferred by cell-wise linear interpolation. We say that a map $f: |K| \rightarrow \mathbb{R}$ linear on cells of K is in *general position* if the values at the vertices are all distinct.

For any map $f_0: \text{Vrt}(K) \rightarrow X$ defined on the vertices of polytopal complex K , there is, by the previous observation, at most one map f linear on cells of K extending f_0 to $|K|$. The existence of f however can not be ensured for arbitrary polytopal complexes, because f_0 might not respect affine dependencies of vertices in a common cell. But for simplicial complexes K , where all vertices of a common cell are affinely independent, an extension of f_0 always exists: Cell-wise linear interpolation of a map $f_0: \text{Vrt}(K) \rightarrow X$ gives a map $f: |K| \rightarrow \text{conv}(X)$ from $|K|$ to the convex hull of X . If additionally the codomain X is the domain of a simplicial complex L and for each simplex S of K , its vertices are mapped to vertices of a simplex in L , that is $f_0(\text{Vrt}(S)) = \text{Vrt}(T)$ for some simplex $T \in L$, then the extension by cell-wise linear interpolation yields a simplicial map $f: |K| \rightarrow |L|$.

By abuse of notation, the vertical bars representing the domain are sometimes dropped to indicate that f is linear on cells or cellular: $f: K \rightarrow X$ for some $X \subseteq \mathbb{R}^m$ denotes a map $f: |K| \rightarrow X$ that is linear on cells of K and $f: K \rightarrow L$ for a complex L denotes a cellular map $f: |K| \rightarrow |L|$. We will use the notation without vertical bars for functions, if at all, only sparingly, because we occasionally consider functions whose arguments are indeed the cells of the complex and not the points in its domain: The combinatorial equivalences $\phi^{\text{comb}}: K \rightarrow L$ considered in Subsection 1.2.7 bijectively assign cells of one complex to cells of another complex; the functions studied in discrete Morse theory, a theory developed by Forman [18] and shortly summarised in Section 1.3 of this thesis, assign a real number to each cell of a complex. Nevertheless, the dropping of the vertical bars is quite common in many other contexts when no ambiguity arises, for example topological properties of the domain of a complex are often attributed to the complex itself.

Table 1.1: The different kinds of closed intervals and for each interval I the associated interval complex formed by the set of its faces $\mathcal{F}(I)$.

I	$\mathcal{F}(I)$	I	$\mathcal{F}(I)$
$[a, b]$	$\{\emptyset, \{a\}, \{b\}, [a, b]\}$	$\{h\}$	$\{\emptyset, \{h\}\}$
$(-\infty, h]$	$\{\emptyset, \{h\}, (-\infty, h]\}$	$[h, \infty)$	$\{\emptyset, \{h\}, [h, \infty)\}$
\mathbb{R}	$\{\emptyset, \mathbb{R}\}$	\emptyset	$\{\emptyset\}$

1.2.2 Interval Complexes

This subsection contains a few remarks on interval complexes. By our definition, the cells of interval complexes are closed intervals, including unbounded intervals, singletons, and the empty interval. There is only one interval complex containing the interval \mathbb{R} , namely $\{\mathbb{R}, \emptyset\}$. Other interval complexes can contain at most one interval half-bounded at the right end and at most one interval half-bounded at the left end, in total at most two half-bounded intervals.

We are mainly concerned with the simplest case of an interval complex, the complex associated with a single closed interval. The cells of the associated interval complex are the interval itself, the empty cell, and for each endpoint of the interval a cell containing only the endpoint. Table 1.1 lists the details for all kinds of closed intervals. Each cell is a face of itself, the empty cell is a face of every other cell, and an endpoint cell is a face of the whole interval. Let us take a look at an example to see why regarding a closed interval as a complex can be useful for describing level sets.

Consider a triangle $T = abc$ (see Figure 1.1) with a linear function f on T defined by linear interpolation between the values at the vertices. Let us assume that these values and the level value $h \in \mathbb{R}$ we are interested in are in the following order: $f(a) < f(b) < h < f(c)$. Then the preimage $f^{-1}(h)$ of the level h is a line segment that divides the triangle into a quadrilateral $f^{-1}((-\infty, h])$ and a triangle $f^{-1}([h, \infty))$. (From now on, we simplify our notation for preimages of intervals and omit the outer parentheses when the interval is denoted by its endpoints in brackets.) The edges of the quadrilateral are the edge ab , the line segment $f^{-1}(h)$ and two line segments contained in the edges ac and bc , namely the part of the edges between a or b and the intersection point of the respective edge with the line segment $f^{-1}(h)$. These line segments can be written as $ac \cap f^{-1}((-\infty, h])$ and $bc \cap f^{-1}((-\infty, h])$. The vertices of the quadrilateral are a , b , $ac \cap f^{-1}(h)$, and $bc \cap f^{-1}(h)$. Similarly, the triangle $f^{-1}[h, \infty)$ has the edges $f^{-1}(h)$, $ac \cap f^{-1}[h, \infty)$, and $bc \cap f^{-1}[h, \infty)$, and the vertices c , $ac \cap f^{-1}(h)$, and $bc \cap f^{-1}(h)$. Note that the vertex c can be represented by the intersection $c \cap f^{-1}[h, \infty)$, that the line segment $f^{-1}(h)$ coincides with $T \cap f^{-1}(h)$, and that similar representations exist for the vertices a and b and for the edge ab together with the preimage $f^{-1}((-\infty, h])$.

When the letter I denotes one of the closed intervals $(-\infty, h]$ or $[h, \infty)$, and is regarded as a polyhedral set, the faces of I are either I itself, or the singleton set of its endpoint h , or the empty set. For each face of the level set $f^{-1}(I) = T \cap f^{-1}(I)$ we have found a representation as the intersection of




Figure 1.1: A triangle $T = abc$ subdivided at level h . The linear function $f: T \rightarrow \mathbb{R}$ assigns to each point of T its height with respect to the vertical axis. Indicated are the quadrilateral level set $f^{-1}(-\infty, h]$ below level h , the triangular level set $f^{-1}[h, \infty)$ above level h , and their common intersection, the line segment forming the level set $f^{-1}(h)$ exactly at level h .

some face of the triangle T and the preimage of some face of the interval I . Conversely, one easily checks that each such intersection is a face of the level set. Hence, we can sum up our observations by the following statement: Each face of $T \cap f^{-1}(I)$ can be written as $F \cap f^{-1}(Y)$ where F is a face of T and Y is a face of I , moreover each such term $F \cap f^{-1}(Y)$ is a face of $T \cap f^{-1}(I)$.

We will see in Chapter 2 that this statement remains true if T and I are replaced by arbitrary polyhedral sets and can be generalised from single cells to complexes. Since we will apply the results of Chapter 2 later only to finite simplicial complexes M , the consideration of unbounded intervals could be avoided by bounding them by $\max_{x \in |M|} f(x)$ from above and by $\min_{x \in |M|} f(x)$ from below. This would even permit to restrict our study completely to the more common class of polytopal complexes. But it seems more intuitive to describe a level set such as $f^{-1}(-\infty, h]$ as the set of points whose f -value is at most a certain threshold h rather than as the set $f^{-1}[\min_{x \in |M|} f(x), h]$ of points whose f -value lies between the global minimum and the threshold.

1.2.3 Operations on Complexes

We recall some frequently used operations on complexes. The operations have geometric and combinatorial aspects. The notion of joins, including as special cases cones and suspensions, is only introduced for polytopal complexes; products and intersections are also defined for polyhedral complexes with unbounded cells.

Joins. The geometric definition of a join of two subsets A and B in \mathbb{R}^n as the union of all line segments connecting a point in A with a point in B can be applied to any pair of subsets [44, p. 1], [21, Exercise IV.5, p. 137]. But such a broad application has the disadvantage that the combinatorial description of

a join given below fails to hold for arbitrary pairs. Therefore we restrict the definition of a join to independent sets as considered in [44, p. 22]. Equivalent restrictions are used for the definition of a join of simplicial complexes in [21, p. 6] and [26, p. 6].

Definition 1.11 (independent sets [44, p. 22]). Two subsets A and B of \mathbb{R}^n are *independent* if they are disjoint and the interiors of all line segments between a point in A and a point in B are pairwise disjoint. In other words, A and B are independent if for every point that is a convex combination $\lambda a + (1 - \lambda)b$ of a point a in A and a point b in B , the representation as such a convex combination is unique in the sense that $\lambda \in [0, 1]$ is unique, $a \in A$ is unique unless $\lambda = 0$, and $b \in B$ is unique unless $\lambda = 1$.

For the combinatorial description of a join, we use the notion of the product order on a Cartesian product of posets. Let (A, \leq_A) and (B, \leq_B) be two posets. Then a poset structure $(A \times B, \leq_\times)$ on the set-theoretic Cartesian product of A and B is given by the rule that $(a, b) \leq_\times (a', b')$ holds if $a \leq_A a'$ and $b \leq_B b'$. Two posets (A, \leq_A) and (B, \leq_B) are isomorphic if there is a bijection $\phi: A \rightarrow B$ such that $\phi(a) \leq \phi(a')$ holds in B if and only if $a \leq a'$ holds in A .

Definition 1.12 (join). For two independent subsets A and B of \mathbb{R}^n , their *join*, denoted by AB is defined as follows: If both sets are non-empty, the join AB is the union of all line segments connecting a point in A to a point in B ; equivalently, AB is the set of all points expressible as convex combinations $\lambda a + (1 - \lambda)b$ of a point a in A and a point b in B . For joins involving the empty set, we define $\emptyset A = A\emptyset = A$.

The join of a p -dimensional and a q -dimensional convex polytope is $(p+q+1)$ -dimensional convex polytope whose faces are exactly the pairwise joins of faces of the original polytopes: $\mathcal{F}(ST) = \{FG \mid F \in \mathcal{F}(A), G \in \mathcal{F}(B)\}$. Adding the observation that the intersection of joins $ST \cap S'T'$ agrees with the join of the intersections $(S \cap S')(T \cap T')$, we can conclude that the following definition of a join of complexes yields a valid complex:

For polytopal complexes K and L with independent domains $|K|$ and $|L|$ in \mathbb{R}^n , the collection of all pairwise joins $\{ST \mid S \in K, T \in L\}$ is a polytopal complex, called the *join* of K and L and denoted by KL , whose domain is the join of the domains: $|KL| = |K||L|$. The relationship of the dimensions carries over from joins of convex polytopes to joins of polytopal complexes. The vertex set $\text{Vrt}(KL)$ of the join KL is the union of the vertex sets of K and L . If K and L are simplicial, then their join KL is also simplicial.

Combinatorially, that is with regard to the poset structure, the join acts like a product: A cell FG in KL is a face of some other cell ST in KL if and only if F is face of S in K and G is a face of T in L . In other words, the face relation \sqsubseteq_{KL} induces a poset on KL that is isomorphic to the product order of the posets (K, \sqsubseteq_K) and (L, \sqsubseteq_L) , where the poset isomorphism is given by assigning ST in KL to (S, T) in the set-theoretic Cartesian product $K \times L$.

The independence of two sets can be enforced by embedding them into independent subspaces of a higher dimensional space. This construction is called the *external join* in [44, p. 23] and also applies to sets embedded in different spaces: Embed the sets $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^m$ as $A \times \{0\}^m \times \{0\}$ and $\{0\}^n \times B \times \{1\}$ into $\mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}$, identified with \mathbb{R}^{n+m+1} in the usual way. The embedded

sets are independent in \mathbb{R}^{n+m+1} and their join is the external join of A and B . Consequently, the external join of two polytopal complexes is the complex consisting of all external joins of cell pairs. Whenever we consider joins of sets or complexes that are non-independent or embedded in different spaces, we are referring to the external join.

Two important special cases of joins are cones and suspensions. A cone is the join of a set with a single point, a suspension is the join of a set with two distinct points. In both definitions, we tacitly assume that the joined sets are independent or an external join is performed.

Definition 1.13 (cone). A *cone* on a set $B \subseteq \mathbb{R}^n$ is the join of B and a singleton set $\{v\} \subseteq \mathbb{R}^n$. We usually write vB instead of $\{v\}B$. For a cone vB , the point v is called the *apex* of the cone and the set B its *base*. For a complex K and a singleton set $\{v\}$, the complex formed by the join of the complexes $\mathcal{F}(\{v\})$ and K is also called a cone on K . Clearly, the domain of this cone complex on K is the cone on the domain of K with apex v . Again we prefer writing vK in place of $\mathcal{F}(\{v\})K$ or $\{v\}K$, and we call v the apex and K the base of the cone. Occasionally we refer to $\mathcal{F}(\{v\})$ or $\{v\}$ as apex as well, and to $|K|$ as the base.

Definition 1.14 (suspension). A *suspension* of a set $B \subseteq \mathbb{R}^n$ is the join of B and a set $\{v, w\}$ of two distinct points. For a complex K , a suspension of K is a join of K with a complex $\{\emptyset, \{v\}, \{w\}\}$ for distinct vertices v and w . We sometimes call the points v and w the apices of the suspension.

Cartesian Products. The Cartesian product of two polyhedral sets, one in \mathbb{R}^m , the other in \mathbb{R}^n , is a polyhedral set in \mathbb{R}^{n+m} . The faces of the product are the pairwise products of faces: $\mathcal{F}(S \times T) = \{F \times G \mid F \in \mathcal{F}(S), G \in \mathcal{F}(T)\}$. Given two sets S and S' in \mathbb{R}^n and two sets T and T' in \mathbb{R}^m , the intersection of the Cartesian products $S \times T \cap S' \times T'$ is the Cartesian product of the intersections $(S \cap S') \times (T \cap T')$. These observations allow to define a product complex of two polyhedral complexes by collecting all pairwise Cartesian products of cells [44, p. 15].

Since this description resembles the description of a join of complexes, with join replaced by product, let us point out some differences. In the case of a join of two complexes, assigning to each pair (S, T) of cells the cell ST in the join establishes a bijection between pairs of cells and cells of the join. For products however, the assignment $(S, T) \mapsto S \times T$ is not injective, because any Cartesian product with an empty cell as factor yields the empty cell. Fortunately, the empty cell is the only cell representable by products of different pairs. Consequently, the combinatorial structures of the face posets of joins and products differ. For joins, we observed that FG is a face of ST if and only if F is a face of S and G is a face of T . For products, the implication that if F is face of S and G is a face of T , then $F \times G$ is a face of $S \times T$ still holds, but the converse fails when the empty cell is involved. A correct characterisation of the poset structure on the product complex is provided by the following statement: The cell $F \times G$ is a face of $S \times T$ if and only if F is empty, or G is empty, or F is face of S and G is a face of T .

Another difference is the dimension: The join of an n -cell and an m -cell is an $(n+m+1)$ -cell, whereas their product is an $(n+m)$ -cell. In particular, the vertex set of a product complex is the Cartesian product of the vertex sets, the vertex

set of a join is the union of the vertex sets. Furthermore the join of simplicial complexes is simplicial, but the product complex of simplicial complexes is not simplicial in general. Actually, a product complex is only simplicial in very trivial cases, when the domain is empty, or one of the complexes is 0-dimensional and the other simplicial; as soon as both factors contain an edge, the product complex contains a quadrilateral and is not simplicial. At least the product complex of polytopal complexes is a polytopal complex.

Note that the set-theoretic Cartesian product $K \times L$ of two complexes, which we used above to describe the poset structure on the join, and the product complex of K and L , which we also denote by $K \times L$, are different objects: The elements of the former are pairs of cells, whereas the elements of the latter are Cartesian products of cells. Set-theoretic Cartesian products of complexes appear here essentially only in the already mentioned context, namely as isomorphic description of poset structures. Product complexes occur in this work mostly in connection with an isotopy or homotopy, in products such as $K \times [0, 1]$ or $K \times [a, b]$, where the second factor is the complex associated with the interval parametrising the isotopy or homotopy. Therefore the meaning of a term $K \times L$ will usually be clear from context, otherwise we will explicitly state whether we refer to the set-theoretic product or the product complex.

Definition 1.15 (product complex). The *product complex* $K \times L$ of two polyhedral complexes K and L is the complex formed by all Cartesian products $S \times T$ where S is a cell in K and T is a cell in L . The domain $|K \times L|$ of the product complex is the Cartesian product $|K| \times |L|$ of the domains. The face poset structure on $K \times L$ is induced by the structures on K and L by the following characterisation: A cell $F \times G$ is a face of $S \times T$ if and only if F is empty, or G is empty, or F is face of S and G is a face of T .

For polytopal complexes K and L , the face poset on the product complex $K \times L$ is isomorphic to the restriction of the face poset on the (external) join complex KL to the subset obtained by removing the cells $S = S\emptyset$ and $T = \emptyset T$ for non-empty cells $S \in K$ and $T \in L$. Furthermore, in the language of structure preserving maps between posets, the map from the join KL to the product complex $K \times L$ given by the assignment $ST \mapsto S \times T$ is a surjective monotone map with respect to the face posets, that is $FG \sqsubseteq ST$ implies $F \times G \sqsubseteq S \times T$. For non-polytopal complexes K and L , where we did not define the notion of a join, we can use the product order to obtain an analogous result: The face poset on $K \times L$ is isomorphic to the product order of (K, \sqsubseteq_K) and (L, \sqsubseteq_L) restricted to the subset of pairs where the members are either both non-empty or both empty. The assignment $(S, T) \mapsto S \times T$ yields a surjective monotone map from the product order to the face poset of $K \times L$.

Intersections. For the intersection of two polyhedral sets, we can observe similar rules as for joins and products, that allow to define an intersection complex as the set of all pairwise intersections of cells. The intersection of two polyhedral sets in \mathbb{R}^n is again a polyhedral set in \mathbb{R}^n . The faces of the intersection are the pairwise intersections of faces: $\mathcal{F}(S \cap T) = \{F \cap G \mid F \in \mathcal{F}(A), G \in \mathcal{F}(B)\}$. Clearly, we also have $(S \cap T) \cap (S' \cap T') = (S \cap S') \cap (T \cap T')$. Therefore, for two polyhedral complexes K and L embedded in \mathbb{R}^n , the set $\{S \cap T \mid S \in K, T \in L\}$ is a polyhedral complex in \mathbb{R}^n . We denote

this intersection complex by $K \cap L$, although in general it differs from the set-theoretic intersection. On the other hand, the set-theoretic intersection of complexes is usually only considered for subcomplexes of the same complex, where both notions, the set-theoretic intersection and the intersection complex, yield the same result.

Definition 1.16 (intersection complex). The *intersection complex* $K \cap L$ of two polyhedral complexes K and L is the complex formed by all pairwise intersections $S \cap T$ of cells S in K and T in L . The domain $|K \cap L|$ of the intersection complex is the intersection $|K| \cap |L|$ of the domains.

The intersection gives yet another operation on complexes whose cells are given by a pairwise operation on cells. What are the differences to the previous two examples, the join and the product complexes? Let us again consider the map assigning to pairs of cells (S, T) the result of the operation, in this case $S \cap T$. As in the case of products, any intersection involving the empty cell results in the empty cell; but in most cases, there are even more disjoint pairs of cells than just the trivial ones. And the empty cell is sometimes not the only cell that is the intersection of several distinct pairs of cells. For example, the intersection of two non-parallel edges sharing the same point in \mathbb{R}^n as vertex agrees with the intersection of one edge with the vertex and with the intersection of the vertex with itself. Hence, in general, the dimension of the intersection of two cells or complexes can range between -1 and the smaller dimension of the two cells or complexes. In particular, not all vertices of the intersection complex are necessarily obtained by an intersection involving a vertex; higher dimensional cells may also intersect in just a single point. Furthermore the intersection of two simplices is not necessarily a simplex, therefore the intersection complex of simplicial complexes is not necessarily simplicial. But it is sufficient that at least one of the two complexes is polytopal to obtain a polytopal intersection complex.

Situations where a cell in the intersection complex is representable as intersection of different pairs of cells also induce counterexamples to a hypothetical implication that $F \cap G$ being a face of $S \cap T$ implies that F is a face of S and G is a face of T . Take for F and G the non-parallel edges of the previous example and for S and T their shared vertex, then $F \cap G$ is a face of $S \cap T$, but neither F is a face of S nor G is a face of T . The converse implication however still holds for the poset structure of intersection complexes: If F is a face of S and G is a face of T , then $F \cap G$ is a face of $S \cap T$.

This leads to two questions: Which pairs of cells have the same intersections? Are there rules characterising the face poset of the intersection complex in terms of the face posets of the original complexes? The answer to the first question clearly depends on the geometric location of the cells in space. It turns out that the second question is related to the first one: Considering two pairs of cells as equivalent, if their intersections agree, and taking a suitable system of canonical representatives for the equivalence classes, the face poset of the intersection complex is isomorphic to the restriction of the product order of the original complexes to this set of representatives. For further details we refer to the study of level set complexes and their combinatorial structure in Chapter 2, because intersection complexes can be considered as special cases of level set complexes.

1.2.4 Piecewise Linear Maps

In this subsection, we review some characterisations and properties of piecewise linear maps. As reflected in the name, one characterisation of piecewise linear maps calls a map piecewise linear, if the domain can be subdivided into suitable pieces, such that the restriction of the map to each piece is a linear map. In particular, a map linear on cells of a polyhedral complex K is piecewise linear in this sense. Our study concerning the notions of regular and critical points of piecewise linear maps in Chapter 3 profits from having other points of view on these maps at hand. In particular, some notions introduced by Rourke and Sanderson [44] for their characterisation of piecewise linear maps are useful for us. Therefore we summarise the basic results on piecewise linear maps from [44]. Throughout the thesis, we will often abbreviate the words “piecewise linear” by the letters PL.

The different points of view on piecewise linear maps already start with different descriptions of the class of objects from which domains and codomains of the maps are taken. Sometimes the objects are described as domains of simplicial complexes, sometimes as polyhedra in the sense explained below. Our setting of considering a function $f: |M| \rightarrow \mathbb{R}$ linear on cells of a complex M that forms combinatorial manifold is already a bit of a mixture of both points of view, because the domain of f is represented as the domain of a simplicial complex, but the codomain \mathbb{R} is not viewed as a complex.

Polyhedra, Cone Neighbourhoods, and Links. Rourke and Sanderson define a *polyhedron* as a subset P of an Euclidean space fulfilling the following local property [44, p. 2]: Each point $p \in P$ has a neighbourhood N in P such that N is the cone with apex p and base L for some compact set L . Such a neighbourhood N is called a *cone neighbourhood* or *star* of p in P , and L is called a *link*. The cone neighbourhood N can always be assumed to be an ϵ -neighbourhood, that is, N is a ball in P of radius ϵ centred at p and L is the corresponding sphere for some sufficiently small ϵ , where distances in P are given by the ∞ -norm of the ambient space [44, p. 3]. Apparently, the argument does not rely on the ∞ -norm being used as norm; but the piecewise linear setting suggests using norms whose ϵ -balls are convex polytopes, with ∞ -norm and 1-norm being the standard examples. Using the facts that Euclidean spaces and half-spaces are polyhedra [44, Examples 1.3 (5) and 1.3 (8), p. 4], and that a finite intersection of polyhedra is a polyhedron [44, Example 1.3 (3), p. 4], we can conclude that polyhedral sets are polyhedra.

Another result states that a locally finite union of a family of compact polyhedra is a polyhedron, where locally finite means that any point has a neighbourhood intersecting only finitely many polyhedra from the family [44, Example 1.3 (5), p. 4]. Conversely, any polyhedron has a representation as a locally finite union of a family of simplices, and for compact polyhedra, a finite family of simplices suffices [44, Theorem 2.2, p. 12]. In particular, convex polytopes can be represented as finite unions of simplices. This implies that polyhedra can be equivalently characterised as locally finite unions of simplices, or as locally finite unions of convex polytopes. Furthermore, this equivalence restricts to an equivalence of compact polyhedra and (globally) finite unions of simplices or convex polytopes: Any finite union of simplices or convex polytopes is a compact polyhedron, because a finite union of compact sets is compact, and any

compact polyhedron is a finite union of simplices, in particular a finite union of convex polytopes.

So far, the family of simplices or convex polytopes does not necessarily form a complex. For the case of compact polyhedra, Rourke and Sanderson show that this additional requirement can be met [44, Theorem 2.11, p. 16]: Any compact polyhedron is the domain of some simplicial complex. On the other hand, the domain of any simplicial or polytopal complex is a compact polyhedron since it is a finite union of compact convex polytopes. Summing up, compact polyhedra, domains of simplicial complexes, and domains of polytopal complexes describe the same class of subsets of Euclidean spaces.

Characterisations of Piecewise Linear Maps. A continuous map $f: P \rightarrow Q$ for two polyhedra P and Q is *piecewise linear*, if each point p in P has a cone neighbourhood $N = pL$ such that the function values of f are obtained by linearly interpolating between the apex p and the base L : For any point q in N with its representation as convex combination $q = \lambda p + (1 - \lambda)r$ for some $r \in L$, the equality $f(q) = \lambda f(p) + (1 - \lambda)f(r)$ holds [44, p. 5]. This characterising property is referred to as f being *locally conical*. If P is represented as a locally finite union of compact polyhedra P_i , then a map f is piecewise linear if and only if the restriction of f to P_i is piecewise linear for each i ; in particular, f being linear on each P_i is sufficient [44, Example 1.5, p. 5]. Conversely, if $f: P \rightarrow Q$ is a piecewise linear map, then there is a representation of P as the locally finite union of simplices such that for each simplex, the restriction of f to that simplex is a linear map [44, Corollary 2.3, p. 12]. Therefore, we obtain an equivalent characterisation of piecewise linear maps in terms of the existence of a representation of the domain as a union of polytopal or simplicial pieces where the map is linear: A map $f: P \rightarrow Q$ between polyhedra P and Q is piecewise linear if and only if there is a representation of P as the locally finite union of simplices such that for each simplex, the restriction of f to that simplex is a linear map. Analogously, f is piecewise linear if and only if P is representable as a locally finite union of convex polytopes such that f is linear on each polytope.

Let us consider now the case that P is a compact polyhedron, in other words the domain of some polytopal or simplicial complex K . It turns out that for this case, piecewise linear maps correspond to maps linear on cells of a suitable subdivision of the complex. From the above characterisations of piecewise linear maps, we can immediately conclude that maps linear on cells of a polytopal complex K are piecewise linear, because the cells of the complex are compact polytopes. As a side note, this result extends to maps linear on cells of polyhedral complexes: The non-compact cells are polyhedra and hence locally finite unions of compact polyhedra; this yields a representation of the domain as a locally finite union of compact polyhedra where the map is linear, implying piecewise linearity.

We want to consider maps from $|K|$ to \mathbb{R} , but the results in the literature for piecewise linear maps on complexes are usually stated for maps between domains of complexes. Therefore, we start with some observations explaining why the codomain can be assumed to be the domain of some other polytopal or simplicial complex, so that the results carry over to arbitrary codomains.

For any piecewise linear map defined on a compact polyhedron, its image is

again a compact polyhedron [44, Corollary 2.5, p. 13]. Thus, given a piecewise linear map $f: |K| \rightarrow X \subseteq \mathbb{R}^m$ on a polytopal complex K , we can restrict the codomain to the image $f(|K|) \subseteq X$ of f . Since $f(|K|)$ is a compact polyhedron, it is the underlying set of some simplicial complex L embedded in \mathbb{R}^m and we may consider f as a piecewise linear map $f: |K| \rightarrow |L|$. Alternatively, we can use the fact that images of compact sets are compact. Hence the image of f is contained in some large enough hypercube, and choosing for L the polytopal complex associated with the hypercube works as well. We also could replace the hypercube by some other large enough convex polytope or simplex, if we prefer.

The relation between piecewise linear maps, maps linear on cells of a subdivision, and cellular or simplicial maps on subdivisions is established by the following results: If the map $f: |K| \rightarrow |L|$ between polytopal complexes K and L is linear on cells of K , then K and L have simplicial subdivisions such that f is simplicial with respect to these subdivisions [44, Lemma 2.13, p. 17]. If $f: |K| \rightarrow |L|$ is piecewise linear, the same consequence follows: The map f is simplicial with respect to some suitable subdivisions of K and L to simplicial complexes [44, Theorem 2.14, p. 17]. Combining these observations with the previous ones shows that the following properties for maps $f: |K| \rightarrow |L|$ between polytopal complexes K and L are equivalent:

1. The map f is piecewise linear.
2. The map f is linear on cells of a suitable subdivision of K .
3. The map f is cellular with respect to suitable subdivisions of K and L .
4. The map f is simplicial with respect to suitable simplicial subdivisions of K and L .

Forgetting about the structure as complex on $|L|$, we obtain the following characterisation of piecewise linear maps on polytopal complexes with arbitrary codomains: A map $f: |K| \rightarrow X \subseteq \mathbb{R}^m$ on a polytopal complex K is piecewise linear if and only if K has a simplicial subdivision such that f is linear on cells of that subdivision; replacing “simplicial subdivision” by “subdivision” yields another variant of the equivalence.

The Cone Construction. A frequent construction in PL topology is the cone construction, which extends PL maps between complexes to cones of the complexes. This construction can be considered as a special case of a join of maps defined as follows [44, Remark 2.22 (4), p. 23]: Let AB and CD be joins, with at least AB adhering to our convention of being a join of independent sets or an external join. For two maps $f: A \rightarrow C$ and $g: B \rightarrow D$, the join $h: AB \rightarrow CD$ of the maps is defined by the following assignment: To a point in AB represented by $\lambda a + (1 - \lambda)b$ for some $a \in A$, $b \in B$, and $\lambda \in [0, 1]$, we assign the value $\lambda f(a) + (1 - \lambda)g(b)$. The join of PL maps is again PL. If A and C are single points v and w , then specifying f explicitly is unnecessary, and $h: vB \rightarrow wD$ defined by $h(\lambda v + (1 - \lambda)b) = \lambda w + (1 - \lambda)g(b)$ is called the cone on g [44, Exercise 1.6 (3), pp. 5–6].

Summary. The following summary recapitulates the basic facts for our setting.

Fact 1.17. Domains of polytopal complexes fulfil the following property characterising polyhedra: Any point p has a *cone neighbourhood* $N = pL$, called a *star* of p , with compact base L , called a *link* of p . A map $f: |K| \rightarrow \mathbb{R}^m$ is *piecewise linear*, if it is continuous and locally conical, that is, any point p in $|K|$ has a link L such that $f(\lambda p + (1 - \lambda)r) = \lambda f(p) + (1 - \lambda)f(r)$ for all $r \in L$ and $\lambda \in [0, 1]$; or equivalently, if K has a simplicial subdivision such that f is linear on cells of that subdivision.

1.2.5 Subcomplexes

We list some frequently occurring types of subcomplexes, namely k -skeletons, links, stars, anti-stars, and deletions.

Definition 1.18 (k -skeleton). For arbitrary complexes K , the subset of all cells of dimension at most k forms a subcomplex of K , called the k -skeleton of K and denoted by $\text{skel}^k(K)$ [44, Example 2.8 (8), p. 15].

The subcomplexes described as *stars* and *links* most often occur as the star or the link of a certain vertex v . The star of a vertex v is the set of all cells containing v and all their faces; the link of v consists of all cells of the star of v that do not contain v . For the different characterisations of regular and critical points of piecewise linear maps in Chapter 3, it is convenient to extend these notions to arbitrary points of the domain. Another generalisation is obtained by considering vertices as 0-dimensional cells and define stars and links for arbitrary cells of a complex. The *deletion* of a vertex is the set of cells not containing the vertex. When we reformulate this condition as the set of cells not intersecting the vertex, we obtain the notion of an *anti-star*, that agrees with the deletion for vertices and when generalised to points in the domain, but differs from it when generalised to cells.

The generalised definitions for cells are taken from [23, p. 40] for the star, the link, and the anti-star, and from [2] for the deletion.

Definition 1.19. Let K be a polyhedral complex. For a vertex $v \in \text{Vrt}(K)$, or an arbitrary point $x \in |K|$ in the domain, or a cell $S \in K$, we define subcomplexes of K called the star, the link, the deletion, and the anti-star of v , x , or S in K as follows.

The *star* of v in K , denoted by $\text{st}_K(v)$, is the subcomplex of K consisting of all cells in K that contain the vertex v or are a face of such a cell. Analogously, the star $\text{st}_K(x)$ of x in K is the set of all cells containing the point x plus all their faces, and the star $\text{st}_K(S)$ of S in K is the set of all cells that contain the cell S or are a face of such a cell. In other words, the star of S is the smallest subcomplex containing all co-faces of S .

The *link* of v in K , denoted by $\text{lk}_K(v)$, is the set of cells in $\text{st}_K(v)$ that do not contain v . The link $\text{lk}_K(x)$ of x in K is the set of cells in $\text{st}_K(x)$ that do not contain x , and the link $\text{lk}_K(S)$ of S in K is the set of cells in $\text{st}_K(S)$ that do not intersect S .

The *deletion* of v in K , denoted by $\text{del}_K(v)$, is the set of cells not containing v . Analogously, $\text{del}_K(x)$ is the set of cells not containing x , and $\text{del}_K(S)$ is the set of cells not containing S as a whole.

The *anti-star* of v in K , denoted by $\text{ast}_K(v)$, is the set of cells not intersecting $\{v\}$. Analogously, $\text{ast}_K(x)$ is the set of cells not intersecting $\{x\}$, and $\text{ast}_K(S)$ is the set of cells not intersecting S .

Table 1.2: Relationships between stars, links, deletions, and anti-stars in a polyhedral complex K for a vertex $v \in \text{Vrt}(K)$, a point x in the domain $|K|$, and a non-empty cell $S \in K$.

$$\begin{array}{lll}
\text{ast}(v) = \text{del}(v) & \text{ast}(x) = \text{del}(x) & \text{ast}(S) \subseteq \text{del}(S) \\
\text{lk}(v) = \text{st}(v) \cap \text{ast}(v) & \text{lk}(x) = \text{st}(x) \cap \text{ast}(x) & \text{lk}(S) = \text{st}(S) \cap \text{ast}(S) \\
\text{lk}(v) = \text{st}(v) \cap \text{del}(v) & \text{lk}(x) = \text{st}(x) \cap \text{del}(x) & \text{lk}(S) \subseteq \text{st}(S) \cap \text{del}(S) \\
K = \text{st}(v) \cup \text{del}(v) & K = \text{st}(x) \cup \text{del}(x) & K = \text{st}(S) \cup \text{del}(S)
\end{array}$$

Table 1.3: Relationships of links and stars in simplicial complexes K for a vertex $v \in \text{Vrt}(K)$, a point x in the domain $|K|$ and the simplex X containing x in its interior, and a simplex $S \in K$.

$$\begin{array}{lll}
\text{st}(v) = v \text{lk}(v) & \text{st}(x) = X \text{lk}(X) & \text{st}(S) = S \text{lk}(S) \\
& \text{lk}(x) = \text{bd}(X) \text{lk}(X) &
\end{array}$$

In all cases, we omit the complex K in the subscript, when it is clear from the context, and write $\text{st}(S)$, $\text{lk}(x)$, $\text{del}(v)$, and so on. For all four notions, the definition for a vertex v and the definition for its corresponding 0-dimensional cell $\{v\}$ agree, and vertices are just a special case of arbitrary elements of the domain. Furthermore, when X denotes the cell containing the point x in its interior, then the equalities $\text{st}(X) = \text{st}(x)$ and $\text{del}(X) = \text{del}(x)$ hold; the link and the anti-star of X however are usually only a subcomplex of the corresponding complex for x .

We conclude this subsection with some observations regarding links, stars, anti-stars, and deletions.

Fact 1.20. The following relationships between the subcomplexes are collected in Table 1.2: For points in the domain, including vertices, deletion and anti-star agree; for non-empty cells S , the anti-star $\text{ast}(S)$ is a subcomplex of the deletion $\text{del}(S)$. The link is always the intersection of star and anti-star [23, p. 40], and for single points, it is also the intersection of star and deletion. Finally, the whole complex K is in all cases representable as the union of the star and the deletion.

For polytopal complexes K , the domain $|\text{st}(x)|$ of the star of a single point x forms a cone neighbourhood of x whose link is the domain of the link, that is $|\text{st}(x)| = x|\text{lk}(x)|$. In other words, as Rourke and Sanderson [44, p. 20] put it, the domains of *the* star and *the* link of a point x in K as defined here are an example of *a* star and *a* link of x in the compact polyhedron $|K|$ in the sense defined in Subsection 1.2.4. For this reason, both definitions of stars and links can coexist peacefully.

In the case of simplicial complexes, the additional relationships between links and stars listed in Table 1.3 hold: For vertices, the representation of the star as cone on the link does not only hold for the domains, but also for the complexes;

the star of a vertex v is the cone on the link of v with apex v . This generalises from vertices to simplices because the star of a simplex is the join of the simplex with its link [44, Remark 2.22 (3), p. 23]. For arbitrary points x in the domain, we can also express $\text{st}(x)$ and $\text{lk}(x)$ as joins using the simplex X containing x in its interior as indicated in Table 1.3.

1.2.6 Subdivisions

In this subsection, we review stellar, derived, and barycentric subdivisions of polytopal complexes. A barycentric subdivision is a special case of a derived subdivision, and a derived subdivision is again a certain stellar subdivision. A stellar subdivision in turn is obtained by a sequence of elementary starrings. Therefore we start with the definitions of starrings and stellar subdivisions, taken from [44, p. 15] and [21, p. 8].

Definition 1.21 (elementary starring). Let K be a polytopal complex, and let p be a point in its domain $|K|$. The following operation is called an *elementary starring* of K at p and produces a subdivision of K : Remove all cells containing p from K , and add for each cell S that is a face of a removed cell but not removed itself the cell pS . In other words, we replace the star of p by the cone on its link with apex p and obtain the complex $\text{del}_K(p) \cup p\text{lk}_K(p)$.

Definition 1.22 (stellar subdivision). A *stellar subdivision* of a polytopal complex K is any subdivision that is obtained by a sequence of elementary starrings starting from K .

The construction of a derived subdivision of a polytopal complex K requires the choice of a family of points: for each non-empty cell S in K , a point v_S in the interior of S ; these points form the vertex set of the subdivision. There are three common ways of characterising a derived subdivision. The first characterisation constructs the derived subdivision as a stellar subdivision obtained by starring at each point v_S in order of decreasing dimension of S [26, p. 9]. The second characterisation constructs the subdivision starting from the complex of the empty cell by iteratively adding a subdivision of S in order of increasing dimension of S , namely the subdivision of S formed by the cone with apex v_S whose base is the derived subdivision of the boundary of S obtained by the previous steps [44, p. 20]; clearly, the choice of vertices for the derived subdivision of the boundary of S has to be the restriction of the choice made for the derived subdivision of the whole complex K . The third method directly characterises the vertex sets whose convex hulls form the cells of the subdivision: The convex hull of the vertex set $\{v_{S_0}, \dots, v_{S_k}\}$ is a cell of the derived subdivision if and only if the corresponding face set $\{S_0, \dots, S_k\}$ forms a chain in the face poset of K , that is $S_{\sigma(0)} \sqsubseteq S_{\sigma(1)} \sqsubseteq \dots \sqsubseteq S_{\sigma(k)}$ for some permutation σ of the indices [21, p. 7].

Definition 1.23 (derived subdivision). A subdivision of a polytopal complex K is called a *derived subdivision* if it is a stellar subdivision obtained by choosing for each non-empty cell S a point v_S in its interior and performing elementary starrings at the points v_S in order of decreasing dimension of S . The cells of the derived subdivision can be characterised as convex hulls of vertex sets $\{v_{S_0}, \dots, v_{S_k}\}$ such that the corresponding face set $\{S_0, \dots, S_k\}$ forms a chain

in the face poset of K . Each non-empty cell S of K gets subdivided into the cone with apex v_S and with the derived subdivision restricted to the boundary of S as base. A derived subdivision of a polytopal complex is always a simplicial complex; it is sometimes called a first derived, and an r -th derived is defined inductively as a derived of a $(r - 1)$ -th derived [44, p. 20].

For constructing a derived subdivision algorithmically, we need a subroutine for finding points in the interior of a cell. When the vertex set $\{v_1, \dots, v_k\}$ of the cell is given, the simplest term that is guaranteed to represent a point in the interior is the coordinate-wise arithmetic mean of the vertices $\frac{1}{k} \sum_{i=1}^k v_i$ [23, Exercise 2.3.6, p. 17]; the resulting point is the barycentre of the vertices, that is the centre of mass with respect to a mass distribution concentrated at the vertices, with equal weights at each vertex. One could also use the barycentre with respect to a mass that is uniformly distributed throughout the volume of the cell. For simplices, both barycentres agree, but for arbitrary convex polytopes, the barycentre with respect to a uniform mass distribution can differ from the barycentre of the vertices and its computation is less straightforward. Therefore we usually prefer using the barycentres of the vertices.

Definition 1.24 (barycentric subdivision). By a *barycentric subdivision* we mean a derived subdivision, where for each cell S the barycentre of its vertices is chosen as v_S .

The fact that any derived subdivision of a polytopal complex is simplicial shows that any polytopal complex can be triangulated, that is, it can be subdivided into a simplicial complex. Interestingly, it is even possible to triangulate a polytopal complex without introducing new vertices; such a triangulation can be obtained by a stellar subdivision that performs elementary starrings at each vertex of the given polytopal complex in some arbitrary order [44, Proposition 2.9 and following Exercise, p. 16].

Fact 1.25. Any polytopal complex can be triangulated without additional vertices. Any derived subdivision yields a triangulation, and a stellar subdivision consisting of elementary starrings at each vertex of the given complex in some arbitrary order yields a triangulation without additional vertices.

A triangulation obtained by elementary starrings at each vertex is sometimes called a pulling triangulation. This notion comes from the following operation described in [23, p. 80]: Start with a convex polytope defined as the convex hull of some vertices. Modify the polytope by “pulling” one of its vertices v , that is, move the vertex in a suitable direction slightly outwards by some small enough amount such that the original position of the vertex is contained in the interior of the modified convex polytope. Then the face poset of the boundary of this modified polytope is isomorphic to the one obtained from the original polytope boundary complex by starring at v , with the isomorphism assigning the resulting modified vertex to v and all other vertices to themselves.

1.2.7 Combinatorial Equivalences and Piecewise Linear Homeomorphisms

Combinatorial Equivalences. In our definition of polyhedral complexes, we noted that the face relation among the cells induces a poset structure on the

complex. The notion of combinatorial equivalence is based on that structure. Grünbaum defines the notion for what we call polytopal complexes [23, p. 199]: Two polytopal complexes are *combinatorially equivalent* if their induced poset structures are isomorphic. Rourke and Sanderson call combinatorially equivalent complexes *abstractly isomorphic* [44, p. 20]. It is a well known fact that combinatorially equivalent simplicial complexes are also PL-homeomorphic, and that, even stronger, two simplicial complexes are PL-homeomorphic if and only if they have combinatorially equivalent subdivisions. One purpose of this subsection is to recall some background of this result and to observe that it generalises to polytopal complexes. Furthermore, we extend the notion of combinatorial equivalence to polyhedral complexes for convenient formulations of our results on level set complexes in Chapter 2.

Let us see what would happen if we adopted the plain characterisation by isomorphic posets to polyhedral complexes. Such a characterisation would consider the complex $\{\mathbb{R}, \emptyset\}$ of the whole real line and the complex $\{\{h\}, \emptyset\}$ of a singleton as equivalent. But we can avoid this undesired effect if we additionally require the poset isomorphism to respect the dimension of the cell.

Definition 1.26 (combinatorial equivalence). Two polyhedral complexes K and L are *combinatorially equivalent* if there is a bijection $\phi^{\text{comb}}: K \rightarrow L$ between them that is a poset isomorphism with respect to the poset structure given by the face relation and that maintains the dimensions of the cells. That means that ϕ^{comb} maps the cells of K to the cells of L such that $\dim(\phi^{\text{comb}}(S)) = \dim(S)$ holds for any cell S of K , and for any pair (S, T) of cells of K , the cell S is a face of T in K if and only if $\phi^{\text{comb}}(S)$ is a face of $\phi^{\text{comb}}(T)$ in L :

$$S \sqsubseteq T \iff \phi^{\text{comb}}(S) \sqsubseteq \phi^{\text{comb}}(T)$$

When referring to the map ϕ^{comb} establishing the combinatorial equivalence of the complexes, we often call the map itself a *combinatorial equivalence* between K and L for short.

Note that in the case of two polytopal complexes, the requirement to respect the dimension is fulfilled by any poset isomorphism [23, p. 38]. Therefore the restriction of our notion of combinatorial equivalence to polytopal complexes agrees with Grünbaum's definition. Moreover, there is yet another characterisation of combinatorial equivalence for polytopal complexes, namely in terms of vertices. This follows from the fact that a convex polytope is the convex hull of its vertices. We say that a subset W of vertices of some complex K spans a cell S in K if S is the convex hull of the vertices in W .

Fact 1.27. Two compact polytopal complexes K and L are combinatorially equivalent if and only if there is a bijection $\phi^{\text{Vrt}}: \text{Vrt}(K) \rightarrow \text{Vrt}(L)$ mapping the vertices of K to the vertices of L such that a subset of vertices spans a cell in one complex if and only if its corresponding subset of vertices spans a cell in the other complex. In other words, the bijection ϕ^{Vrt} fulfils the condition that for all subsets $W \subseteq \text{Vrt}(K)$ of vertices of K , the convex hull of W is a cell of K if and only if the convex hull of the image of W under ϕ^{Vrt} is a cell of L :

$$\text{conv}(W) \in K \iff \text{conv}(\phi^{\text{Vrt}}(W)) \in L$$

For a proof, one easily checks that such a bijection ϕ^{Vrt} of the vertex sets can be transformed into a combinatorial equivalence ϕ^{comb} and vice versa as follows:

Assume first that a bijection ϕ^{Vrt} as above exists. In order to define $\phi^{\text{comb}}(S)$ for a cell S of K , consider the vertex set of S and its image under ϕ^{Vrt} . Because of the condition on ϕ^{Vrt} , the image spans a cell of L , and we can define ϕ^{comb} by assigning this cell to S :

$$\phi^{\text{comb}}(S) = \text{conv}(\phi^{\text{Vrt}}(\text{Vrt}(S)))$$

On the other hand, the restriction of a combinatorial equivalence $\phi^{\text{comb}}: K \rightarrow L$ to the set of vertices $\text{Vrt}(K)$ gives the corresponding vertex set bijection ϕ^{Vrt} . There is one subtlety to keep in mind here: Vertices are elements of the ambient Euclidean space, whereas cells are subsets of the space. Hence, formally speaking, we should distinguish between a vertex $v \in |K|$ in the domain of a complex and its corresponding cell $\{v\} \in K$ in the complex formed by the singleton set containing the vertex. The equation

$$\{\phi^{\text{Vrt}}(v)\} = \phi^{\text{comb}}(\{v\})$$

is the formally correct way to express ϕ^{Vrt} in terms of ϕ^{comb} . By abuse of notation, we may identify vertices with their singleton set cells and write $\phi^{\text{Vrt}}(v) = \phi^{\text{comb}}(v)$ instead. We will continue using this identification of singleton cells with their element whenever it seems preferable.

The close connection between vertex bijections and cell bijections establishing a combinatorial equivalence of complexes allows to switch freely between the two points of view. Therefore we also refer to the vertex bijection ϕ^{Vrt} between combinatorially equivalent complexes as a combinatorial equivalence. When we want to distinguish the two points of view, we call ϕ^{Vrt} a combinatorial equivalence in terms of vertices and ϕ^{comb} a combinatorial equivalence in terms of cells or in terms of face posets.

Recall from Definition 1.10 that a map $f: |K| \rightarrow |L|$ from the domain of a polytopal complex K to the domain of a polytopal complex L is called cellular if it maps cells linearly to cells. If such a cellular map is bijective, then it also induces a bijection between the cells. Hence, assigning to each cell S in K its image $f(S)$ in L under some bijective cellular map $f: |K| \rightarrow |L|$ yields a combinatorial equivalence in terms of cells from K to L , and the restriction of f to the vertex set of K yields the corresponding combinatorial equivalence in terms of vertices.

Notions Maintaining Combinatorial Equivalences. For several of the constructions in the previous subsections, the combinatorial structure of the resulting complex depends only on the combinatorial structures of the underlying complexes. Therefore, applying these constructions to combinatorially equivalent input complexes yields combinatorially equivalent output complexes. We collect some statements of this kind.

Fact 1.28. This collection of statements covers operations on complexes, sub-complexes, and subdivisions.

If K and K' are combinatorially equivalent polyhedral complexes, and so are L and L' , then the Cartesian products $K \times L$ and $K' \times L'$ are combinatorially equivalent, and for polytopal complexes, the joins KL and $K'L'$ are combinatorially equivalent as well. In particular, all cones on a polytopal complex K are combinatorially equivalent to each other and to each cone on a complex

Table 1.4: Corresponding stars, links, deletions, and anti-stars of combinatorially equivalent complexes K and L . By abuse of notation, we use the letter ϕ to denote both the combinatorial equivalence in terms of cells $\phi^{\text{comb}}: K \rightarrow L$ and the corresponding combinatorial equivalence in terms of vertices $\phi^{\text{Vrt}}: \text{Vrt}(K) \rightarrow \text{Vrt}(L)$. The results are stated for a vertex $v \in \text{Vrt}(K)$ in the first column and for a cell $S \in K$ in the last column. The middle column states the results for a point $x \in |K|$ together with a point $y \in |L|$ in a corresponding cell, that is, the cell $X \in K$ containing x in its interior corresponds via ϕ^{comb} to the cell $Y = \phi^{\text{comb}}(X)$ containing y in its interior.

$$\begin{array}{lll} \phi(\text{st}_K(v)) = \text{st}_L(\phi(v)) & \phi(\text{st}_K(x)) = \text{st}_L(y) & \phi(\text{st}_K(S)) = \text{st}_L(\phi(S)) \\ \phi(\text{lk}_K(v)) = \text{lk}_L(\phi(v)) & \phi(\text{lk}_K(x)) = \text{lk}_L(y) & \phi(\text{lk}_K(S)) = \text{lk}_L(\phi(S)) \\ \phi(\text{del}_K(v)) = \text{del}_L(\phi(v)) & \phi(\text{del}_K(x)) = \text{del}_L(y) & \phi(\text{del}_K(S)) = \text{del}_L(\phi(S)) \\ \phi(\text{ast}_K(v)) = \text{ast}_L(\phi(v)) & \phi(\text{ast}_K(x)) = \text{ast}_L(y) & \phi(\text{ast}_K(S)) = \text{ast}_L(\phi(S)) \end{array}$$

K' combinatorially equivalent to K . Analogously, for combinatorially equivalent polytopal complexes K and K' , all suspensions of K are combinatorially equivalent to each other and to each suspension of K' . In all cases, the combinatorial equivalence between the resulting complexes arises naturally from the original equivalences: If $\phi^{\text{comb}}: K \rightarrow K'$ and $\psi^{\text{comb}}: L \rightarrow L'$ denote the original combinatorial equivalences, then for cells $S \in K$ and $T \in L$ the assignment $S \times T \mapsto \phi^{\text{comb}}(S) \times \psi^{\text{comb}}(T)$ gives a combinatorial equivalence from $K \times L$ to $K' \times L'$ and the assignment $ST \mapsto \phi^{\text{comb}}(S)\psi^{\text{comb}}(T)$ gives a combinatorial equivalence from KL to $K'L'$.

Now let K and L be combinatorially equivalent polyhedral complexes with corresponding combinatorial equivalences $\phi^{\text{comb}}: K \rightarrow L$ and $\phi^{\text{Vrt}}: \text{Vrt}(K) \rightarrow \text{Vrt}(L)$. Then corresponding subcomplexes of K and L are combinatorially equivalent and certain subdivision steps can be performed in an analogous way on both complexes so that the resulting subdivisions are combinatorially equivalent. In detail, we have the following results:

Any subcomplex K' of K is combinatorially equivalent to its corresponding subcomplex $\phi^{\text{comb}}(K')$ of L with the restriction of ϕ^{comb} to the subcomplexes being the combinatorial equivalence. For example, the k -skeleton of L corresponds to the k -skeleton of K under the combinatorial equivalence: $\text{skel}^k(L) = \phi^{\text{comb}}(\text{skel}^k(K))$. Similar results hold for stars, links, deletions, and anti-stars of vertices, cells, or points that are corresponding under the combinatorial equivalence, where points can be considered as corresponding, if the cells containing the points in their interiors are corresponding; see Table 1.4 for a list of such results.

Combining the previous observations with the fact from Definition 1.21 that an elementary starring of K at x produces the complex $\text{del}_K(x) \cup x \text{lk}_K(x)$, we can conclude that elementary starring at corresponding points of combinatorially equivalent polytopal complexes yields combinatorially equivalent subdivisions. In particular, if x and x' are in the interior of the same cell of a polytopal complex K , the subdivision obtained by starring K at x and the subdivision obtained by starring K at x' are combinatorially equivalent. De-

note by $K' = \text{del}_K(x) \cup x \text{lk}_K(x)$ the result of starring K at a point x and by $L' = \text{del}_L(y) \cup y \text{lk}_L(y)$ the result of starring L at a point y corresponding to x . Then both subdivisions contain at most one additional vertex compared to the original complexes, namely x and y ; in other words, the vertex sets of K' and L' can be written as $\text{Vrt}(K') = \text{Vrt}(K) \cup \{x\}$ and $\text{Vrt}(L') = \text{Vrt}(L) \cup \{y\}$. Therefore the combinatorial equivalence between K' and L' can be easily expressed in terms of vertices: Extend the original combinatorial equivalence ϕ^{Vrt} between the vertex sets of K and L to the vertex sets of K' and L' by assigning the additional vertex $y \in \text{Vrt}(L')$ to the additional vertex $x \in \text{Vrt}(K')$.

In Definition 1.23, the cells of a derived subdivision were characterised by chains in the original face posets. This characterisation shows that the combinatorial structure of a derived subdivision is uniquely determined by the combinatorial structure of the original complex and is independent of the choice of the points v_S used as vertices of the subdivision. Hence, derived subdivisions of combinatorially equivalent polytopal complexes are combinatorially equivalent. This applies in particular to all derived subdivisions of a single polytopal complex obtained by different choices of points used as vertices. A combinatorial equivalence in terms of vertices between a derived of K with chosen points v_S for each $S \in K$ and a derived of L with chosen points u_T for each $T \in L$ is given by assigning $u_{\phi^{\text{comb}}(S)}$ to v_S .

More on Subdividing and Triangulating Combinatorially Equivalent Complexes. Regarding the fact that elementary starrings at corresponding points of combinatorially equivalent complexes produces combinatorially equivalent complexes, if we want to generalise this fact from single starrings to stellar subdivisions by induction, we have to be careful. It is not sufficient in general, that the points at which a starring is performed lie in corresponding cells of the starting complexes, as the counterexample in Figure 1.2 shows. Rather they have to lie in corresponding cells of the subdivisions obtained by the previous starrings. Therefore we can at least infer by induction that for any choice of points in one complex there is a choice of corresponding points in the other complex leading to combinatorially equivalent stellar subdivisions: Assuming we have a combinatorial equivalence ϕ_{j-1}^{comb} after the first $j-1$ elementary starrings and the given point for the next starring in one complex is interior to some cell S_j , then the point for the next starring in the other complex can be chosen arbitrarily in the interior of the cell corresponding to S_j via ϕ_{j-1}^{comb} .

Fact 1.29. Let K and L be combinatorially equivalent polytopal complexes, and let K' be a stellar subdivision of K obtained by a sequence of starrings at points $(p_i)_{i=1}^n$ in $|K|$. Then there is a choice of a corresponding sequence of points $(q_i)_{i=1}^n$ in $|L|$ such that the stellar subdivision of L obtained by the corresponding sequence of starrings at the points $(q_i)_{i=1}^n$ is combinatorially equivalent to K' .

Although derived subdivisions are a special case of stellar subdivisions, we already observed in Fact 1.28, relying on a different characterisation of derived subdivisions, that we can choose the points for the starrings freely in the interior of each cell of the original complex, without changing the combinatorial type of the results. In this special case of a stellar subdivision, the problem that points originally belonging to corresponding cells fail to do so after some subdivision steps does not occur because of the order of the starrings: A cell is only sub-




Figure 1.2: Example of three combinatorially different stellar subdivisions that use points lying originally in the interior of the same cell. The top row shows the original complex of two triangles sharing an edge with a point on a boundary edge for the first starring and a point in the interior of the adjacent triangle for the second starring; only the position of the second point inside the triangle differs in each column. The positions are chosen in such a way that the points belong to the interior of different cells after the first starring. The bottom row shows the resulting complexes after performing the second starring. The results are pairwise non-equivalent, which can be verified by comparing the number of vertices with exactly three incident edges.

divided into several strictly smaller cells if the starring is performed at a point which is contained in the cell as interior or boundary point; but the starring at the chosen interior point of the cell is always performed before all other starings potentially affecting the cell. The problem does also not occur for the other method suggested in Fact 1.25 for triangulating polytopal complexes, namely starring at the original vertices. Hence we can conclude that combinatorially equivalent polytopal complexes have combinatorially equivalent triangulations.

Lemma 1.30. *Let K and L be combinatorially equivalent polytopal complexes; then there are combinatorially equivalent simplicial subdivisions K' of K and L' of L . Furthermore, it is not necessary to introduce new vertices.*

Proof. Fact 1.25 mentions two kinds of subdivisions that produce simplicial complexes. We may construct combinatorially equivalent triangulations with both methods. The first method of taking derived subdivisions of both complexes introduces new vertices; as already stated in Fact 1.28, the resulting triangulations are combinatorially equivalent. The second method of starring at each vertex in corresponding orders does not introduce new vertices and hence proves the additional statement; it produces combinatorially equivalent subdivisions as observed above.

As a special case of the first method, we can use the barycentric subdivisions from Definition 1.24. For each cell S of any dimension of the polytopal complexes K and L with vertices v_1, \dots, v_k , we choose for v_S the point $\frac{1}{k}(v_1 + \dots + v_k)$ in the interior of the cell and obtain derived subdivisions K' of K and L' of L . The combinatorial equivalence of K and L in terms of vertices $\phi^{\text{Vrt}}: \text{Vrt}(K) \rightarrow \text{Vrt}(L)$ extends to a combinatorial equivalence $\psi^{\text{Vrt}}: \text{Vrt}(K') \rightarrow \text{Vrt}(L')$ of K' and L' by setting $\psi^{\text{Vrt}}(\frac{1}{k}(v_1 + \dots + v_k)) = \frac{1}{k}(\phi^{\text{Vrt}}(v_1) + \dots + \phi^{\text{Vrt}}(v_k))$.

For the second method, we perform an elementary starring at each vertex for some fixed order v_1, \dots, v_k of the vertices of K , and the corresponding order $\phi^{\text{Vrt}}(v_1), \dots, \phi^{\text{Vrt}}(v_k)$ of the vertices of L . By induction, since all starrings are performed at vertices, ϕ^{Vrt} remains a combinatorial equivalence in terms of vertices between each intermediate pair of subdivisions obtained by performing the first $j \leq k$ starrings on both complexes, and for each index i from 1 to k , the vertices v_i and $\phi^{\text{Vrt}}(v_i)$ remain interior to the same pair of corresponding cells, namely $\{v_i\}$ and $\{\phi^{\text{Vrt}}(v_i)\}$. Therefore, the observations from Fact 1.28 regarding elementary starrings at corresponding points of combinatorially equivalent complexes, applied iteratively, imply that the obtained simplicial subdivisions of K and L are combinatorially equivalent with the same equivalence in terms of vertices as the original one. \square

We observed in Fact 1.29 that any stellar subdivision of a polytopal complex can be imitated on any combinatorially equivalent complex by an analogous subdivision. For arbitrary subdivisions however, there are some limitations. Assume that K and L are combinatorially equivalent polytopal complexes and K' is a subdivision of K . In general, it is not true that a subdivision L' of L exists such that K' and L' are combinatorially equivalent. A counterexample is a regular octahedron and a slightly perturbed combinatorially equivalent octahedron whose vertices are in general position, that is, no four of them in a common plane. The latter has subdivision containing a tetrahedron whose corresponding vertices in the regular octahedron lie in a common plane. Therefore the regular octahedron cannot be subdivided in such a way.

PL Homeomorphisms. Now we turn to the relationship between combinatorial equivalences and PL homeomorphisms. Recall that a homeomorphism is defined as a continuous bijection with continuous inverse map. For a PL homeomorphism it is additionally required that both maps, the bijection and its inverse, are piecewise linear. Since piecewise linear maps are continuous, it suffices to characterise a PL homeomorphism as a bijective piecewise linear map with piecewise linear inverse. For our usual setting, we can even prove that any piecewise linear bijection between polytopal complexes has a piecewise linear inverse: If K and L are polytopal complexes and $\phi: |K| \rightarrow |L|$ is a PL bijection, then K and L have subdivisions K' and L' such that ϕ is simplicial. Now consider ϕ restricted to some simplex S of K' and its inverse ϕ^{-1} restricted to the corresponding simplex $T = \phi(S)$ of L' . Then the restricted map ϕ extends to an affine linear map of full rank from the affine hull of S to the affine hull of T and this extension has an affine linear inverse that agrees with ϕ^{-1} on T . This proves that ϕ^{-1} is linear on each cell of L' and thus a piecewise linear map. Note that analogous arguments show that a cellular bijection has a cellular inverse and that a simplicial bijection has a simplicial inverse. Two polyhedra P and Q are called PL-homeomorphic, if there is a PL homeomorphism $\phi: P \rightarrow Q$, and two polyhedral complexes are called PL-homeomorphic if their domains are.

Let us make a remark on injective PL maps. If $\phi: P \rightarrow Q$ is a PL injection and P is a compact polyhedron, then the image $\phi(P)$ of ϕ in Q is a compact polyhedron as well because images of compact polyhedra under PL maps are compact polyhedra [44, Corollary 2.5, p. 13]. Then P and $\phi(P)$ are domains of simplicial complexes and ϕ establishes a PL bijection between the domains, hence P and $\phi(P)$ are PL-homeomorphic. But if P is non-compact, the image is

not necessarily a polyhedron [44, Exercise 1.7 (1) (c)]. Therefore the definition of a PL embedding in [44, p. 7] explicitly requires the image to be a polyhedron: A PL embedding is a PL map $\phi: P \rightarrow Q$ such that the image $\phi(P)$ is a polyhedron and $\phi: P \rightarrow \phi(P)$ is a PL homeomorphism. By the above, any PL injection defined on a compact polyhedron, in particular on the domain of a polytopal complex, is a PL embedding.

PL Homeomorphisms and Combinatorially Equivalent Subdivisions.

For combinatorially equivalent simplicial complexes K and L , the combinatorial equivalence in terms of vertices $\phi^{\text{Vrt}}: K \rightarrow L$ naturally induces a PL homeomorphism $\phi^{\text{PL}}: |K| \rightarrow |L|$ between the domains of the complexes by simplex-wise linear interpolation. Note that the image of each simplex under this homeomorphism is exactly the simplex assigned to it by the combinatorial equivalence in terms of cells: $\phi^{\text{PL}}(S) = \phi^{\text{comb}}(S)$. For arbitrary polytopal complexes however, it is in general not possible to find a cell-wise linear homeomorphism without subdividing. For example, all convex quadrilaterals are combinatorially equivalent, but a linear bijection maps parallelograms only to parallelograms. By subdividing the quadrilaterals into two triangles by a diagonal, we obtain combinatorially equivalent simplicial complexes and can define a PL homeomorphism by interpolating linearly in both triangles individually. Lemma 1.30 shows that any two combinatorially equivalent polytopal complexes have combinatorially equivalent simplicial subdivisions, which allow to define a PL homeomorphism by simplex-wise linear interpolation. This yields the following close relation between PL homeomorphisms and combinatorially equivalent subdivisions:

Theorem 1.31. *Two polytopal complexes are PL-homeomorphic if and only if they have combinatorially equivalent subdivisions.*

Proof. Let K and L be polytopal complexes and let $\phi^{\text{PL}}: |K| \rightarrow |L|$ be a PL homeomorphism. For the proof that K and L have combinatorially equivalent subdivisions, we use [44, Theorem 2.14, p. 17] to obtain subdivisions K' and L' such that $\phi^{\text{PL}}: |K'| \rightarrow |L'|$ is simplicial. Hence, with respect to these subdivisions, ϕ^{PL} is a cellular bijection, and we may apply our earlier observation that a cellular bijection induces a combinatorial equivalence: The restriction of ϕ^{PL} to the vertex set of K' gives the desired combinatorial equivalence in terms of vertices $\phi^{\text{Vrt}}: \text{Vrt}(K') \rightarrow \text{Vrt}(L')$.

For the other direction, let K' and L' be combinatorially equivalent subdivisions of K and L . By Lemma 1.30, we can assume without loss of generality that K' and L' are simplicial. The combinatorial equivalence $\phi^{\text{Vrt}}: \text{Vrt}(K') \rightarrow \text{Vrt}(L')$ can be extended to a PL map $\phi^{\text{PL}}: |K'| \rightarrow |L'|$ by simplex-wise linear interpolation. The map ϕ^{PL} maps each simplex S of K' bijectively onto its corresponding simplex $\phi^{\text{comb}}(S)$ in L' , where ϕ^{comb} denotes the combinatorial equivalence in terms of cells. Furthermore, for each simplex the bijection restricts to a bijection between the interiors of the corresponding simplices. Since the interiors of the simplices are disjoint and together cover the domain of their complex, the map ϕ^{PL} is also a global bijection between $|K'|$ and $|L'|$. Hence ϕ^{PL} is a piecewise linear bijection between the domains of polytopal complexes and thus, as observed above, a PL homeomorphism. Its inverse is the PL map obtained by simplex-wise linear interpolation of the inverse of ϕ^{Vrt} . \square

Most of our considerations are based only on combinatorial and topological

features of a complex and carry over to all combinatorially equivalent complexes, independently of the geometric location of their cells. Therefore we sometimes do not distinguish between geometrically different instances of combinatorially equivalent complexes. In particular, we identify the results of operations such as taking cones, suspensions, or derived subdivisions, where the geometry depends on a choice of vertices, but all choices lead to combinatorially equivalent results. This point of view justifies speaking about *the* cone, *the* suspension, or *the* derived subdivision of a complex K and writing $\text{cone}(K)$, $\text{susp}(K)$, and $\text{sdv}^{\text{drv}}(K)$ respectively to denote such a resulting complex. This approach is also reflected in the usage of abstract simplicial complexes, which are defined as collections of subsets of a finite base set, closed under taking subsets. One obtains an abstract simplicial complex from a geometric one by taking the vertex set as finite base set, and including for each simplex the subset of the vertices that spans the simplex into the collection of subsets. The face relation is represented by the inclusion relation. It is well known that each abstract simplicial complex can be realised by some geometric simplicial complex. Since all realisations of an abstract simplicial complex are combinatorially equivalent, they are PL-homeomorphic. For a similarly defined notion of abstract polytopal complexes however, there are abstract complexes not realisable in Euclidean space [23, Exercise 11.1.8, pp. 206–210].

1.2.8 Pseudo-Radial Projection

An important basic result of piecewise linear topology is the fact that all polyhedral links of a point in a polyhedron are PL-homeomorphic, proven by the so called *pseudo-radial projection* [44, pp. 20–21]. Recall from Subsection 1.2.4 that for some point p in a polyhedron P , any compact subset L such that pL is a cone neighbourhood of p in P is called a link of p . For any two links of a fixed point p , radial projection with p as centre point establishes a topological homeomorphism between the two links [44, Exercise 1.7 (2), p. 6]. Of course, we primarily consider links that are itself polyhedra and thus domains of simplicial complexes, and we would preferably apply the strengthened result mentioned above, claiming the existence of PL homeomorphisms between any two such polyhedral links of some fixed point. Since radial projection itself is not piecewise linear in general, even if both links are polyhedra, another solution is necessary, namely pseudo-radial projection, which can be perceived as a piecewise linear approximation to the radial projection. Clearly, any PL homeomorphism between links of points extends to a PL homeomorphism between the corresponding cone neighbourhoods by the cone construction, because the cone on a PL homeomorphism is again a PL homeomorphism [44, Exercise 1.6 (3), pp. 5–6]. Hence, all stars with polyhedral link of a fixed point in a polyhedron are PL-homeomorphic.

Link Complexes. From now on, whenever we consider a link of a point in a polyhedron, we tacitly assume that the link is a polyhedron. We also implicitly make this assumption, when we talk about cone neighbourhoods and stars, namely that the link forming the base and consequently the cone neighbourhood itself are polyhedra. In the context of links of points in polyhedra, we call a polytopal complex whose domain is a link of some point a *link complex* of the considered point.

PL Invariance of Links and Stars. The above results can be generalised from two links of a single point to links of two points that correspond to each other by some PL homeomorphism. This generalised result is called “PL invariance of links and stars” by Rourke and Sanderson in [44, pp. 6–7, 20–21]; it states for PL-homeomorphic polyhedra P and Q with PL homeomorphism $\phi: P \rightarrow Q$ and a point $p \in P$ that any link of p in P is PL-homeomorphic to any link of $\phi(p)$ in Q , and that any star of p in P is PL-homeomorphic to any star of $\phi(p)$ in Q . This applies in particular to *the* links and *the* stars if P and Q are domains of polytopal complexes K and L : The links $\text{lk}_K(p)$ and $\text{lk}_L(\phi(p))$ are PL-homeomorphic, and so are the stars $\text{st}_K(p)$ and $\text{st}_L(\phi(p))$.

Now consider two points with PL-homeomorphic links. From a local point of view, since all their cone neighbourhoods are PL-homeomorphic, the points are similar. Therefore we may regard the property of having PL-homeomorphic links as the defining property of an equivalence relation on the class of all pairs (p, P) where P is a polyhedron and p a point in P . We will refine this equivalence relation in Chapter 3 by considering additionally the local behaviour near p of a real-valued PL map defined on P . In preparation of this, we explain some of the underlying ideas as far as they already apply to the unrefined relation and are based on observations associated with the PL invariance of links and stars and its proof by pseudo-radial projection.

Pseudo-Radial Projection. We start with some remarks regarding pseudo-radial projection based on the descriptions in [44, Lemma 2.19, p. 21] and [26, Lemma 1.12, pp. 20–21]. Consider a point p in a polyhedron P and two polytopal link complexes K and L forming links of p in P . This situation occurs for example when we are given two polytopal complexes \mathcal{K} and \mathcal{L} with the same domain $|\mathcal{K}| = |\mathcal{L}|$ and we look at the links $\text{lk}_{\mathcal{K}}(p)$ and $\text{lk}_{\mathcal{L}}(p)$ of a point p in that domain. The aim is to construct a PL homeomorphism between K and L by pseudo-radial projection. Since all points of a link of p lie on distinct rays emanating from p , a standard radial projection from the centre p assigns to each point in some given link of p a corresponding point on the same ray in some other given link of p . Changing roles of both links, radial projection also yields the inverse map of this assignment. In fact, the obtained bijection between the two links is a topological homeomorphism.

Denote by $\pi: K \rightarrow L$ the projection in direction from K to L and by $\pi^{-1}: L \rightarrow K$ its inverse projection in the opposite direction. Now each cell S in K is mapped by π to the union of cells $\bigcup_{T \in L} (T \cap \pi(S))$. Inversely, each cell T in L is mapped by the inverse to the union $\bigcup_{S \in K} (S \cap \pi^{-1}(T))$. We obtain a subdivision K' of K formed by all cells of the form $S \cap \pi^{-1}(T)$ for $S \in K$ and $T \in L$, and analogously a subdivision L' of L formed by all cells of the form $T \cap \pi(S)$ for $S \in K$ and $T \in L$. The cell $S \cap \pi^{-1}(T)$ projects to $T \cap \pi(S)$ by π , and inversely $T \cap \pi(S)$ projects to $S \cap \pi^{-1}(T)$ by π^{-1} . See Figure 1.3 for an example abbreviating $S \cap \pi^{-1}(T)$ by S_T and $T \cap \pi(S)$ by T_S . Furthermore, faces of such cells induced by $S \in K$ and $T \in L$ are induced by the faces of S and T :

$$\begin{aligned} \mathcal{F}(S \cap \pi^{-1}(T)) &= \{F \cap \pi^{-1}(G) \mid F \sqsubseteq S, G \sqsubseteq T\} \\ \mathcal{F}(T \cap \pi(S)) &= \{G \cap \pi(F) \mid F \sqsubseteq S, G \sqsubseteq T\} \end{aligned}$$

Using these observations, one easily checks that the mutual assignment of $S \cap \pi^{-1}(T)$ and $T \cap \pi(S)$ yields a well defined combinatorial equivalence in terms

of cells between the subdivisions K' and L' . Since the assignment is based on the cell-wise images and preimages under π and π^{-1} , we can think of this combinatorial equivalence as the combinatorial equivalence induced by radial projection.

This already implies, using Theorem 1.31, that the link complexes K and L are PL-homeomorphic. The construction of the PL homeomorphism involves the choice of combinatorially equivalent simplicial subdivisions of K' and L' . We would prefer a choice such that the combinatorial equivalence is still induced by radial projection. This can be achieved easily: For any convex polytope contained in a cell S' of K' , its radial projection to $|L'|$ is a combinatorially equivalent convex polytope contained in the projected cell $\pi(S')$ of L' . In particular simplices contained in S' are mapped to simplices contained in $\pi(S')$. Hence the family of all simplices $\pi(X)$, where X ranges over the simplices of a simplicial subdivision of K' , forms a simplicial subdivision of L' . Clearly, this also works with π^{-1} and the roles of K' and L' reversed. In particular, we can adapt the methods from Lemma 1.30 to obtain compatible simplicial subdivisions whose combinatorial equivalence is still given by radial projection: The simplicial subdivisions obtained from K' and L' by starring at each vertex in corresponding orders are mapped to each other by radial projection. If we construct a derived subdivision, we have to choose corresponding points in the interior of the cells to obtain a combinatorial equivalence by projection: A chosen point $v_{S'}$ in the interior of a cell S' in K' corresponds to $u_{T'} = \pi(v_{S'})$, which is chosen as interior point of the cell $T' = \pi(S')$ in L' . After having chosen the simplicial subdivisions, we assign to each vertex in one subdivided link its radially projected corresponding vertex in the other subdivided link and obtain a PL homeomorphism by simplex-wise linear interpolation. This homeomorphism is called a pseudo-radial projection, because it approximates the radial projection by piecewise linear interpolation between the values of the radial projection at the vertices of the combinatorially equivalent simplicial subdivisions.

Pseudo-Radial Projection via Polyhedral Fans. Let us describe one of the construction step for the pseudo-radial projection again from a slightly different point of view, namely the construction for obtaining the combinatorially equivalent subdivisions K' and L' . Consider one of the links, say K . For each cell $S \in K$, construct the convex polyhedron $S^+ = \{\lambda p + (1-\lambda)q \mid q \in S, \lambda \leq 1\}$ formed by all the rays emanating from p and intersecting S [44, Proof of Lemma 2.19, p. 21]. In other words, S^+ is obtained from the cone pS by extending each line segment connecting the apex p with some point in the base S to a ray emanating from p (see Figure 1.3). We can also think of S^+ as the cone pS with “the base removed”, which is a short expression for the following more precise construction: Take a minimal collection of half-spaces whose intersection defines pS ; remove from the collection the half-space whose boundary contains S and whose interior contains p ; the intersection of the remaining half-spaces is S^+ . The polyhedron S^+ is the convex hull of the rays emanating from p and intersecting a vertex of S . Its non-empty faces are obtained from applying the same construction to the faces T of S , yielding faces T^+ of S^+ . A polyhedron that is generated as the convex hull of finitely many rays emanating from a single point a , like S^+ or its non-empty faces T^+ , is called an *extended convex polyhedral cone* from the apex a or an *extended cone* for short.




Figure 1.3: Two links of a point in \mathbb{R}^2 and their combinatorially equivalent subdivisions found by radial projection. The left and the central figure show each a link of a point in the plane. Both links have a highlighted and labelled edge, S on the left and T in the centre. The induced extended cones S^+ and T^+ are also indicated. The figure on the right depicts the combinatorially equivalent subdivisions of both links found by radial projection. The cells of the subdivisions obtained by projecting T to S and vice versa are highlighted, together with the corresponding extended cone $S^+ \cap T^+$. They are labelled by S_T and T_S , where S_T denotes the cell $S \cap \pi^{-1}(T) = S^+ \cap T^+ \cap S$ and T_S denotes the cell $T \cap \pi(S) = S^+ \cap T^+ \cap T$, with π being the map from the first to the second link given by radial projection.

The collection of all S^+ with S ranging over the cells of K forms a polyhedral complex. A polyhedral complex formed by extended convex polyhedral cones with common apex is called a *polyhedral fan* in [16, Definition 2.1.7, p. 46], at least if the common apex is the origin. Let us apply this notion also for the general case with a common apex different from the origin. Back to our construction, we obtained a polyhedral fan K^+ with apex p and the cells from K as removed bases of the extended cones. We can think of the construction of this fan as extending the local situation at p to infinity, with the local situation being described by those cells in the complex obtained from K by starring at p that contain p . Analogously, we may construct a polyhedral fan L^+ with apex p using the cells from L . The intersection complex $K^+ \cap L^+$ is again a polyhedral fan with apex p and a common subdivision of both K^+ and L^+ . We can find for each extended cone cell in this intersection complex a corresponding pair of a cell S in K and a cell T in L such that the extended cone consists exactly of those rays emanating from p that intersect K in S and that intersect L in T . When we now take again the intersection complex of $K^+ \cap L^+$ with K , we obtain K' because the cells $S^+ \cap T^+ \cap S$ and $S \cap \pi^{-1}(T)$ coincide; analogously, the intersection complex of $K^+ \cap L^+$ with L is L' , where $S^+ \cap T^+ \cap T$ coincides with $\pi(S) \cap T$ (see Figure 1.3, where $S^+ \cap T^+ \cap S$ is labelled S_T and $S^+ \cap T^+ \cap T$ is labelled T_S).

Let us draw some conclusions from the above construction. Suppose that K and L induce the same polyhedral fans $K^+ = L^+$. Then K' agrees with K and L' agrees with L which means that K and L are already combinatorially equivalent without subdividing and with the equivalence induced by radial projection. Figure 1.4 depicts three such combinatorially equivalent links of a point in the plane inducing the same polyhedral fan. Since we obtained K^+ by removing the bases from the cones pS for S in K , we can find other links of p




Figure 1.4: Three links of a point in \mathbb{R}^2 inducing the same polyhedral fan. The first and second link from left to right are scaled copies of the combinatorially equivalent subdivisions obtained in the example from Figure 1.3. The link on the right side was obtained by choosing on each ray forming a 1-dimensional cell of the polyhedral fan some arbitrary vertex.

combinatorially equivalent to K by suitably reinserting bases in the extended cones S^+ such that the bases form a valid complex with a domain contained in the polyhedron P . We can visualise this as tilting or translating the half-spaces and their bounding hyperplanes originally defining the bases in such a way that the hyperplanes still intersect all rays of the extended cone inside P and that the resulting movements of bases of neighbouring cones agree on their common intersection. Another way to generate tilted or translated bases is to move the vertices spanning the cells forming the bases, namely the vertices of K : Each vertex of K is allowed to move along the interior of the ray emanating from p and passing through the respective vertex. Then we obtain a polyhedral link of p in P that is combinatorially equivalent to K , provided that for each subset of vertices spanning a base, that is, a cell of K , the moved counterparts still span a cell contained in P combinatorially equivalent to the cell spanned before the moves. If K is simplicial, the requirement of combinatorial equivalence is always fulfilled, because for each subset of vertices spanning a base, the points of this subset and the point p are affinely independent, and this property still holds when the points different from p move along their rays emanating from p . The requirement of being contained in P on the other hand is always fulfilled if all vertices move only in direction towards p because in this case for each original base cell $S \in K$, its moved base cell is contained in the cone pS , which again is contained in P by the assumption that K is a link of p in P .

Summary. A short summary of our observations concludes this subsection.

Fact 1.32. Let $\phi: P \rightarrow Q$ be a PL homeomorphism between polyhedra P and Q mapping $p \in P$ to $q \in Q$. Then all polyhedral links of p in P and all polyhedral links of q in Q are PL-homeomorphic to each other. By the cone construction, all stars of p and q obtained as cones of these polyhedral links are also PL-homeomorphic. The different links of the point $p \in P$ are PL-homeomorphic by pseudo-radial projection. For a cell S of a link complex K of p , let S^+ be the extended cone from p with base S . The collection of all S^+ with $S \in K$ forms a polyhedral fan K^+ centred at p . For two link complexes K and L of p , the intersection complex $K' = (K^+ \cap L^+) \cap K$ is a subdivision of K and the intersection complex $L' = (K^+ \cap L^+) \cap L$ is a subdivision of L combinatorially

equivalent to K' . The combinatorial equivalence is induced by cell-wise radial projection, assigning $(S^+ \cap T^+) \cap S \in K'$ to $(S^+ \cap T^+) \cap T \in L'$ for $S \in K$ and $T \in L$. Simplicial subdivisions maintaining cell-wise radial projection define a pseudo-radial projection by simplex-wise linear interpolation of the radial projection at the vertices.

Two link complexes inducing the same polyhedral fan are combinatorially equivalent already without subdivisions by cell-wise radial projection. For a given link complex K of p , we can construct combinatorially equivalent links with the same polyhedral fan by moving or tilting the bases of the cones in pK slightly, or equivalently, moving the vertices of K along rays emanating from p , with these movements being subject to some restrictions ensuring that we obtain a valid polyhedral link in P . These restricting conditions are met for example if K is simplicial and the vertices are moved in direction towards p from their original position.

1.2.9 Combinatorial Manifolds

There are several ways of characterising manifolds in piecewise linear topology. First of all, there are characterisations defining manifolds in the category of polyhedra, which are called *piecewise linear manifolds* [26, p. 20], [44, p. 7]; and there are characterisations defining manifolds in the category of simplicial complexes, which are called *combinatorial manifolds* [21, p. 18], [26, p. 26]. Since the domain of a simplicial complex is a compact polyhedron and vice versa, both categories overlap. Somewhat unsurprisingly, both notions of manifolds agree in that case: A simplicial complex is a combinatorial manifold if and only if its domain is a piecewise linear manifold [26, p. 26], [44, Exercise 2.21 (1), p. 22].

Piecewise Linear Manifolds. We start with characterisations of PL manifolds. Rourke and Sanderson [44, p. 7] define a piecewise linear d -manifold without boundary as a polyhedron in which any point has a neighbourhood PL-homeomorphic to an open set in \mathbb{R}^d . For a d -manifold with boundary, also points with neighbourhoods PL-homeomorphic to open subsets of a half-space in \mathbb{R}^d are allowed. The boundary of such a manifold consists of the points whose neighbourhood is PL-homeomorphic to an open subset of a half-space in such a way that the point in question is mapped to the boundary of the half-space. The other points form the interior of the manifold. This definition strongly resembles the common characterisation of a topological manifold, with topological spaces replaced by polyhedra and continuous maps replaced by piecewise linear ones.

For our study, this characterisation has the drawback that it deals with open neighbourhoods, whereas we prefer to consider cone neighbourhoods, namely the stars of points in simplicial complexes, which are closed neighbourhoods. A characterisation of PL manifolds in terms of closed neighbourhoods is given by Hudson [26, p. 20]. It starts by defining PL balls and PL spheres:

Definition 1.33 (PL ball, PL sphere). A *piecewise linear d -ball* is a polyhedron PL-homeomorphic to a d -simplex. A *piecewise linear d -sphere* is a polyhedron PL-homeomorphic to the boundary of a $(d + 1)$ -simplex.

Clearly, instead of taking a simplex as reference for defining balls and spheres, we could replace it by any other PL ball of the same dimension. Every compact

convex polytope of dimension d is a PL d -ball [26, Lemma 1.12, p. 20], [44, Corollary 2.20, p. 21]. Most notable examples are cross-polytopes as balls of the 1-norm and hypercubes as balls of the ∞ -norm, the latter being used for the definition of PL balls and spheres in [44, p. 8]. Having defined the notions of PL balls and spheres, Hudson [26, p. 20] characterises a PL manifold as a polyhedron such that any point has a neighbourhood that is a PL ball. This characterisation includes manifolds with boundaries, because the point may lie on the boundary of the ball. Thus, for characterising manifolds without boundary we would additionally require for each point a ball neighbourhood containing the point in the interior of the ball.

Since the characterisation so far allows arbitrary closed neighbourhoods, we would like to show that we can restrict our attention to cone neighbourhoods without loss of generality. Along the way we also show a possible characterisation in terms of links instead of stars. So let us assume a point p in a polyhedron P has a closed neighbourhood N PL-homeomorphic to a d -simplex Δ^d . Denote by $\phi: N \rightarrow \Delta^d$ the PL homeomorphism and let $\phi(p) = q$. Because N is a PL ball, there is some simplicial complex K with domain N such that ϕ is simplicial with respect to K and some subdivision of the simplex Δ^d . We show that the cone neighbourhood $\text{st}_K(p)$ of p in N with polyhedral link $\text{lk}_K(p)$ as base is a PL d -ball and a cone neighbourhood of p in P . Note that the simplex Δ^d itself is a star of q in Δ^d , with the corresponding link being the union of all proper faces of Δ^d not containing q . In particular, the link of q is a PL $(d-1)$ -sphere if q is in the interior of Δ^d , and it is a PL $(d-1)$ -ball if q is on the boundary of Δ^d . By PL invariance of links and stars, any cone neighbourhood of p in N is a PL d -ball, and the corresponding link is either a PL $(d-1)$ -sphere if p is in the interior of the ball N or a PL $(d-1)$ -ball if p is on the boundary of N . Because N is a neighbourhood of p in P , any cone neighbourhood of p in N is also a cone neighbourhood of p in P . Applying PL invariance of links and stars again, we see that not only $\text{st}_K(p)$, but actually all cone neighbourhoods of p in P are PL d -balls. Furthermore, any link of p in P is either a PL $(d-1)$ -sphere if p is interior to its ball neighbourhood, or a PL $(d-1)$ -ball if p is on the boundary of its ball neighbourhood. Therefore, having a PL $(d-1)$ -sphere or a PL $(d-1)$ -ball as link is necessary for having a PL d -ball as neighbourhood; and since the cone on a PL $(d-1)$ -sphere or a PL $(d-1)$ -ball is a PL d -ball [44, Proposition 2.23, p. 23], this condition is also sufficient.

Let us summarise the different criteria defining a piecewise linear manifold. For the summary and in the following, we denote by \mathbb{R}_+^d the half-space of \mathbb{R}^d consisting of all points $x = (x_1, \dots, x_d)$ whose last coordinate x_d is non-negative.

Definition 1.34 (PL manifold). A polyhedron P is called a *piecewise linear d -manifold without boundary* if for each point p in P one of the following equivalent statements holds:

- The point p has a neighbourhood in P that is PL-homeomorphic to an open set in \mathbb{R}^d .
- The point p has a neighbourhood in P that is a PL d -ball containing p in its interior.
- The point p has a cone neighbourhood in P that is a PL d -ball containing p in its interior.

- The point p has a link in P that is a PL $(d - 1)$ -sphere.

A *piecewise linear d -manifold with boundary* is a polyhedron P such that for each point p in P one of the following equivalent statements holds:

- The point has a neighbourhood in P that is PL-homeomorphic to an open set in \mathbb{R}_+^d .
- The point has a neighbourhood in P that is a PL d -ball.
- The point has a cone neighbourhood in P that is a PL d -ball.
- The point has a link in P that is a PL $(d - 1)$ -sphere or a PL $(d - 1)$ -ball.

The *boundary* of a manifold with boundary consists of those points whose link is a PL ball; they lie on the boundary of their ball neighbourhood. The remaining points, whose links are PL spheres and which are interior to their respective ball neighbourhoods, form the *interior* of the manifold.

Combinatorial Manifolds. Now we turn to the characterisation of combinatorial manifolds. Following Glaser [21, p. 18], we also start by defining balls and spheres:

Definition 1.35 (combinatorial ball, combinatorial sphere). A *combinatorial d -ball* is a simplicial complex whose domain is a PL d -ball. A *combinatorial d -sphere* is a simplicial complex whose domain is a PL d -sphere.

By the common abuse of notation that attributes properties of the domain of a complex to the complex itself, we occasionally call a complex a PL ball or sphere if it is a combinatorial ball or sphere. We denote a combinatorial or PL d -ball by B^d and a combinatorial or PL d -sphere by S^d .

Definition 1.36 (combinatorial manifold). A simplicial complex K is a *combinatorial d -manifold without boundary* if for each vertex v the link $\text{lk}_K(v)$ is a combinatorial $(d - 1)$ -sphere. It is a *combinatorial d -manifold with boundary* if each vertex link $\text{lk}_K(v)$ is a combinatorial $(d - 1)$ -sphere or a combinatorial $(d - 1)$ -ball.

By PL invariance of links and stars, it is clear that the domain of a simplicial complex K is a PL d -manifold if and only if $\text{lk}_K(x)$ is a combinatorial $(d - 1)$ -sphere or -ball for each point x in the domain, where balls only occur for manifolds with boundary. It is less obvious that checking the links of vertices suffices, so that a simplicial complex is a combinatorial manifold if and only if its domain is a PL manifold. This can be deduced starting from a property of combinatorial manifolds stated in [21, Theorem II.2, p. 19], which is used as the defining property of combinatorial manifolds in [26, p. 26]:

Fact 1.37. Let S be an r -simplex in a combinatorial d -manifold K . Then its link $\text{lk}_K(S)$ is a combinatorial $(d - r - 1)$ -sphere or -ball, with balls occurring only for manifolds with boundary.

Using for a point x in the interior of S the equality $\text{lk}_K(x) = \text{bd}(S) \text{lk}_K(S)$ from Table 1.3, we see that $\text{lk}_K(x)$ is the join of a combinatorial $(r - 1)$ -sphere and a combinatorial $(d - r - 1)$ -sphere or -ball. Hence, completing the proof sketch, $\text{lk}_K(x)$ is a combinatorial $(d - 1)$ -sphere or -ball because joins of balls and spheres obey the following rules [21, Proposition II.1, p. 18], [26, Lemma 1.13, p. 22], [44, Proposition 2.23, p. 23]:

Fact 1.38. For combinatorial as well as piecewise linear balls and spheres the following rules concerning their joins hold:

- The join of a d -ball and an r -ball is a $(d + r + 1)$ -ball.
- The join of a d -ball and an r -sphere is a $(d + r + 1)$ -ball.
- The join of a d -sphere and an r -sphere is a $(d + r + 1)$ -sphere.

Fact 1.37 implies that each $(d - 1)$ -simplex in a combinatorial d -manifold has either one or two d -simplices as proper co-faces. In manifolds without boundary only $(d - 1)$ -simplices with exactly two proper co-faces occur. Otherwise, the $(d - 1)$ -simplices with only one d -simplex as co-face and all their faces form the boundary of the manifold [21, p. 17]. It consists of the simplices whose link is a combinatorial ball [26, Lemma 1.18, p. 27], and it is a combinatorial $(d - 1)$ -manifold without boundary [21, Corollary II.4, p. 21].

Another important fact about PL balls and spheres is the following [44, Lemma 1.10, p. 8]:

Fact 1.39. Assume that for two PL balls B_1 and B_2 , there is a PL homeomorphism between their boundary spheres. Then the boundary homeomorphism extends to a PL homeomorphism between B_1 and B_2 by the cone construction.

The fact that polytopal complexes can be triangulated by starring at its vertices and PL invariance of links and stars shows that the domain of a polytopal complex is a PL manifold if and only if the link of every vertex is a PL sphere or ball. Therefore one might be tempted to extend the notions of combinatorial manifolds, spheres, and balls to polytopal complexes, because many properties would also hold for the extended notions. But an analogue of Fact 1.37 fails for polytopal cells. For this reason, when we occasionally call a polytopal complex a combinatorial manifold, sphere, or ball, we actually mean that some triangulation of it is a combinatorial manifold, sphere, or ball, which is independent of the choice of the triangulation.

1.2.10 Homotopy and Homology

Several different parts of the thesis deal with notions that are related to homotopy and homology. Both concepts are central for the broad field of algebraic topology. Here we only sketch the basic notions and state the definitions and results relevant for our study. For further information, we refer to textbooks on algebraic topology, for example Hatcher [25].

Homotopies. A *homotopy* is a continuous map $h: X \times I \rightarrow Y$ for topological spaces X and Y and some closed interval I . The interval is usually imagined to model a certain time interval, and is often assumed to be the unit interval $[0, 1]$. This point of view leads to a family of maps consisting of maps h_t for each point t in time interval I , where $h_t: X \rightarrow Y$ is defined by $h_t(x) = h(x, t)$ for each x in X . The homotopy is viewed as a continuous transition from the map h_a at the left endpoint of the interval $I = [a, b]$ to the map h_b at the right endpoint. The maps h_a and h_b are called *homotopic*. When for some subspace Z of X , the restrictions of all maps h_t to Z agree, the homotopy h is called a *homotopy relative to Z* .

Homotopy Equivalences and Deformation Retractions. Two topological spaces X and Y are *homotopy equivalent* if there are continuous maps $f: X \rightarrow Y$ and $g: Y \rightarrow X$ such that the compositions $f \circ g$ and $g \circ f$ are homotopic to the identities on X and Y respectively. The maps f and g are called homotopy equivalences. A special case of homotopy equivalences is induced by deformation retractions. A *retraction* from a topological space X to a subspace Z is a continuous map $f: X \rightarrow Z$ such that the restriction to that subspace is the identity on Z . A *deformation retraction* from X to Z as above is a homotopy from the identity on X to a retraction from X to Z ; it is a *strong deformation retraction* if it is a homotopy relative to Z . For a deformation retraction from X to Z , the spaces X and Z are homotopy equivalent with the retraction from X to Z homotopic to the identity on the one hand, and the inclusion map from Z to X on the other hand being the homotopy equivalences. We also say that X deformation retracts to Z or that Z is a deformation retract of X , if there is a deformation retraction from X to Z .

Collapses. Typical examples of deformation retracts arise from collapses. For simplicial complexes, the operation of an *elementary collapse* is defined as the removal of two distinct simplices $T \subseteq S$ from a complex K with the property that T and S are the only co-faces of T in K [21, p. 49]. The domain of the resulting complex is a strong deformation retract of $|K|$. The cell S in such an elementary collapse is necessarily a maximal element of the face poset, and T is called a *free face* of S in K . The operation of an elementary collapse can be defined more generally for a pair of cells in a CW complex. A detailed definition is given for example by Cohen in [14, p. 14]. Important for our purposes is only the implied fact that the definition for simplicial complexes works without change for polytopal complexes as well: If T is a proper face of S such that T and S are the only co-faces of T in some polytopal complex K , then we call the removal of S and T from K an elementary collapse, and the domain of the resulting polytopal complex $K \setminus \{S, T\}$ is a strong deformation retract of $|K|$. We sometimes say that a pair (S, T) fulfilling the requirements for an elementary collapse “can be collapsed” or “defines a valid elementary collapse”.

Applying the operation of an elementary collapse iteratively yields the notion of a collapse: Let $K_1 \supseteq K_2 \supseteq \dots \supseteq K_n$ be a sequence of subcomplexes such that K_{i+1} is obtained from K_i by an elementary collapse $K_{i+1} = K_i \setminus \{S_i, T_i\}$. Then we say that K_1 *collapses to* K_n and write $K_1 \searrow K_n$; the sequence $((S_1, T_1), (S_2, T_2), \dots, (S_{n-1}, T_{n-1}))$ is said to describe a sequence of valid elementary collapses. If K_1 collapses to K_n , then $|K_n|$ is a strong deformation retract of $|K_1|$. A complex that collapses to a single vertex is called *collapsible*.

Isotopies. An *isotopy* is a homotopy $h: X \times I \rightarrow Y$ such that each map h_t for $t \in I$ is a topological embedding. When X and Y are polyhedra, we could additionally require that the isotopy h is piecewise linear and each map h_t is a PL embedding, and call such an isotopy a *piecewise linear isotopy*. Another characterisation defines a PL isotopy as a level-preserving PL embedding $H: X \times I \rightarrow Y \times I$, where level-preserving means that there is a function $h: X \times I \rightarrow Y$ such that the equality $H(x, t) = (h(x, t), t)$ holds for all x in X and t in I , in other words, H leaves the I -part unchanged [44, Definitions 3.21, p. 37]. Hudson [26, p. 128] remarks that both characterisations are equiva-

lent. We will later define a special kind of PL isotopies called f -level-preserving PL isotopies. One could consider such f -level-preserving PL isotopies in terms of the second characterisation, but let us point out already here that being level-preserving in the sense of that characterisation and being f -level-preserving are distinct and independent requirements.

Collars and Bicollars. Other special cases of homotopies are collars and bicollars, which play only a minor role in this thesis, when we mention that certain level sets are bicollared with a bicollar given by an f -level-preserving PL isotopy. We use these notions in the following sense, based on the definitions by Rourke and Sanderson [44, p. 24] and by Rolfsen [40, p. 34]: Let $h: X \times I \rightarrow Y$ be a PL embedding for polyhedra $X \subseteq Y$ such that h_t is the inclusion map of X into Y for some t in the interval I ; then h , or sometimes also its image, is called a *collar* on X in Y if t is an endpoint of the interval, and it is called a *bicollar* if t is in the interior of the interval. The polyhedron X is said to be collared or bicollared in Y . By piecewise linear reparametrisation of the interval, one can assume defaults for the interval I and the value t at which h_t is the inclusion map, namely $I = [0, 1]$ and $t = 0$ for collars, and $I = [-1, 1]$ and $t = 0$ for bicollars. Rourke and Sanderson additionally require that the image of a collar $h: X \times [0, 1] \rightarrow Y$ is a neighbourhood of $X = h(X \times \{0\})$ in Y with $h(X \times [0, 1))$ being open in Y . The bicollars on level sets occurring here as images of f -level-preserving PL isotopies also have the property that they form a neighbourhood of the level set.

Homology Groups. The main application area of homology in this thesis are computational aspects because the homology of a simplicial complex is computable and a topological invariant of its domain. Note that we can apply simplicial homology also to polytopal complexes by triangulating them. Therefore we can avoid resorting to more general concepts of homology for the most part, such as cellular homology for CW-complexes, which include polytopal complexes, or singular homology for arbitrary topological spaces. For the definition of homology groups, different abelian groups can be used as coefficients. Common choices for the coefficients are first of all the integers \mathbb{Z} , the rationals \mathbb{Q} , and the field with two elements \mathbb{Z}_2 .

Homology groups can be defined based on a chain complex, a sequence of abelian groups C_n with homomorphisms $\partial_n: C_n \rightarrow C_{n-1}$ called boundary maps, such that the composite of two successive boundary maps is the trivial homomorphism $\partial_n \circ \partial_{n+1} = 0$, or equivalently, the image $\text{im}(\partial_{n+1})$ is contained in the kernel $\ker(\partial_n)$. The n -th homology group H_n of the chain complex is the quotient group $\ker(\partial_n)/\text{im}(\partial_{n+1})$.

In the case of simplicial homology, the group of n -chains $C_n(K)$ for a simplicial complex K is a direct sum of copies of the coefficient group, one copy for each n -simplex in K . The elements are usually written as formal sums $\sum_{\sigma} g_{\sigma} \sigma$ where σ ranges over the n -simplices in K and the g_{σ} are elements from the coefficient group. For coefficients from \mathbb{Z} , this means that $C_n(K; \mathbb{Z})$ is the free abelian group generated by the n -simplices of K . The chain groups $C_n(K)$ for usual homology are trivial for $n > \dim(K)$ or $n < 0$, but there is a variant called reduced homology where the chain complex differs at a single chain group in the sequence: In the so-called augmented chain complex of K the

chain group $C_{-1}(K)$ is not trivial but generated by the empty simplex, hence isomorphic to the coefficient group. The homology groups with respect to the augmented chain complex of K are called reduced homology groups and denoted by $\tilde{H}_n(K)$.

It remains to define the boundary maps ∂_n used in simplicial homology. It suffices to specify the results for the generating elements, that is an $(n-1)$ -chain $\partial_n(g\sigma)$ for each n -simplex σ and element g of the coefficient group. This is particularly simple for coefficients from \mathbb{Z}_2 : The boundary map maps an n -simplex $\sigma = 1 \cdot \sigma$ to the formal sum $\sum_{\tau \sqsubseteq \sigma, \dim(\tau)=n-1} \tau$ of its $(n-1)$ -faces. For other coefficient groups, the $(n-1)$ -chain $\partial_n(g\sigma)$ is the sum $\sum g_\tau \tau$ of the $(n-1)$ -faces τ of σ with coefficients $g_\tau = g$ or $g_\tau = -g$ depending on orientation, for details see Hatcher [25, pp. 102–105, 153], Kaczynski, Mischaikow, and Mrozek [28, pp. 382–384], or Zomorodian [48, pp. 31, 71]. For the definition of the boundary map ∂_0 in reduced homology, the empty simplex is considered a (-1) -face of a vertex, for the definition in usual homology, ∂_0 is the trivial map. This way, we obtain for each n an n -th homology group $H_n(K)$ of K and an n -th reduced homology group $\tilde{H}_n(K)$.

For arbitrary topological spaces X , singular homology groups $H_n(X)$ and $\tilde{H}_n(X)$ are defined based on the following chain complex: The chain group $C_n(X)$ is generated as direct sum from the set of continuous maps $f: \Delta^n \rightarrow X$ from the n -simplex Δ^n to X . For reduced homology, the empty map $f: \emptyset \rightarrow X$ generates $C_{-1}(X)$, for unreduced homology, $C_{-1}(X)$ is trivial. The boundary map is defined for a generating n -chain $f: \Delta^n \rightarrow X$ by taking the formula for the simplicial boundary of Δ^n and replacing each face of Δ^n in the formal sum by the restriction of f to that face. The simplicial homology groups, reduced and unreduced, of a simplicial complex are isomorphic to their correspondent singular homology groups of the domain [25, Theorem 2.27, p. 128]. We say that two topological spaces X and Y “have the same homology” if corresponding homology groups are isomorphic. Homotopy equivalent spaces have the same homology [25, Corollary 2.11, p. 111]. A topological space is *acyclic* if all its reduced homology groups are trivial [25, p. 142].

It is easy to switch between reduced and unreduced homology because of the following rules, here stated for the case of integer coefficients [25, p. 110]: For non-empty spaces X , the unreduced homology group $H_0(X)$ is isomorphic to the direct sum of \mathbb{Z} with the reduced homology group $\tilde{H}_0(X)$, and for all other $n \neq 0$, the unreduced homology group $H_n(X)$ is isomorphic to the reduced homology group $\tilde{H}_n(X)$; for the special case of the empty set, $H_n(\emptyset)$ and $\tilde{H}_n(\emptyset)$ are isomorphic for all $n \neq -1$, and for the remaining case $n = -1$, the unreduced homology group $H_{-1}(\emptyset) = 0$ is the trivial group, whereas the reduced homology group $\tilde{H}_{-1}(\emptyset)$ is isomorphic to \mathbb{Z} .

For a subspace Y of a topological space X we may define *relative homology groups* $H_n(X, Y)$ as follows [25, p. 115]: The chain groups $C_n(X, Y)$ are the quotient groups $C_n(X)/C_n(Y)$. Since the usual boundary map $\partial_n: C_n(X) \rightarrow C_{n-1}(X)$ maps $C_n(Y)$ to $C_{n-1}(Y)$, it induces naturally a well defined boundary map on the quotient groups.

Betti Numbers and Euler Characteristic. The n -th Betti number $\beta_n(X)$ of a topological space X is the rank of the n -th homology group $H_n(X)$ for integer coefficients. One can also define Betti numbers with respect to coefficients

from a field F as the dimension of the F -vector space $H_n(X; F)$. The Betti numbers with respect to the rational coefficients agree with the standard Betti numbers for integer coefficients [25, Corollary 3A.6, p. 266]. Moreover, we can define reduced Betti numbers $\tilde{\beta}_n(X)$ analogously as the rank of the n -th reduced homology group $\tilde{H}_n(X)$. The above results concerning the relation between reduced and unreduced homology groups imply $\beta_n(X) = \tilde{\beta}_n(X)$ for all n and X but the following exceptions: If $n = 0$ and $X \neq \emptyset$, then $\beta_0(X) = \tilde{\beta}_0(X) + 1$ holds, and if $n = -1$ and $X = \emptyset$, then $\beta_{-1}(\emptyset) = \tilde{\beta}_{-1}(\emptyset) - 1$ holds. If all but finitely many homology groups $H_n(X)$ are trivial and thus all but finitely many $\beta_n = 0$, the *Euler characteristic* $\chi(X)$ is defined as the alternating sum $\sum_n (-1)^n \beta_n(X)$ of the Betti numbers. For simplicial complexes or even finite CW-complexes, the Euler characteristic can also be computed by the alternating sum $\sum_{n \geq 0} (-1)^n c_n$ where c_n is the number of n -cells in the complex [25, Theorem 2.44, p. 146].

Homology Computation, Homology Spheres, and Novikov’s Undecidability Result for Sphere Recognition. As already mentioned, homology groups of simplicial complexes are computable for various coefficient groups, including integers, rationals, and finite fields. The Betti numbers can also be read off from the computed groups. The standard way to compute the homology is based on representing the boundary maps between the freely generated chain groups by matrices and transforming them into their Smith normal forms, as described by Kaczynski, Mischaikow, and Mrozek [28, Chapter 3, pp. 93–141] and Zomorodian [48, Section 7.3.1, p. 137–139].

A manifold that has the same homology as a d -sphere is commonly called a *homology d -sphere*. We usually restrict our attention to homology spheres that are combinatorial manifolds. The notion of a *homology ball* is defined and used analogously. A result of Novikov implies that there is no algorithm for distinguishing combinatorial d -spheres from other homology d -spheres in dimension $d \geq 5$. Novikov’s proof of his “theorem on the impossibility of recognizing whether an n -dimensional manifold is an n -dimensional sphere ($n \geq 5$) or whether a contractible domain in $(n + 1)$ -dimensional Euclidean space with smooth boundary is an $(n + 1)$ -dimensional disc” is included in a paper by Volodin, Kuznetsov, and Fomenko [47]. An exposition of Novikov’s proof adapted to a piecewise linear setting is given by Chernavsky and Leksine [12]. Another variant of the result is stated by Matoušek, Tancer, and Wagner [32, Theorem 3.1]: For each $d \geq 5$, there is a computable sequence of combinatorial d -manifolds K_i , each a homology sphere, such that no algorithm exists deciding for each K_i whether it is a combinatorial sphere or not. For future reference, we restate the result as follows:

Fact 1.40 (Novikov). For $d \geq 5$, it is algorithmically undecidable whether a given simplicial complex is a combinatorial d -sphere or not. The problem remains undecidable if the input complexes are restricted to homology d -spheres.

1.2.11 Pachner Moves and Recognition of 3-Spheres

Sphere recognition in dimensions greater than four is algorithmically undecidable by Novikov’s result just stated above. In this subsection we review known

facts on sphere recognition in small dimensions with a focus on the 3-dimensional case.

Sphere Recognition and Ball Recognition. Note first that sphere recognition is computationally equivalent to ball recognition in the sense that both problems can be reduced to each other. This observation is based on the following two facts. Firstly, for a d -ball contained in a d -sphere, the closure of its complement is again a d -ball [44, Corollary 3.13, p. 35]. Secondly, a d -ball with a cone on its boundary attached is a d -sphere, which can easily be seen by looking at the boundary of the $(d + 1)$ -dimensional cross-polytope subdivided into two hemispheres by the equator: Both hemispheres are balls and cones on their common boundary, hence the whole sphere is one hemisphere with a cone attached to its boundary. A homeomorphism from a ball to the hemisphere extends by the cone construction to a homeomorphism from the ball with a cone on its boundary to the whole sphere. Therefore we can reduce ball recognition to sphere recognition by attaching a cone to the boundary of the given complex. The resulting complex is a combinatorial d -sphere if and only if the original complex is a combinatorial d -ball. On the other hand, d -sphere recognition can be reduced to d -ball recognition by removing a single, arbitrary d -simplex from the input complex. The resulting complex is a combinatorial d -ball if and only if the input complex is a combinatorial d -sphere.

Before we turn to the recognition of 3-spheres, let us settle matters with the other remaining dimensions. A combinatorial 0-sphere consists of exactly two vertices. Only slightly less trivial is the recognition of 1-spheres: In the language of graphs, a combinatorial 1-sphere is a non-empty connected graph in which every vertex has degree two. For the recognition of combinatorial 2-spheres, one starts with checking whether the link of each vertex is a 1-sphere and whether the 1-skeleton is connected. If both tests succeed, the considered simplicial complex is a connected compact combinatorial 2-manifold without boundary. By the classification of compact surfaces, such a manifold is a combinatorial 2-sphere if and only if its Euler characteristic is two. The question whether combinatorial 4-spheres are algorithmically recognisable or not is still open [24].

The first correct algorithm for 3-sphere recognition was proposed by Rubinstein [45]. Thompson described a variant of the algorithm with a different point of view on the correctness proof [46]. Based on the results leading to this algorithm, Mijatović showed that a conceptually simple algorithm using Pachner moves allows to recognise 3-spheres as well [34]. We summarise the facts important for us from Mijatović's paper [34], but note that the treatment there allows for non-combinatorial abstract simplicial complexes, that is, the same set of vertices may span several distinct simplices. For example, one obtains a triangulation of a d -sphere by gluing together two copies of a d -simplex by identifying corresponding boundary cells. Mijatović remarks that in dimension 3, any triangulation can be made combinatorial by taking the second derived subdivision, which can also be obtained by suitable Pachner moves.

Pachner Moves. For the definition of a d -dimensional Pachner move, consider the boundary of a $(d + 1)$ -simplex Δ^{d+1} consisting of $d + 2$ d -simplices. Each collection of k such d -simplices for $1 \leq k \leq d + 1$ defines a combinatorial d -ball, and the remaining $d + 2 - k$ d -simplices form another combinatorial

d -ball with the same boundary. Up to combinatorial equivalence the two triangulations depend only on k . A Pachner move is a switch between two such triangulations of a d -ball with the same boundary as follows. Denote by B_k^d the subcomplex of $\text{bd}(\Delta^{d+1})$ with k d -simplices and by \bar{B}_k^d the opposite subcomplex containing the other $d+2-k$ d -simplices. By Fact 1.39, the identity on the common boundary $B_k^d \cap \bar{B}_k^d$ extends to a homeomorphism between the balls B_k^d and \bar{B}_k^d . Therefore, if a combinatorial d -manifold K contains B_k^d as a subcomplex, up to combinatorial equivalence of course, we can replace the simplices from $B_k^d \setminus B_k^d \cap \bar{B}_k^d$ by the simplices from $\bar{B}_k^d \setminus B_k^d \cap \bar{B}_k^d$ without changing the topology of K . In other words, the triangulation in the interior of the ball B_k^d is replaced by the interior triangulation of the ball \bar{B}_k^d with the boundary left unchanged. Such a modification is called a $(k - (d - k + 2))$ Pachner move. It can be reversed by the corresponding $((d - k + 2) - k)$ Pachner move. The simplest example of a Pachner move is a $(1 - (d + 1))$ move: It amounts to an elementary starring at a point in the interior of a d -simplex, which subdivides the d -simplex into $(d + 1)$ smaller d -simplices.

Mijatović's Algorithm for 3-Sphere Recognition The key facts making Mijatović's algorithm possible are the following:

Fact 1.41 ([34, Theorem 1.1]). Two combinatorial d -manifolds K and L are PL-homeomorphic if and only if there is a finite sequence of d -dimensional Pachner moves transforming K into L up to combinatorial equivalence.

Fact 1.42 ([34, Theorem 1.2]). Let S be the the 3-sphere triangulation formed by two 3-simplices with identified boundaries. Any triangulation of a 3-sphere with t 3-simplices can be transformed into S by at most $6 \cdot 10^6 t^2 2^{5 \cdot 10^4 t^2}$ Pachner moves.

Fact 1.43 ([34, Proposition 1.3]). If there is a computable function $f: \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ such that for any two triangulations K and L of a compact PL manifold M , with K containing at most k and L containing at most l maximal simplices, the number of Pachner moves needed to transform K into L is bounded by $f(k, l)$, then the recognition problem for the manifold M is algorithmically decidable.

The construction of an algorithm proving the last fact is simple: Fix a triangulation K of M with k maximal simplices and let L be the input complex of the algorithm with l maximal simplices. Check for all possible sequences of Pachner moves starting from L with length at most $f(k, l)$ whether the resulting complex is combinatorially equivalent to K . Since the second fact provides a computable bound for the case that M is the 3-sphere, such an algorithm allows to recognise combinatorial 3-spheres.

The argument used for Fact 1.43 can be used in many other variants. We sketch a general proof template that can be applied to obtain other similar results. Let \mathcal{C} be a class of pairs of simplicial complexes. Suppose a computable function $f: \mathcal{C} \rightarrow \mathbb{N}$ exists such that whenever a pair (K, L) in \mathcal{C} is a pair of PL-homeomorphic manifolds, then the number of Pachner moves needed to transform K into L is bounded by $f(K, L)$. Then the problem whether a pair of manifolds from class \mathcal{C} is PL-homeomorphic is algorithmically decidable. Note that for bounded dimension, a function f bounding the number of necessary Pachner moves in terms of the size of the complexes, measured in number of maximal simplices or total number of simplices, always exists: Up to combinatorial equivalence, there are only finitely many simplicial complexes of bounded

dimension of a certain size, so there are only finitely many homeomorphic manifold pairs. For each homeomorphic manifold pair, there is a finite sequence of Pachner moves by Fact 1.41, hence for each given input size, the value of f can be chosen as the maximum of a finite set. By contraposition, the above observation implies that if the problem whether a manifold pair is homeomorphic is algorithmically undecidable for class \mathcal{C} , then this function f or any other function bounding the number of necessary Pachner moves for that class is not computable.

1.2.12 Schoenflies Conjecture

In Chapter 4 we will develop and study several equivalent characterisations of regular points in PL Morse theory. One of the characterisations for a point x in a combinatorial d -manifold requires that the intersection $|\text{lk}(x)| \cap f^{-1}(f(x))$ of $|\text{lk}(x)|$ with the level set at $f(x)$ is embedded in $|\text{lk}(x)|$ in the same way as the standard $(d-2)$ -sphere is embedded in the standard $(d-1)$ -sphere. Since the Schoenflies conjecture, in a nutshell, states that any embedding of a $(d-2)$ -sphere into the $(d-1)$ -sphere is equivalent to the standard one, the status of the conjecture has influence on the characterisation of regular vertices as well. We give a short survey on the notions and results that are interesting for us, following the treatment in [44].

Definitions. At first, we list some definitions for pairs [44, p. 50]: For a PL manifold Q of dimension k contained in a PL manifold P of dimension d , the pair (P, Q) is called a (d, k) -manifold pair. The pair is a *sphere pair* if both manifolds are PL spheres, and it is a *ball pair* if both manifolds are PL balls. The standard ball and sphere pairs are given by the hypercubes formed by the ∞ -norm unit balls and spheres and their restriction to a subspace: The *standard (d, k) -ball pair* is $([-1, 1]^d, [-1, 1]^k \times \{0\}^{d-k})$, the *standard (d, k) -sphere pair* is $(\text{bd}([-1, 1]^{d+1}), \text{bd}([-1, 1]^{k+1}) \times \{0\}^{d-k})$. Two manifold pairs (P, Q) and (P', Q') are said to be homeomorphic as pairs if there is a PL homeomorphism $h: P \rightarrow P'$ such that $h(Q) = Q'$. A ball or sphere pair is called *unknotted* if it is homeomorphic as a pair with the standard ball or sphere pair. A manifold pair (P, Q) is *proper* if $\text{bd}(Q) = Q \cap \text{bd}(P)$. A proper (d, k) -manifold pair (P, Q) is called *locally flat* if each point in Q has a neighbourhood N in P such that the pair $(N, N \cap Q)$ is homeomorphic as a pair with the standard (d, k) -ball pair. It is common to say that Q is locally flat in P or that Q is unknotted in P if the pair (P, Q) has that property. Similarly, if a locally flat or unknotted pair (P, Q) arises from a PL embedding $f: M \rightarrow P$ with image Q , then the embedding f of M into P is also called locally flat or unknotted respectively.

Clearly, we can apply all these definitions to combinatorial manifolds as well. Note that if the domain of a combinatorial manifold L is contained in the domain of another compact manifold K , then there are subdivisions K' and L' of K and L such that L' is a subcomplex of K' . This is proven by the following fact [44, Addendum 2.12, p. 16]: For polytopal complexes K and L_1, \dots, L_n such that $|L_i| \subseteq |K|$ for each i , there are simplicial subdivisions K' of K and L'_i of L_i such that L'_i is a subcomplex of K' for each i . Since compact polyhedra can be triangulated, we can find for any manifold pair (P, Q) of compact manifolds, in particular for ball and sphere pairs, a combinatorial manifold pair (K, L) such that $|K| = P$, $|L| = Q$, and L is a subcomplex of K .

The Conjecture and Its Status. Now the d -dimensional Schoenflies conjecture for the PL category can be stated as follows:

Conjecture 1.44 (Schoenflies conjecture). *Every $(d, d - 1)$ -sphere pair is unknotted.*

The conjecture is known to be true for $d \leq 3$, and it is still open for $d \geq 4$ [44, p. 47]. It is either true for all dimensions or false for all dimensions $d \geq 4$, see Luft [31]. A weaker positive result is the following [31, Theorem 5]: If S^{d-1} is a combinatorial $(d - 1)$ -sphere contained in a combinatorial d -sphere S^d , then S^{d-1} divides S^d into two complexes B_1^d and B_2^d with $B_1^d \cup B_2^d = S^d$ and $B_1^d \cap B_2^d = S^{d-1}$ and with the decomposition being different from the trivial one with this property, namely the decomposition into S^d and S^{d-1} . This decomposition is unique [31, Theorem 3], and the components B_1^d and B_2^d are both topological d -balls. Furthermore, if one of them is a combinatorial ball, then the other component is a combinatorial ball as well [31, Theorem 2], see also [44, Corollary 3.13, p. 35]. Using the cone construction, one can prove that both components being balls is necessary and sufficient for the pair being unknotted. Locally flat $(d, d - 1)$ -sphere pairs are known to be unknotted for all dimensions except $d = 4$, where the Schoenflies conjecture is still open even if local flatness is already assumed [44, p. 47].

1.3 Discrete Morse Theory

A completely different approach to translate the ideas of Morse theory to simplicial complexes called discrete Morse theory was developed by Forman [18, 19]. Discrete Morse theory studies functions that assign values to the cells of a complex, in contrast to the piecewise linear setting, where the domain of the considered function is the domain of the complex. Furthermore, discrete Morse functions can be defined for arbitrary finite CW complexes, whereas the definition of PL Morse functions used here requires that the considered complex is a combinatorial manifold. In spite of the differences, there is a result linking discrete Morse functions on combinatorial manifolds to piecewise linear Morse functions on the derived subdivision. Therefore we give a short introduction to the basic notions. We omit the subtleties necessary for CW complexes and focus on discrete Morse theory for polytopal complexes.

Definitions. Forman defines a discrete Morse function on a polytopal complex K as follows [18, Definition 2.1]:

Definition 1.45 (discrete Morse function). A function $f: K \rightarrow \mathbb{R}$ defined on the cells of a polytopal complex K is a *discrete Morse function* if for each d -cell S in K the number of $(d - 1)$ -faces $T \sqsubseteq S$ with $f(T) \geq f(S)$ is at most one, and the number of $(d + 1)$ -co-faces $U \supseteq S$ with $f(U) \leq f(S)$ is at most one.

In discrete Morse theory, the empty cell is usually not considered to belong to the domain of a discrete Morse function. The definition requires that each cell S has at most one exceptional immediate face with greater or equal f -value, and at most one exceptional immediate co-face with smaller or equal f -value. A basic result is that for each cell S at most one of the two exceptional cases

occurs [18, Lemma 2.4]. The presence or absence of exceptional faces or co-faces is used to distinguish between critical and non-critical cells:

Definition 1.46 (critical cell). A cell S in K is a *critical cell* of a discrete Morse function $f: K \rightarrow \mathbb{R}$ if S has no exceptional faces or co-faces, that is all immediate faces have strictly smaller f -value, and all immediate co-faces have strictly larger f -value.

The non-critical cells of a discrete Morse function always appear as immediate face/co-face pairs $T \sqsubseteq S$ with $f(T) \geq f(S)$. These pairs can be viewed as a partial matching in the Hasse diagram of the face poset: The *Hasse diagram* of a finite poset is a graph whose vertices are the elements of the poset and whose edges are given by the cover relation, that is, two elements x and y are connected by an edge if $x < y$ or $y < x$ holds and there is no third element z with $x < z < y$ or $y < z < x$. In our case, the cover relation of the face poset is the immediate face relation. Since the empty cell is neglected in discrete Morse theory, the Hasse diagrams are also considered without a vertex for the empty cell. Chari [11] showed that matchings in the Hasse diagram coming from a discrete Morse function can be characterised as acyclic matchings in the following sense: Consider the face poset of a polytopal complex with the empty cell removed and a matching in the corresponding Hasse diagram without empty cell. The matching is called *cyclic* if there is a non-trivial simple cycle in the Hasse diagram such that the vertices are alternatingly face and co-face of its predecessors and every other edge belongs to the matching. In other words, there is a sequence of distinct cells $(S_1, S_2, \dots, S_{2n-1}, S_{2n})$ with $n > 1$ such that $S_{2k-1} \sqsubseteq S_{2k}$ belongs to the matching for each k with $1 \leq k \leq n$, the cell S_{2k+1} is an immediate face of S_{2k} for each k with $1 \leq k \leq n-1$, and the cell S_{2k} is an immediate face of S_{2k-1} . Otherwise the matching is called *acyclic*. With these notions, the following holds [11, Proposition 3.3]:

Fact 1.47. For any discrete Morse function, the induced matching of the non-critical cells in the Hasse diagram is acyclic. For any acyclic matching in the Hasse diagram, there is a discrete Morse function inducing that matching.

For this reason acyclic matchings are also called *Morse matchings*.

From a given Morse matching, it is not hard to construct discrete Morse functions inducing that matching that fulfil even stronger requirements than demanded by Forman's definition. A popular variant for the definition of discrete Morse functions is as follows, see for example [1]: A discrete Morse function on a polytopal complex K is a function $f: K \rightarrow \mathbb{R}$ such that $T \sqsubseteq S$ implies $f(T) \leq f(S)$, the preimage of each real number contains at most two cells and if the preimage contains two cells $S \neq T$, then one is the immediate face of the other. In particular, under this definition a discrete Morse function is monotone with respect to the poset ordering on K and the usual ordering on \mathbb{R} ; the non-critical cells are immediate face/co-face pairs $T \sqsubseteq S$ where monotonicity is not strict, that is $f(T) = f(S)$. Discrete Morse functions inducing the same Morse matching are often considered as equivalent [1]; so every discrete Morse function in Forman's sense is equivalent to some monotone discrete Morse function in the sense just described.

For our approach in Chapter 3 leading to a link from discrete Morse functions to piecewise linear Morse functions on the derived subdivision, we need slightly different additional assumptions on a discrete Morse function $f: K \rightarrow \mathbb{R}$. The

assumptions are: All values $f(S)$ for $S \in K$ are distinct; the pairs of the induced Morse matching are the only exceptions to monotonicity, that is, $T \sqsubseteq S$ implies $f(T) \leq f(S)$ unless the pair (T, S) belongs to the induced Morse matching, in which case $f(T) > f(S)$ holds. A discrete Morse function f' fulfilling these assumptions can be constructed from an equivalent monotone discrete Morse function f as above by perturbing the values of the matched pairs. For example, setting $f'(S) = f(S)$ if S is critical or the larger cell of a matched pair $T \sqsubseteq S$ with $f(T) = f(S)$ and setting $f'(T) = f(T) + \epsilon$ for sufficiently small ϵ if T is the smaller cell of a matched pair yields such a discrete Morse function f' .

Results. Forman shows that with the notions of a discrete Morse function and its critical cells results analogous to those of smooth Morse theory can be deduced. The role of the index of a critical point is taken over by the dimension of a critical cell. Instead of a level set $M_{\leq h}$, one considers a level subcomplex $K(h)$ consisting of all cells whose f -value is at most h and all their faces. There are results relating the change in homotopy type between $K(a)$ and $K(b)$ for $a < b$ to the critical cells with f -value in the interval $(a, b]$. The first result is an analogue to the last statement of Fact 1.2 from smooth Morse theory.

Fact 1.48 ([18, Lemma 2.6]). Let f be a discrete Morse function on a polytopal complex K . If there are no critical cells with f -value in the interval $(a, b]$, then the level subcomplex $K(b)$ collapses to the level subcomplex $K(a)$.

Then there is an analogue to Fact 1.4:

Fact 1.49 ([18, Lemma 2.7]). Assume that for some discrete Morse function $f: K \rightarrow \mathbb{R}$ the preimage $f^{-1}(a, b]$ contains exactly one critical cell. Let λ be the dimension of that cell. Then $K(b)$ is homotopy equivalent to $K(a)$ with a λ -cell attached.

Based on this, one obtains the main theorem of discrete Morse theory, which is an analogue of Fact 1.5 for smooth Morse theory.

Fact 1.50 ([18, Theorem 2.5]). Let f be a discrete Morse function on polytopal complex K . Then K is homotopy equivalent to a CW-complex whose cells are in bijection with the critical cells of f : one λ -cell for each critical cell of dimension λ .

There are also discrete versions of the Morse inequalities (Fact 1.6) and the Reeb theorem (Fact 1.7).

Fact 1.51 (Morse inequalities [18, Theorem 2.11]). Let K be a polytopal complex with a discrete Morse function f . Denote by c_λ the number of critical cells of dimension λ of f , and by β_λ the λ -th Betti number of $|K|$. Then the following inequalities and equalities hold:

1. The Weak Morse Inequalities. For each λ :

$$\beta_\lambda \leq c_\lambda$$

2. An equality representing the Euler characteristic $\chi(K)$ of K in two ways:

$$\chi(K) = \sum_{\lambda} (-1)^\lambda c_\lambda = \sum_{\lambda} (-1)^\lambda \beta_\lambda$$

3. The Strong Morse Inequalities. For each λ :

$$\sum_{i=0}^{\lambda} (-1)^i \beta_{\lambda-i} \leq \sum_{i=0}^{\lambda} (-1)^i c_{\lambda-i}$$

Fact 1.52 (Reeb theorem [18, Theorem 8.2]). If a combinatorial manifold has a discrete Morse function with exactly two critical cells, then it is a combinatorial sphere.

Chapter 2

Level Set Complexes

The study of level sets is an important tool in smooth Morse theory. The classical results collected in Section 1.1 contain statements involving various kinds of level sets of a Morse function $f: M \rightarrow \mathbb{R}$: the preimage $M_{\leq h} = f^{-1}(-\infty, h]$ of a closed interval $(-\infty, h]$, half-bounded to the right, the preimage $f^{-1}(a)$ of a single value $a \in \mathbb{R}$, and the preimage $f^{-1}[a, b]$ of a closed bounded interval $[a, b]$. When we consider analogous preimages for a function $f: |M| \rightarrow \mathbb{R}$ linear on cells of a combinatorial manifold M , it turns out that they are the domain of a polytopal complex naturally induced by the cells of M . These complexes are typical examples of the level set complexes defined and studied in this chapter.

The results of the chapter are mainly quite basic observations on the combinatorial structure of the level set complexes. The observations serve as a toolbox for many proofs in the subsequent chapters, so that we can avoid repeating similar simple arguments that distract from the actual proof ideas. We develop the toolbox for a more general setting than we actually need later on. In this general setting we consider the domains of two polyhedral complexes, the first serves as the domain of the function f linear on cells, the second as the “level range”, that is the set whose preimage defines the desired level set. This allows for more uniform and simple formulations of the results than restricting the focus to real-valued functions on combinatorial manifolds.

The first section of the chapter defines level set complexes in Definition 2.6 and starts with first simple observations on how their cells can be represented by the cells of the domain complex and the cells of the level range complex. The notion of a canonical representation by such pairs introduced in the next section by Definition 2.8 permits refined results on the combinatorial structure of the level set complex. In fact, Lemma 2.10 yields the fundamental observation that the combinatorial structure of the original complexes and the canonical representations determine the combinatorial structure of the level set complex completely. A convenient alternative characterisation of canonical representations is given in Lemma 2.15. Finally, we return to the special case of real-valued functions: The last section provides results regarding combinatorial equivalence of certain level set complexes and possible collapses in level set complexes. For example, the absence of combinatorial changes in level set complexes when the varying level does not pass a value attained at a vertex is expressed by combinatorial equivalences in Lemmata 2.35 and 2.36. The collapse described in Lemma 2.46 shows that a level set $f^{-1}(-\infty, h]$ of a polytopal complex M is

homotopy equivalent to the subcomplex of M consisting of those cells that are completely contained in the level set.

2.1 Definitions and Basic Observations

The main objective of this thesis is to translate classical smooth Morse theory to a piecewise linear setting with combinatorial manifolds and functions linear on cells of them. Therefore it is unavoidable to study the properties of level sets in this setting. For analogues of the results from Section 1.1, we need to consider the following types of level sets for a function f linear on cells of a combinatorial manifold M : The set of points in $|M|$ whose f -value is at most a given real value h , the set of points whose f -value is exactly a given value a , and the set of points whose f -value lies in the range of some interval $[a, b]$. Level sets of the first type are often called a *lower level set* at level h or a *sublevel set* of h . The dual concept to such sets are *upper level sets* which consist of points whose f -value is at least a given value. Switching from f to its negative $-f$ turns upper level sets into lower level sets and vice versa.

The common feature of the level sets described above is that they are preimages of closed intervals. Since f is linear on cells, these preimages can be obtained from the combinatorial manifold M by intersecting each cell S of M with suitable hyperplanes or half spaces. These intersections are convex polytopes that coincide only in common faces. Therefore, the level sets can be viewed as polytopal complexes. In the following, we work out the details of this point of view. We gain more insight by extending some of the ideas to a more general situation: The combinatorial manifold M can be replaced by an arbitrary polyhedral complex, the map $f: |M| \rightarrow \mathbb{R}^r$ is still linear on cells, and as level sets we can allow any set of points in $|M|$ that is the preimage under f of the domain of some polyhedral complex embedded in \mathbb{R}^r .

2.1.1 Notation

Before we start, some remarks on the notation used for level sets seem to be appropriate. From an abstract point of view, we can use the term *level set* for any subset of the domain D of a function $f: D \rightarrow C$ that is the preimage $f^{-1}(R)$ of some subset R of the codomain C of f . We call the corresponding subset R of the codomain the *level range*. This way, the characterisation of a level set can be rephrased by saying that a level set is a subset of the domain consisting of all the points that are mapped to a given level range in the codomain. In this abstract setting, we denote the level set $f^{-1}(R) \subseteq D$ by $D_{f \in R}$. When the function f is clear from the context, we omit it and write simply $D_{\in R}$.

From this general notation, we derive several variants for special level ranges, namely intervals, and additional structure on the level set as a complex. All these variants have in common as a basic scheme, that the level set is denoted by the domain followed by a subscript expression indicating the level range. Occasionally, the considered function is included in the subscript, but often it will be omitted since only a single function is considered.

The derived expressions for describing level set complexes adhere to the pattern illustrated by the following example: Let $I \subseteq \mathbb{R}$ be a closed interval, let M be a polyhedral complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells, and let

Table 2.1: A list of abbreviated notations for level sets using the domain of a polyhedral complex M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells as example. The first line represents the standard notation that can be used not only for intervals, but also for arbitrary subsets $I \subseteq \mathbb{R}$. The other lines define abbreviations for special intervals typically occurring as level ranges: A real number h defines level sets consisting of all points with f -value strictly smaller than, strictly larger than, equal to, at most, or at least the given value h . In each equation, the first term is the abbreviated expression without function symbol, the second term is the level set expression including the function symbol, the third term describes the level set in set-builder notation, and the last term represents the level set as a preimage under f of the level range.

$$\begin{aligned}
|M|_{\in I} &= |M|_{f \in I} = \{x \in |M| \mid f(x) \in I\} = f^{-1}(I) \\
|M|_{< h} &= |M|_{f < h} = \{x \in |M| \mid f(x) < h\} = f^{-1}(-\infty, h) \\
|M|_{> h} &= |M|_{f > h} = \{x \in |M| \mid f(x) > h\} = f^{-1}(h, \infty) \\
|M|_{= h} &= |M|_{f = h} = \{x \in |M| \mid f(x) = h\} = f^{-1}(h) \\
|M|_{\leq h} &= |M|_{f \leq h} = \{x \in |M| \mid f(x) \leq h\} = f^{-1}(-\infty, h] \\
|M|_{\geq h} &= |M|_{f \geq h} = \{x \in |M| \mid f(x) \geq h\} = f^{-1}[h, \infty)
\end{aligned}$$

S be a cell of the complex M . Then for the domain $|M|$, the notation $|M|_{f \in I}$ denotes the level set $f^{-1}(I) = \{x \in |M| \mid f(x) \in I\}$, which is the underlying set of the level set complex $M_{f \in I}$, and $S_{f \in I} = \{x \in S \mid f(x) \in I\}$ denotes a cell of this complex.

Ideally, the expression for a level set complex would consist of complexes for domain and range, the expression for a cell of a level set complex would consist of cells for domain and range, and the expression for the underlying set would consist of the underlying sets for domain and range. But we allow the usual abuse of notation for these level set expressions as well, in particular for the level range expression, that is, identifying complexes with their domain or cells with the complex formed by their faces. Otherwise only $M_{f \in \mathcal{F}(I)}$ would be a formally correct expression for the aforementioned level set complex.

One instance of the special notation for intervals occurred already in Section 1.1 on classical Morse theory: For the lower level set at level h of a smooth manifold with Morse function f we used the term $M_{\leq h}$. Table 2.1 lists possible subscript expressions, using as domain the underlying set $|M|$ of a polyhedral complex M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. We consider such alternative subscript expressions as abbreviations for the corresponding standard expressions, for example $|M|_{f \leq h}$ is a shortened term for $|M|_{f \in (-\infty, h]}$.

When we consider a level set $|M|_{\in I}$ for a polyhedral complex M as above and choose a closed interval I as level range, the level set can be viewed as the domain of a polyhedral complex whose combinatorial structure is induced by the complex M , as we will see in the following. This complex is a typical example of a level set complex, and we denote it by $M_{\in I}$. Apart from the trivial cases $I = \mathbb{R}$ and $I = \emptyset$, a closed interval I is one of $(-\infty, h]$, $[h, \infty)$, $[a, b]$, or $\{h\} = [h, h]$ for some real values h , a , and b . Table 2.2 lists the abbreviated notations for level set complexes associated with these intervals, which are the

Table 2.2: Notation for level set complexes typically considered in Morse theory. It can be used more generally for arbitrary polyhedral complexes M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Listed are the level set complexes whose domains consist of all the points whose f -value is, in that order, equal to, at most, or at least a given real number h , or is contained in the closed interval between two real numbers a and b . In each row, the first term is the notation without function symbol, the second term the notation including function symbol.

$$\begin{aligned} M_{=h} &= M_{f=h} && \text{denotes the complex induced on } |M|_{f=h} \\ M_{\leq h} &= M_{f \leq h} && \text{denotes the complex induced on } |M|_{f \leq h} \\ M_{\geq h} &= M_{f \geq h} && \text{denotes the complex induced on } |M|_{f \geq h} \\ M_{\in[a,b]} &= M_{f \in [a,b]} && \text{denotes the complex induced on } |M|_{f \in [a,b]} \end{aligned}$$

level ranges commonly considered in Morse theory.

2.1.2 Level Set Cells

For the remainder of this chapter, I denotes a closed interval and M a polyhedral complex with a map $f: |M| \rightarrow \mathbb{R}^r$ linear on cells unless explicitly stated otherwise. In the following examinations, we usually begin with closed intervals as level ranges for real valued maps f as motivating example, and extend the observations to more general maps and level ranges afterwards. As mentioned above, the cells $S_{\in I}$ for S a cell of M belong to the level set complex $M_{\in I}$. Therefore we start our study of level set complexes by examining such level set cells.

Definition and Examples. In its most general form, a *level set cell* is the preimage of a polyhedral set R in some \mathbb{R}^r as level range under a linear map g defined on a polyhedral set S in some \mathbb{R}^d . Usually R and S are cells of polyhedral complexes and g is the restriction to S of a map f linear on cells of the complex containing S . For the simple case that S is an affine space with a non-constant linear map $g: S \rightarrow \mathbb{R}$, it is a basic fact from linear algebra, that the preimage of a single value $h \in \mathbb{R}$, in other words $S_{=h}$, forms a hyperplane in S . Furthermore $S_{\leq h}$ and $S_{\geq h}$ are the closed half-spaces separated by this hyperplane. For an arbitrary polyhedral set S with linear function g , we can always think of g as the restriction of a linear function defined on the affine hull of S . Since this map is uniquely determined by g , we do not strictly distinguish between linear maps defined on S and linear maps defined on its affine hull. This way, if g is a linear map on the affine hull of S , then $S_{=h}$ is the intersection of S with the hyperplane $g^{-1}(h)$, and $S_{\leq h}$ and $S_{\geq h}$ are the intersections of S with the respective closed half-spaces defined by that hyperplane.

In the usual setting, the considered polyhedral set S is a cell of M . Then the restriction $f|_S$ of f to S is a linear map, and the unique linear extension of $f|_S$ to the affine hull of S will be denoted by f_S . In order to use the same notation for the abstract study of level set cells, we assume more generally that S is a polyhedral set and $f_S: \text{aff}(S) \rightarrow \mathbb{R}^r$ is an affine linear function defined on the affine hull of S . When we focus on a single cell, it is convenient to restrict

Table 2.3: Faces of a level set cell $S_{\in I}$, where I belongs to one of the non-trivial types of closed intervals $(-\infty, h]$, $[h, \infty)$, $\{h\}$, and $[a, b]$ for real numbers h , a , and b .

Faces of $S_{\leq h}$:	$T_{\leq h}$ and $T_{=h}$	for T a face of S
Faces of $S_{\geq h}$:	$T_{\geq h}$ and $T_{=h}$	for T a face of S
Faces of $S_{=h}$:	$T_{=h}$	for T a face of S
Faces of $S_{\in [a,b]}$:	$T_{\in [a,b]}$, $T_{=a}$, and $T_{=b}$	for T a face of S

our attention to the affine hull of the cell and forget about its embedding in a possibly higher dimensional ambient space. In particular, when we say that certain preimages of f_S are hyperplanes or half-spaces, as in the example above with the map g , these notions are used with respect to the affine hull of S and not with respect to the ambient space.

Under these assumptions, we can represent any level set of S as the intersection of S with a preimage of f_S : For any level range $R \subseteq \mathbb{R}^r$, the equality $S_{\in R} = S \cap f_S^{-1}(R)$ holds. Let us resume the example outlined above and consider a real-valued, non-constant map f_S and the typical level ranges formed by closed intervals. The lower and upper level sets $S_{\leq h}$ and $S_{\geq h}$ are intersections of S with half-spaces, namely $S_{\leq h} = S \cap \{x \mid f_S(x) \leq h\}$ and $S_{\geq h} = S \cap \{x \mid f_S(x) \geq h\}$. The level set $S_{\in [a,b]}$ is the intersection of S with the two half-spaces $\{x \mid f_S(x) \geq a\}$ and $\{x \mid f_S(x) \leq b\}$, in other words, $S_{\in [a,b]}$ is the intersection of S with the range between the two parallel hyperplanes $f_S^{-1}(a)$ and $f_S^{-1}(b)$. Finally, for $S_{=h} = S_{\in [h,h]}$ the hyperplanes coincide, and $S_{=h}$ is the intersection of S and the hyperplane $f_S^{-1}(h)$. Summing up, for closed intervals I , the level set $S_{\in I}$ is the intersection of S with at most two half-spaces, hence a polyhedral set, and if S is bounded, $S_{\in I}$ is also bounded. Since S is a polyhedral set, we can conclude that the level set $S_{\in I}$ is a polyhedral set as well, and if S is a convex polytope, then so is $S_{\in I}$.

The Face Lattice of a Level Set Cell. As a next step, we consider the faces of such level set cells. Recall from Definition 1.8 that all non-empty faces of a polyhedral set P can be written as the intersection of P with hyperplanes that bound the half-spaces used to define P . Furthermore each such intersection is a face of P , possibly empty, but different intersections might represent the same face. Since the level set $S_{\in I}$ is defined by the half spaces used to define S and at most two additional half-spaces, the hyperplanes we have to consider for faces of $S_{\in I}$ are the hyperplanes used for S and at most two additional hyperplanes, namely $f_S^{-1}(h)$ for $S_{\leq h}$, $S_{\geq h}$, and $S_{=h}$ or $f_S^{-1}(a)$ and $f_S^{-1}(b)$ for $S_{\in [a,b]}$. The possible choices of hyperplanes lead to the list given in Table 2.3. Note that this enumeration can be summed up by the statement that the faces of $S_{\in I}$ are the cells $T_{\in Y}$ where T is a face of S and Y is either I itself or the singleton set of one of its endpoints. Regarding the interval I as a polyhedral set, we can describe Y as a face of I .

Now let us rephrase and prove our observations on faces of level set cells using the notion of interval complexes as suggested in Subsection 1.2.2. Recall

that we denote by $\mathcal{F}(P)$ the set of faces of a polyhedral set P .

Lemma 2.1. *Assume that S is a polyhedral set, $f_S: \text{aff}(S) \rightarrow \mathbb{R}$ is a real valued affine linear function defined on the affine hull of S , and the level range $I \subseteq \mathbb{R}$ is a closed interval. Then the level set $S_{\in I}$ is a polyhedral set. Its set of faces $\mathcal{F}(S_{\in I})$ is the set of all level sets $T_{\in X}$ where T is a face of S and X is a face of I in the interval complex $\mathcal{F}(I)$:*

$$\mathcal{F}(S_{\in I}) = \{T_{\in X} \mid T \in \mathcal{F}(S), X \in \mathcal{F}(I)\}$$

In other words, if T is a face of S and X is face of I , then $T_{\in X}$ is a face of $S_{\in I}$ and every face of $S_{\in I}$ is of this type. Note that if S is a convex polytope, then $S_{\in I}$ is a convex polytope, but simplicial S does not necessarily lead to a simplicial $S_{\in I}$.

Proof. We start with the degenerate case that f_S is constant. If the constant function value does not belong to the interval I , then $S_{\in I}$ is the empty cell, and for each face T of S and each face X of I , the set $T_{\in X}$ is empty as well. If the constant function value belongs to I , then $S_{\in I}$ agrees with S , and each face T of S agrees with $T_{\in I}$. Likewise, for faces X of I either $T_{\in X}$ is empty or agrees with T . From these observations the desired results follow easily for constant maps. For the trivial intervals \mathbb{R} and \emptyset , similar observations yield the results: For each face T of S , including S itself, $T_{\in \mathbb{R}}$ agrees with T and $T_{\in \emptyset}$ is empty.

Now let us deal with the more interesting cases where f_S is non-constant and I is non-trivial. As before, half-spaces and hyperplanes are taken with respect to the affine hull of S . So let $(H_k^+)_{k \in K}$ be a family of half-spaces in $\text{aff}(S)$ defining S , and let $(H_k)_{k \in K}$ be the corresponding family of hyperplanes that bound those half-spaces. Then $S_{\in I}$ can be represented as the intersection of S with one or two half-spaces, namely the half-space $f_S^{-1}(-\infty, b]$ if I has a right endpoint b , and the half-space $f_S^{-1}[a, \infty)$ if I has a left endpoint a . Combined with the representation of S as intersection of half-spaces, we obtain a representation as intersection of half-spaces for $S_{\in I}$, showing that $S_{\in I}$ is a polyhedral set.

Now we prove the statements regarding the faces of $S_{\in I}$. First we show that each face of $S_{\in I}$ is of the desired form. The empty face is represented by any $T_{\in X}$ where $T = \emptyset$ or $X = \emptyset$. So let F be a non-empty face of $S_{\in I}$. Using the facts from Definition 1.8 mentioned above, F is the intersection of $S_{\in I}$ with some of the H_k and possibly the level hyperplane $f_S^{-1}(b)$ bounding the half-space $f_S^{-1}(-\infty, b]$ if I has a right endpoint b , or the level hyperplane $f_S^{-1}(a)$ bounding the half-space $f_S^{-1}[a, \infty)$ if I has a left endpoint a . Let J be the subset of K such that $(H_k)_{k \in J}$ are the hyperplanes involved in this intersection. Then $T = S \cap \bigcap_{k \in J} H_k$ is a face of S , and we have:

$$S_{\in I} \cap \bigcap_{k \in J} H_k = f_S^{-1}(I) \cap S \cap \bigcap_{k \in J} H_k = f_S^{-1}(I) \cap T = T_{\in I} \quad (2.1)$$

If no level hyperplane $f_S^{-1}(h)$ for h an endpoint of I is involved, we already have the desired form. Since two different level hyperplanes $f_S^{-1}(a)$ and $f_S^{-1}(b)$ for $a \neq b$ are disjoint, at most one level hyperplane $f_S^{-1}(h)$ is involved. For $X = \{h\} \in \mathcal{F}(I)$ we obtain:

$$F = S_{\in I} \cap \bigcap_{k \in J} H_k \cap f_S^{-1}(X) = f_S^{-1}(I) \cap T \cap f_S^{-1}(X) = T_{\in X} \quad (2.2)$$

It remains to show that each $T_{\in X}$ is a face of $S_{\in I}$, that is, $T_{\in X}$ is empty or the intersection of $S_{\in I}$ with hyperplanes bounding half-spaces defining $S_{\in I}$. The set of such hyperplanes includes the hyperplanes H_k and the level hyperplanes $f_S^{-1}(h)$ for h an endpoint of I . We can assume $T \neq \emptyset$ and $X \neq \emptyset$. Then $T = S \cap \bigcap_{k \in J} H_k$ for some $J \subseteq K$. If $X = I$, then $T_{\in I} = S_{\in I} \cap \bigcap_{k \in J} H_k$ by Equation (2.1), and if $X = \{h\}$ for some endpoint of I , then $T_{\in X} = S_{\in I} \cap \bigcap_{k \in J} H_k \cap f_S^{-1}(X)$ by Equation (2.2) as desired.

The observations in the last remark are fairly trivial: Any polyhedral set contained in a convex polytope is a convex polytope; the triangle with a quadrilateral level set depicted earlier in Figure 1.1 is a simple example illustrating the second part. \square

The lemma and its proof can be generalised to functions f_S taking values in higher dimensional spaces and level ranges R that are polyhedral sets instead of an interval I :

Lemma 2.2. *Assume that S and R are polyhedral sets, and that f_S is an affine linear function defined on the affine hull $\text{aff}(S)$ of S taking values in some $\mathbb{R}^r \supseteq R$. Then the level set $S_{\in R} := \{x \in S \mid f_S(x) \in R\}$ is a polyhedral set. Its set of faces $\mathcal{F}(S_{\in R})$ is the set of all level sets $T_{\in X}$ where T is a face of S and X is a face of R :*

$$\mathcal{F}(S_{\in R}) = \{T_{\in X} \mid T \in \mathcal{F}(S), X \in \mathcal{F}(R)\}$$

In other words, if T is a face of S and X is face of R , then $T_{\in X}$ is a face of $S_{\in R}$ and every face of $S_{\in R}$ is of this type. Note that if S is a convex polytope, then $S_{\in R}$ is a convex polytope as well, whereas the range R being a convex polytope is not sufficient for such a conclusion.

Proof. Again as in the previous proof, let $(H_k^+)_{k \in K}$ be a family of half-spaces in $\text{aff}(S)$ defining S , and let $(H_k)_{k \in K}$ be the corresponding family of hyperplanes bounding those half-spaces. Analogously, let $(h_\kappa^+)_{\kappa \in L}$ be a family of half-spaces in \mathbb{R}^r defining R with the corresponding family $(h_\kappa)_{\kappa \in L}$ of bounding hyperplanes. First we show that $S_{\in R}$ is a polyhedral set. We can rewrite $S_{\in R}$ as $S \cap \bigcap_{\kappa \in L} f_S^{-1}(h_\kappa^+)$. It suffices to show that each $f_S^{-1}(h_\kappa^+)$ is an intersection of half spaces. More precisely, we prove that $f_S^{-1}(h_\kappa^+)$ is either empty, the whole space $\text{aff}(S)$, or a half-space bounded by $f_S^{-1}(h_\kappa)$.

Recall the following basic facts from linear algebra: Any half-space of an affine space A can be written as the solution set of some inequality $g(y) \geq 0$ where g is a real-valued affine linear function defined on A . Moreover, the corresponding bounding hyperplane is formed by the points y where equality, that is $g(y) = 0$, is attained. On the other hand, any such solution set $g(y) \geq 0$ is either empty, the whole space, or a half-space bounded by the hyperplane $g(y) = 0$, with the first two options occurring for constant maps g only. If $g(y) \geq 0$ is empty, then so is $g(y) = 0$; and if $g(y) \geq 0$ is the whole space then $g(y) = 0$ is either empty or the whole space. So assume that $g(y) \geq 0$ characterises h_κ^+ and hence $g(y) = 0$ characterises h_κ . Then $g(f_S(x)) \geq 0$ characterises $f_S^{-1}(h_\kappa^+)$ and $g \circ f_S$ is an affine linear map, showing that $f_S^{-1}(h_\kappa^+)$ is either empty, the whole space $\text{aff}(S)$, or a half-space bounded by the hyperplane $g(f_S(x)) = 0$, which can be described equivalently as $f_S(x) \in h_\kappa$ or $f_S^{-1}(h_\kappa)$ as desired.

Note that, from a technical point of view, the representation of $S_{\in R}$ as $\bigcap_{k \in K} H_k^+ \cap \bigcap_{\kappa \in L} f_S^{-1}(h_\kappa^+)$ is not necessarily a representation as intersection of half-spaces since some of the $f_S^{-1}(h_\kappa^+)$ might not be half-spaces. Nevertheless one easily checks that the relation from Definition 1.8 between the set of faces of $S_{\in R}$ and the set of its intersections with subfamilies of its bounding hyperplanes extends to the set of intersections $S_{\in R} \cap \bigcap_{k \in J} H_k \cap \bigcap_{\kappa \in J'} f_S^{-1}(h_\kappa)$ for subfamilies indexed by $J \subseteq K$ and $J' \subseteq L$, although some of the $f_S^{-1}(h_\kappa)$ might not be bounding hyperplanes of $S_{\in R}$, namely, each non-empty face of $S_{\in R}$ can be represented as such an intersection, and any such intersection is a face of $S_{\in R}$. This is obvious if $S_{\in R}$ is empty because its only face, the empty cell, is represented by all intersections $S_{\in R} \cap \bigcap_{k \in J} H_k \cap \bigcap_{\kappa \in J'} f_S^{-1}(h_\kappa)$. Otherwise all preimages $f_S^{-1}(h_\kappa^+)$ are either half-spaces or the whole space $\text{aff}(S)$. Preimages that are the whole space can be neglected when taking intersections, hence $S_{\in R}$ has a representation as intersection of half-spaces $\bigcap_{k \in K} H_k^+ \cap \bigcap_{\kappa \in L'} f_S^{-1}(h_\kappa^+)$ where $L' \subseteq L$ is the subset of indices κ such that $f_S^{-1}(h_\kappa^+)$ is a half-space. For this representation the characterisation from Definition 1.8 applies directly. This yields for each non-empty face of $S_{\in R}$ a representation as $S_{\in R} \cap \bigcap_{k \in J} H_k \cap \bigcap_{\kappa \in J'} f_S^{-1}(h_\kappa)$ where only indices κ from L' occur. Allowing more subfamilies by permitting indices κ from L will only increase the possible representations. But the preimages $f_S^{-1}(h_\kappa)$ for κ outside of L' are either empty or the whole space. If an empty preimage occurs, the whole intersection represents the empty face; the terms representing the whole space can be omitted without affecting the resulting intersection.

With this prerequisite we can verify the claimed description of the set of faces. We first show that for each face T of S and each face X of R , the set $T_{\in X}$ is a face of $S_{\in R}$. We can assume $T \neq \emptyset$ and $X \neq \emptyset$, otherwise $T_{\in X}$ is the empty face. Then $T = S \cap \bigcap_{k \in J} H_k$ for some index set $J \subseteq K$ and $X = R \cap \bigcap_{\kappa \in J'} h_\kappa$ for some index set $J' \subseteq L$. Substituting these representations for T and X in $T_{\in X}$ and rearranging terms yields the following equation:

$$\begin{aligned} T_{\in X} &= S \cap \bigcap_{k \in J} H_k \cap f_S^{-1}(R \cap \bigcap_{\kappa \in J'} h_\kappa) = S \cap \bigcap_{k \in J} H_k \cap f_S^{-1}(R) \cap \bigcap_{\kappa \in J'} f_S^{-1}(h_\kappa) \\ &= S \cap f_S^{-1}(R) \cap \bigcap_{k \in J} H_k \cap \bigcap_{\kappa \in J'} f_S^{-1}(h_\kappa) = S_{\in R} \cap \bigcap_{k \in J} H_k \cap \bigcap_{\kappa \in J'} f_S^{-1}(h_\kappa) \quad (2.3) \end{aligned}$$

The last term of the equation is a face of $S_{\in R}$ as we have checked just above.

Conversely, for showing that any face F of $S_{\in R}$ is of the form $T_{\in X}$, we can apply the characterisation of faces as intersections and Equation (2.3) backwards: The empty face is represented by any term $T_{\in X}$ where $T = \emptyset$ or $X = \emptyset$. So let F be a non-empty face of $S_{\in R}$. By our previous observation, F has a representation as intersection $S_{\in R} \cap \bigcap_{k \in J} H_k \cap \bigcap_{\kappa \in J'} f_S^{-1}(h_\kappa)$. Furthermore $T = S \cap \bigcap_{k \in J} H_k$ is a face of S , and $X = R \cap \bigcap_{\kappa \in J'} h_\kappa$ is a face of R . Reading Equation (2.3) backwards proves that F agrees indeed with $T_{\in X}$ for this choice of T and X .

Concerning the simple observations of the last remark, the first part follows again from the fact that polyhedral subsets of convex polytopes are convex polytopes. The simplest examples for the second part are constant maps on an unbounded polyhedral set S : For the constant result c of the function, $S_{\in \{c\}} = S$ is unbounded although the singleton set $\{c\}$ is a convex polytope. \square

Vertices of Level Set Cells. Besides their characterisation as bounded polyhedral sets, convex polytopes are also characterised as convex hulls of finite point sets. In fact, each convex polytope is the convex hull of its vertices. Regarding the representation of $S_{\in I}$ as polyhedral set, we already observed that the cell $S_{\in I}$ can be described as the intersection of the half-spaces used to define S and the half-space $f_S^{-1}[a, \infty)$ if a is a left endpoint of I and the half-space $f_S^{-1}(-\infty, b]$ if b is a right endpoint of I . The resulting set of half-spaces is not necessarily minimal, even if the subset used to define S is minimal for S . For the representation of $S_{\in I}$ as convex hull of its vertices, we want to identify the vertices of the cell $S_{\in I}$. Since vertices are 0-dimensional faces, we can use our observations from above.

Lemma 2.3. *Let S be a convex polytope with a real valued affine linear function f_S defined on the affine hull of S , and let I be a closed interval. Then the vertices of $S_{\in I}$ are either vertices of S lying in $f_S^{-1}(I)$ or the intersection of an edge uv of S with a hyperplane $f_S^{-1}(h)$ where h is an endpoint of I with $f_S(u) < h$ and $f_S(v) > h$.*

Proof. Vertices of S lying in $f_S^{-1}(I)$ and intersections of edges uv of S with a hyperplane $f_S^{-1}(h)$ that separates u and v are 0-dimensional faces and hence vertices of $S_{\in I}$. It remains to show that each vertex of $S_{\in I}$ is of one of the two kinds.

Let $(H_k)_{k \in K}$ be the family of hyperplanes that bound half-spaces defining S . Then vertices of $S_{\in I}$ are intersections of some of the H_k and possibly a hyperplane $f_S^{-1}(h)$ where h is an endpoint of I , provided this intersection is a single point w and lies in $S_{\in I}$. If only the hyperplanes H_k are involved, that is $w = \bigcap_{k \in J \subseteq K} H_k$, the intersection point w is a vertex of S lying in $f_S^{-1}(h)$. Otherwise, if a hyperplane $f_S^{-1}(h)$ is necessary in the intersection, that is $w = \bigcap_{k \in J \subseteq K} H_k \cap f_S^{-1}(h) \subsetneq \bigcap_{k \in J \subseteq K} H_k$, the intersection of the remaining hyperplanes $\bigcap_{k \in J \subseteq K} H_k$ is a line ℓ . If the intersection of ℓ with S is not an edge uv of S with $f_S(u) < h$ and $f_S(v) > h$, which implies the second alternative $w = uv \cap f_S^{-1}(h)$, then either $S \cap \ell$ is a vertex of S or $S \cap \ell$ is an edge uv intersecting $f_S^{-1}(h)$ at one of its endpoints. In the first case w is the vertex $S \cap \ell$, in the latter case one of the vertices u or v . Thus, in both cases w is already covered by the first alternative. \square

We refrain from generalising this observation to higher dimensional codomains because characterising 0-dimensional faces gets more involved when several level hyperplanes $f_S^{-1}(h_\kappa)$ and their possible linear dependencies among each other and with other bounding hyperplanes have to be taken into account.

2.1.3 Definition of Level Set Complexes

Now that we have studied level set cells and their faces, we are in a position to formally define level set complexes by proving that suitable collections of level set cells fulfil the requirements for being a complex. We start with a version focusing on the typical case with a polytopal complex as domain and a closed interval as level range. The proof however will only be given for the general version following directly afterwards.

Lemma 2.4. *Let M be a polytopal complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells and let I be a closed interval. The set of all level set cells $S_{\in Y}$ where S is a cell of M and Y is a face of I forms a polytopal complex denoted by $M_{\in I}$ with underlying set $|M|_{\in I}$.*

$$M_{\in I} = \{S_{\in Y} \mid S \in M, Y \in \mathcal{F}(I)\}$$

For the face poset structures (M, \sqsubseteq) , $(\mathcal{F}(I), \sqsubseteq)$, and $(M_{\in I}, \sqsubseteq)$ the following implication holds:

$$T \sqsubseteq S \wedge X \sqsubseteq Y \implies T_{\in X} \sqsubseteq S_{\in Y}$$

Lemma 2.5. *Let M be a polyhedral complex embedded in some \mathbb{R}^m with a map $f: |M| \rightarrow \mathbb{R}^r$ linear on cells and let R be a polyhedral complex embedded in \mathbb{R}^r . The set of all level set cells $S_{\in Y}$ where S is a cell of M and Y is a cell of R forms a polyhedral complex denoted by $M_{\in R}$ with underlying set $|M|_{\in R}$.*

$$M_{\in R} = \{S_{\in Y} \mid S \in M, Y \in R\}$$

For the face poset structures (M, \sqsubseteq) , (R, \sqsubseteq) , and $(M_{\in R}, \sqsubseteq)$ the following implication holds:

$$T \sqsubseteq S \wedge X \sqsubseteq Y \implies T_{\in X} \sqsubseteq S_{\in Y} \quad (2.4)$$

If M is a polytopal complex, then $M_{\in R}$ is polytopal as well. If $R = \mathcal{F}(P)$ is the face lattice of a polyhedral set P , we usually abbreviate the complex $M_{\in \mathcal{F}(P)}$ by $M_{\in P}$.

Proof. We first show that the family of sets $M_{\in R}$ as defined above is a polyhedral complex. Large parts of the proof follow directly from Lemma 2.2: We have proven that all sets of the form $S_{\in Y}$ for $S \in M$ and $Y \in R$ are polyhedral sets. Furthermore, each face of such a cell is of the form $T_{\in X}$ for T a face of S and X a face of Y . Because M and R are complexes, T is a cell of M and X a cell of R , yielding $T_{\in X} \in M_{\in R}$. Therefore the family $M_{\in R}$ contains with each cell also all its faces. For completing the proof that the cells in $M_{\in R}$ form a complex, it remains to show that the intersection of two cells in $M_{\in R}$ is a face of both. Using again the fact that M and R are complexes, it follows for $S, T \in M$ and $Y, X \in R$ that the intersection $S \cap T \in M$ is a face of S and T and that the intersection $Y \cap X \in R$ is a face of X and Y . Therefore, the equality

$$S_{\in Y} \cap T_{\in X} = S \cap f^{-1}(Y) \cap T \cap f^{-1}(X) = (S \cap T)_{\in Y \cap X}$$

shows that the intersection of two cells of $M_{\in R}$ is again a cell of $M_{\in R}$ and a face of both by Lemma 2.2.

The underlying set of $M_{\in R}$ is the union of all the cells $S_{\in Y}$. We obtain

$$\begin{aligned} |M_{\in R}| &= \bigcup_{S \in M, Y \in R} S_{\in Y} = \bigcup_{S \in M, Y \in R} (S \cap f^{-1}(Y)) \\ &= \bigcup_{Y \in R} f^{-1}(Y) \cap \bigcup_{S \in M} S = f^{-1}(|R|) \cap |M| = |M|_{\in R} \end{aligned}$$

as desired. Implication (2.4) follows from Lemma 2.2 as well, just as the fact that $M_{\in R}$ is a polytopal complex if M is, because polytopal cells S generate only polytopal level set cells $S_{\in Y}$. \square

The above results justify the following definition of level set complexes.

Definition 2.6 (level set complex). Let M and R be polyhedral complexes with a map $f: |M| \rightarrow \mathbb{R}^r$ linear on cells of M , where R is embedded in \mathbb{R}^r . Then the complex $M_{\in R}$ consisting of the cells $S_{\in Y} = S \cap f^{-1}(Y)$ for each $S \in M$ and $Y \in R$ is called a *level set complex* of M with level range R .

Intersection Complexes as Level Set Complexes. Before we go on, let us resume some discussion from Subsection 1.2.3. The first topic is the connection between level set complexes and intersection complexes that we foreshadowed. Recall that an intersection complex $K \cap L$ of two complexes K and L embedded in \mathbb{R}^d was defined as the set of cells $S \cap T$ where S is some cell of K and T is some cell of L . When we define $f: |K| \rightarrow \mathbb{R}^d$ as inclusion map, the resulting level set complex $K_{\in L}$ agrees with the intersection complex $K \cap L$. We can obviously switch the roles of K and L and obtain another representation of $K \cap L$ as level set complex $L_{\in K}$ with respect to the inclusion map of $|L|$ into \mathbb{R}^d . Therefore intersection complexes are a special case of level set complexes.

Monotone Maps to the Face Poset of a Level Set Complex. The other topic to which we add some additional remarks addresses order preserving maps between the poset structures of certain complexes. We observed in Subsection 1.2.3 that for polyhedral complexes M and R , there is a surjective monotone map, let us call it ψ , from the product order of two face posets (M, \sqsubseteq_M) and (R, \sqsubseteq_R) to the product complex $M \times R$. The map ψ assigns to a pair (S, Y) of faces in the product order the cell $S \times Y$ in $M \times R$, and gives a poset isomorphism when restricted to the pairs (S, Y) whose members are either both non-empty or both empty. Now if $R \subseteq \mathbb{R}^r$ is the considered level range for some map $f: |M| \rightarrow \mathbb{R}^r$ linear on cells, then Implication (2.4) states that a similar map ϕ assigning to each pair (S, Y) of faces in the product order of (M, \sqsubseteq_M) and (R, \sqsubseteq_R) the level set cell $\phi(S, Y) = S_{\in Y}$ in the level set complex $M_{\in R}$ is order preserving as well. Furthermore ϕ factors through the product complex $M \times R$ via ψ in the sense that there is a monotone map $\phi': M \times R \rightarrow M_{\in R}$ from the face poset of the product complex $M \times R$ to the face poset of the level set complex $M_{\in R}$ such that $\phi = \phi' \circ \psi$. Obviously, ϕ' is given by the assignment $\phi'(S \times Y) = S_{\in Y}$, which only requires a short note on its well-definedness: The only cell in $M \times R$ with ambiguous representation is the empty cell because $\emptyset = S \times \emptyset = \emptyset \times Y$ holds for all S in M and Y in R . But the corresponding cells $S_{\in \emptyset}$ and $\emptyset_{\in Y}$ in the level set complex represent the empty cell as well. Since all cells of the level set complex $M_{\in R}$ are representable as level set cell $S_{\in Y}$ for some $S \in M$ and $Y \in R$, the maps ϕ and ϕ' are surjective.

Observations. The ambiguity of the representation of cells in a level set complex is in general not limited to the empty cell. We will investigate this topic in more detail in the next section. The availability of alternative representations for the empty cell allows to omit all terms $S_{\in \emptyset}$ in the definition of the complex $M_{\in I}$ for a closed interval $I \neq \emptyset$ without losing a cell. Doing so yields the descriptions given in Table 2.4 for the typical interval level ranges. From these descriptions, we can directly read off that $M_{=h}$ is a subcomplex of $M_{\leq h}$ and $M_{>h}$, and that $M_{=a}$ and $M_{=b}$ are subcomplexes of $M_{\in [a,b]}$. More generally, the

Table 2.4: The typically considered level set complexes and their cells. The letter M denotes a polyhedral complex, and h , a , and b are real numbers.

$$\begin{aligned} M_{=h} &= \{S_{=h} \mid S \in M\} \\ M_{\leq h} &= \{S_{=h} \mid S \in M\} \cup \{S_{\leq h} \mid S \in M\} \\ M_{>h} &= \{S_{=h} \mid S \in M\} \cup \{S_{>h} \mid S \in M\} \\ M_{\in[a,b]} &= \{S_{=a} \mid S \in M\} \cup \{S_{=b} \mid S \in M\} \cup \{S_{\in[a,b]} \mid S \in M\} \end{aligned}$$

following facts are immediate consequences of the definitions of subcomplexes and subdivisions.

Fact 2.7. If K is a subcomplex of M and Y is a face or subcomplex of R , then $K_{\in Y}$ is a subcomplex of $M_{\in R}$. Similarly, if K is a subdivision of M and Y is a subdivision of R , then $K_{\in Y}$ is a subdivision of $M_{\in R}$. In particular, if $|R|$ contains the image of f , then $M_{\in R}$ is a subdivision of M . Since any interval complex \mathcal{J} can be regarded as a subcomplex of some subdivision \mathcal{I} of \mathbb{R} , we obtain for any interval complex \mathcal{J} an associated subcomplex $M_{\in \mathcal{J}}$ of some subdivision $M_{\in \mathcal{I}}$ of M .

2.2 Canonical Representations

In the previous section, we have shown that each cell of a level set complex $M_{\in R}$ is represented by a term $S_{\in Y}$ containing a cell S from M and a cell Y from R as parameters. We even formalised this representation by a pair of cells as an order preserving map from the product order of the face posets of M and R to the face poset of the level set complex. Moreover we observed that the representation of a cell of $M_{\in R}$ by a term $S_{\in Y}$ is not always unique, in other words the order preserving map ϕ from the product order of (M, \sqsubseteq_M) and (R, \sqsubseteq_R) to $M_{\in R}$ is not injective. In this section, we investigate in more detail how the ambiguity of the representation affects the combinatorial structure of the level set complex. We will see that among the terms representing the same cell, there is a natural candidate to be considered as the canonical representation of the cell. These canonical representations allow a description of the face poset of the level set complex up to isomorphism, as shown in Lemma 2.10. Furthermore we develop some alternative criteria for characterising canonical representations.

2.2.1 Definition of Canonical Representations and Importance for Determining the Combinatorial Structure of Level Set Complexes

Examples. We already discussed the empty cell as the most obvious example of a cell with ambiguous representation. Whenever the cells S and $f_S^{-1}(Y)$ do not intersect, in particular when $S = \emptyset$ or $Y = \emptyset$, the cell $S_{\in Y}$ is empty. Therefore the empty cell has always several representations, unless the degenerate level set complex $\{\emptyset\}_{\in \{\emptyset\}}$ is considered. It is possible to create examples of level set complexes $M_{\in R}$ where the empty cell is the only ambiguously represented

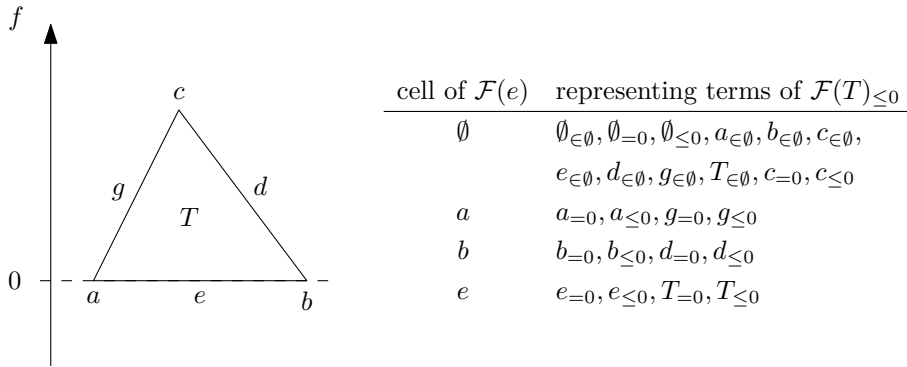


Figure 2.1: An example of a level set complex with ambiguously represented non-empty cells. Consider the triangle T depicted on the left and the complex $\mathcal{F}(T)$ formed by its faces, the non-trivial ones being its edges e , d , and g and its vertices a , b , and c . Let f be the linear map assigning to each point its height with respect to the vertical axis. Then the level set complex $\mathcal{F}(T)_{\leq 0}$ agrees with $\mathcal{F}(e)$. The edge e is an example of an ambiguously represented non-empty cell of $\mathcal{F}(T)_{\leq 0}$: The terms $T_{\leq 0}$ and $e = 0$ are two out of four possible representations. The table on the right hand side lists for each cell of $\mathcal{F}(T)_{\leq 0} = \mathcal{F}(e)$ all terms of the form $S_{\in Y}$ for some $S \in \mathcal{F}(T)$ and some $Y \in \mathcal{F}((-\infty, 0])$ representing the cell.

cell and terms of the form $S_{\in \emptyset}$ and $\emptyset_{\in Y}$ are the only terms representing it. Such examples are equivalently characterised by the property that the monotone map ϕ' from the product complex $M \times R$ to the level set complex $M_{\in R}$ assigning $S_{\in Y}$ to $S \times Y$ is injective; in fact, ϕ' is even a poset isomorphism in this case, as we will see later. An easy example of this kind is given by the intersection complex formed by two complexes in \mathbb{R}^2 , one consisting of horizontal lines, the other of vertical lines; each pair of a horizontal and a vertical line intersects at a unique vertex of the intersection complex. On the other hand, such examples are somewhat untypical, as illustrated by the following sufficient condition for a non-trivial representation of the empty cell: Whenever R contains at least two vertices u and w , and M contains at least one vertex v , then at least one of $v_{\in u}$ and $v_{\in w}$ is empty.

An example of an ambiguously represented non-empty cell in a level set complex is depicted in Figure 2.1. This example apparently relies on choosing a linear map f on a triangle T in such a way that an edge of T is parallel to the isolines of f and choosing exactly the value of f at this edge as a vertex of the level range. One easily checks that slight perturbations of the vertices of the triangle, or of the map f , or of the endpoints of the interval considered as level range suffice to spoil the property of having a non-empty ambiguously represented cell. We will prove later in Corollary 2.16 a characterisation for canonical terms implying that for terms $S_{\in Y}$ that are not the canonical representation of their cell, either $f_S(S)$ intersects Y in a proper face of Y or $f_S^{-1}(Y)$ intersects S in a proper face of S . Therefore the property that distinct terms represent the same non-empty cell is in general sensitive to perturbations. In this sense one could say that non-empty ambiguously represented cells are a degenerate

feature of level set complexes.

The ambiguity in the representation of a cell $S_{\in Y}$ in a level set complex $M_{\in R}$ has another consequence. If $S_{\in Y}$ and $T_{\in X}$ represent the same cell, then it is also represented by the intersection $(S \cap T)_{\in Y \cap X}$. In this situation, we have $T_{\in X} \sqsubseteq (S \cap T)_{\in Y \cap X}$, but neither $T \sqsubseteq S \cap T$ nor $X \sqsubseteq Y \cap X$ is necessarily true. A concrete instance of this situation can be found in the example of Figure 2.1: The cell $T_{\leq 0}$ agrees with and thus is a face of $e_{=0}$, but neither the triangle T is a face of its edge e nor the interval $(-\infty, 0]$ is a face of its endpoint 0. Therefore, the reverse implication of the statement $T \sqsubseteq S \wedge X \sqsubseteq Y \implies T_{\in X} \sqsubseteq S_{\in Y}$ made in Lemma 2.5, Implication (2.4) is not true in general. By introducing the notion of a canonical representation of a level set cell, we obtain a description of level set cells that avoids double counting and allows for a partial converse of that implication.

Definition of Canonical Representations. The observation, that a cell with two representations can also be represented by the intersection, suggests to use the smallest possible representation according to the face poset ordering in M and R , which is well-defined by that observation.

Definition 2.8 (canonical representation). Let $T_{\in X}$ be a level set cell of a level set complex $M_{\in R}$ for some $T \in M$ and $X \in R$. The term $T_{\in X}$ *represents* the cell $T_{\in X}$ *canonically*, if it is the smallest representation of the cell, that is, for any other term $U_{\in Z}$ with $U \in M$ and $Z \in R$ such that $U_{\in Z} = T_{\in X}$, the cell T is a face of U in M and X is a face of Z in R . Equivalently, $T_{\in X}$ is a *canonical representation* of $T_{\in X}$, if $T = \bigcap_{j \in J} U(j)$ and $X = \bigcap_{j \in J} Z(j)$ where the family $(U(j)_{\in Z(j)})_{j \in J}$ is the collection of all representations of $T_{\in X}$ by cells $U(j) \in M$ and $Z(j) \in R$. In particular, each cell of a level set complex $M_{\in R}$ has a unique canonical representation. We will frequently abbreviate the statement that a term $T_{\in X}$ is a canonical representation of the cell $T_{\in X}$ by “ $T_{\in X}$ is canonical” or similar statements.

Determining the Combinatorial Structure of a Level Set Complex from the Terms Canonically Representing Its Cells. The next lemma is the partial converse of the implication $T \sqsubseteq S \wedge X \sqsubseteq Y \implies T_{\in X} \sqsubseteq S_{\in Y}$ from Lemma 2.5 mentioned before.

Lemma 2.9. *Assume that M and R are a polyhedral complexes. Let S and T be cells of M and let Y and X be cells of R . If $T_{\in X}$ is a face of $S_{\in Y}$ and the term $T_{\in X}$ represents the cell canonically, then T is a face of S and X is a face of Y .*

Proof. If $T_{\in X}$ is a face of $S_{\in Y}$, then the intersection of both is the face $T_{\in X}$ itself. Hence, $T_{\in X}$ is also represented by $(S \cap T)_{\in Y \cap X}$. By assumption, $T_{\in X}$ is the representation by the smallest possible cells, thus $T \sqsubseteq S \cap T \sqsubseteq S$ and $X \sqsubseteq Y \cap X \sqsubseteq Y$. \square

Summing up the observations we obtain the following conclusion which shows that the combinatorial structure of a level set complex $M_{\in R}$, that is, its face poset up to isomorphism, is uniquely determined by the combinatorial structures of M and R and by the subset of the set-theoretic Cartesian product $M \times R$ consisting of those pairs (S, Y) whose corresponding term $S_{\in Y}$ canonically represents the corresponding cell $S_{\in Y}$.

Lemma 2.10. *For polyhedral complexes M and R , let S and T be cells of M and let Y and X be cells of R . If T is a face of S and X is a face of Y , then $T_{\in X}$ is a face of $S_{\in Y}$. Moreover, if $T_{\in X}$ is a canonical representation, the converse is also true: $T_{\in X}$ is a face of $S_{\in Y}$ if and only if T is a face of S and X is a face of Y .*

$$\begin{aligned} T \sqsubseteq S \wedge X \sqsubseteq Y &\implies T_{\in X} \sqsubseteq S_{\in Y} \text{ for } S, T \in M \text{ and } X, Y \in R \\ T \sqsubseteq S \wedge X \sqsubseteq Y &\iff T_{\in X} \sqsubseteq S_{\in Y} \text{ if additionally } T_{\in X} \text{ is canonical} \end{aligned}$$

Proof. The first observation is shown in Lemma 2.2, the converse for canonically represented cells in Lemma 2.9. \square

The equivalence $T \sqsubseteq S \wedge X \sqsubseteq Y \iff T_{\in X} \sqsubseteq S_{\in Y}$ for canonical terms $T_{\in X}$ shows that the monotone map ϕ from the product order of (M, \sqsubseteq_M) and (R, \sqsubseteq_R) assigning $S_{\in Y}$ to the pair (S, Y) defines a poset isomorphism when restricted to the set of pairs (S, Y) such that $S_{\in Y}$ is a canonical term. This observation carries over to the monotone map $\phi': M \times R \rightarrow M_{\in R}$ from the product complex to the level set complex: When restricted to the Cartesian products $S \times Y$ such that $S_{\in Y}$ is canonical, the map ϕ' yields a poset isomorphism. In particular, for the special case that $S_{\in \emptyset}$ and $\emptyset_{\in Y}$ for non-empty S and Y are the only non-canonical cells, the posets of the product complex $M \times R$ and the level set complex $M_{\in R}$ are isomorphic via ϕ' . Therefore we can restate the result of Lemma 2.10 informally as follows: The face poset of a level set complex $M_{\in R}$ is isomorphic to the product order of the face posets of M and R restricted to the canonical terms, that is, the pairs (S, Y) such that $S_{\in Y}$ is canonical, with the isomorphism being described by the assignments $(S, Y) \mapsto S_{\in Y}$. In particular, the canonical terms enumerate the cells of a level set complex without double-counting: $M_{\in R} = \{S_{\in Y} \mid S \in M, Y \in R, S_{\in Y} \text{ is canonical}\}$

2.2.2 Criteria Characterising Canonical Representations

In the following, we want to develop some criteria to identify those terms $S_{\in Y}$ that are canonical representations of their cells. Clearly $\emptyset_{\in \emptyset}$ is the canonical representation of the empty cell, and no term $S_{\in \emptyset}$ or $\emptyset_{\in Y}$ for $S \neq \emptyset$ and $Y \neq \emptyset$ is canonical. As a first step, we aim for a general characterisation of canonical terms applying to arbitrary level set complexes, stating that $S_{\in Y}$ is canonical if and only if both are empty or some value in the interior of Y is attained at some point in the interior of S . From this characterisation, which is presented as last criterion in Lemma 2.15, we deduce specialised criteria for closed intervals as level ranges.

General Criteria for Arbitrary Level Ranges

Slightly Strengthening the Minimality Condition in the Definition.

Let us start with an observation regarding those terms that do not canonically represent a cell. Reversing the characterisation of canonical terms as face-minimal representations yields that for any non-canonical term $S_{\in Y}$, the cell $S_{\in Y}$ has another representation $U_{\in Z}$ with faces $U \sqsubseteq S$ and $Z \sqsubseteq Y$ where at least one of the two face relations is proper. A simple argument shows that there even is a representation where exactly one of the two face relations is proper and one of U or Z agrees with S or Y respectively; in other words either

$U_{\in Y}$ for some proper face U of S , or $S_{\in Z}$ for a proper face Z of Y represents the cell $S_{\in Y}$. This allows to treat the two alternatives independently in case distinctions, with one case assuming a representation by a smaller cell of the domain, the other case assuming a representation by a smaller cell in the level range. The key observation for this result relies on the transitivity of the face relation:

Lemma 2.11. *Assume that a cell in a level set complex is represented by terms $S_{\in Y}$ and $T_{\in X}$ such that $T \sqsubseteq S$ and $X \sqsubseteq Y$. Then any intermediate term, that is a term $U_{\in Z}$ for faces U and Z with $T \sqsubseteq U \sqsubseteq S$ and $X \sqsubseteq Z \sqsubseteq Y$, is a representation of the cell as well.*

Proof. By Lemma 2.10, the cell $T_{\in X}$ is a face of $U_{\in Z}$ and $U_{\in Z}$ is a face of $S_{\in Y}$. Thus, if $T_{\in X}$ and $S_{\in Y}$ agree, $U_{\in Z} = S_{\in Y}$ as desired. \square

Now we can state the result as an alternative criterion for non-canonical terms and consequently, via contraposition, as an alternative characterisation of canonical terms.

Corollary 2.12. *Let $M_{\in R}$ be a level set complex and consider a cell represented by a term $S_{\in Y}$ for some $S \in M$ and $Y \in R$.*

1. *The term $S_{\in Y}$ is non-canonical if and only if $U_{\in Y} = S_{\in Y}$ for some proper face U of S or $S_{\in Z} = S_{\in Y}$ for some proper face Z of Y holds.*
2. *The term $S_{\in Y}$ is canonical if and only if for all proper faces U of S the inequality $U_{\in Y} \neq S_{\in Y}$, and for all proper faces Z of Y the inequality $S_{\in Z} \neq S_{\in Y}$ holds. In other words, the term $S_{\in Y}$ is canonical if and only if replacing one of S or Y in the term by a proper face always yields a proper face of the level set cell $S_{\in Y}$.*

In both characterisations, we can restrict our attention to immediate faces U of S and Z of Y .

Proof. We only prove the first statement, since the second statement is an immediate consequence resulting from negating both sides of the equivalence. The “if”-part follows from the definition of a canonical representation: The assumed alternative representations $U_{\in Y}$ or $S_{\in Z}$ render the term $S_{\in Y}$ a non-minimal representation of its cell. For the “only if”-part, assume that $S_{\in Y}$ is not canonical and the cell is also represented by a term $T_{\in X}$ for some faces $T \sqsubseteq S$ and $X \sqsubseteq Y$, where at least one of the face relations is proper. If T is a proper face of S , we choose $U = T$, otherwise $Z = X$. By Lemma 2.11, the intermediate term $U_{\in Y}$ or $S_{\in Z}$ represents $S_{\in Y}$. For the strengthened version, we modify our choice of U or Z . If T is a proper face of S , we choose for U an intermediate face $T \sqsubseteq U \sqsubseteq S$ that is an immediate face of S , otherwise we choose for Z an intermediate face $X \sqsubseteq Z \sqsubseteq Y$ that is an immediate face of Y . \square

The Role of the Interiors. The next lemma establishes the connection between the requirements regarding representations by proper faces described in the previous observation and the requirements involving interior points and values that we are aiming for. For the proof we need a fact on convex sets stated for example in [23, Proposition 2.2.4, p. 12] and an immediate consequence of it.

Fact 2.13. Assume that C is a convex set contained in the boundary of another convex set $K \subseteq \mathbb{R}^d$. Then there is a supporting hyperplane of K containing C . Since the intersection of the supporting hyperplane with K defines a proper face of K , this implies that C is contained in a proper face of K . Furthermore, since the intersection of two convex sets is convex, we can conclude that a convex set does not intersect the interior of a convex set K if and only if it intersects K at most in a proper face of K .

Lemma 2.14. *Let $f_S: \text{aff}(S) \rightarrow \mathbb{R}^r$ be an affine linear map on a polyhedral set S and Y a polyhedral set in \mathbb{R}^r . Then $S_{\in Y}$ does not intersect the interior of S but its boundary if and only if $S_{\in Y}$ is non-empty and contained in a proper face of S . Furthermore, f_S does not attain a value in the interior of Y but in its faces if and only if $S_{\in Y}$ and hence $f_S(S) \cap Y$ are non-empty and $f_S(S) \cap Y$ is contained in a proper face of Y .*

Proof. For the first statement, note that $S_{\in Y}$ is a convex subset of S . If it is not intersecting the interior of S , this implies by Fact 2.13 that $S_{\in Y}$ is contained in a proper face of S . The other direction is obviously true since a proper face does not intersect the interior. The second statement can be shown similarly: The set $f_S(S) \cap Y$ is a convex subset of Y , and it is empty if and only if $S_{\in Y}$ is empty. If $f_S(S) \cap Y$ is not intersecting the interior of Y , it is implied by Fact 2.13 once more that $f_S(S) \cap Y$ is contained in a proper face of Y . The other direction follows again because proper faces do not intersect the interior. \square

Characterisation by Interior Points Attaining Interior Values. Now we are ready to prove the characterisation suggested above.

Lemma 2.15. *The following criteria can be used to characterise canonical representations $S_{\in Y}$. The first two equivalences consider separately the two conditions that replacing in the term $S_{\in Y}$ the cell S on the one hand, and cell Y on the other hand by a proper face yields a proper face of the level set cell. The third equivalence is the criterion for canonical terms obtained as immediate consequence by requiring both conditions. The last equivalence combines the two conditions into a single one.*

1. *The inequality $U_{\in Y} \neq S_{\in Y}$ holds for all proper faces U of S if and only if S is empty or the interior of S contains a point x with $f_S(x) \in Y$.*
2. *The inequality $S_{\in Z} \neq S_{\in Y}$ holds for all proper faces Z of Y if and only if Y is empty or S contains a point x such that $f_S(x)$ is in the interior of Y .*
3. *A term $S_{\in Y}$ is canonical if and only if S and Y are empty or the interior of S contains a point x_1 such that $f_S(x_1) \in Y$ and S contains a point x_2 such that $f_S(x_2)$ is in the interior of Y .*
4. *A term $S_{\in Y}$ is canonical if and only if S and Y are empty or the interior of S contains a point x such that $f_S(x)$ is in the interior of Y .*

Proof. The right hand sides of the first two equivalences are obviously sufficient conditions for the corresponding left hand sides: If a point or its value lies in the interior of the domain or level range, then the point is not contained in the level set cell defined by a proper face of the domain or level range respectively.

We continue by showing that the criteria on the right hand sides of the first two equivalences are necessary.

For the first statement, assume that S is non-empty and does not contain an interior point x with $f_S(x) \in Y$. If $S_{\in Y}$ is empty, then choosing $U = \emptyset$ gives a representation of $S_{\in Y} = U_{\in Y}$ by a proper face U of S . Otherwise, $S_{\in Y}$ intersects S but not the interior of S . By Lemma 2.14, it is contained in some proper face U of S , which fulfils $S_{\in Y} = U_{\in Y}$.

For the second statement, assume that no value in the interior of Y is attained. If $S_{\in Y}$ is empty, then choosing $Z = \emptyset$ gives a representation of $S_{\in Y} = S_{\in Z}$ by a proper face Z of Y . Otherwise $f_S(S) \cap Y$ intersects Y but not the interior of Y . By Lemma 2.14, it is contained in some proper face Z of Y , which fulfils $S_{\in Z} = S_{\in Y}$.

The third equivalence is an immediate consequence of Corollary 2.12 and the first two equivalences. Its right hand side criterion is obviously implied by the right hand side criterion of the last equivalence. It remains to show that the right hand side of the last equivalence is a necessary condition for a canonical term. If $S_{\in Y}$ is canonical, then by the third statement, there is an interior point x_1 of S with $f_S(x_1) \in Y$ and a point $x_2 \in S$ such that $f_S(x_2)$ is in the interior of Y . Then the midpoint of x_1 and x_2 lies in the interior of S and its value under f_S is the midpoint of $f_S(x_1)$ and $f_S(x_2)$, hence lies in the interior of Y . Therefore the right hand side of the last statement is a necessary condition for $S_{\in Y}$ to be canonical. \square

Characterisations by Intersecting Interiors. Rephrasing the above criteria gives the following characterisation that was mentioned in the discussion on which kinds of ambiguously represented cells typically appear.

Corollary 2.16. *A term $S_{\in Y}$ is canonical if and only if S and Y are both empty or $f_S(S)$ intersects the interior of Y and $f_S^{-1}(Y)$ intersects the interior of S . Conversely, a term $S_{\in Y}$ is non-canonical if and only if at least one of S and Y is non-empty and $f_S(S)$ intersects Y at most in a proper face of Y or $f_S^{-1}(Y)$ intersects S at most in a proper face of S .*

Proof. The statement that $f_S(S)$ intersects the interior of Y is just another way of saying that S contains a point x such that $f_S(x)$ is in the interior of Y . Similarly, $f_S^{-1}(Y)$ intersects the interior of S if and only if $f_S(x) \in Y$ for some interior point x of S . For the converse, recall from Fact 2.13 that for convex sets, not intersecting the interior of another convex set is the same as intersecting at most a proper face of the other set. \square

For the subsequent discussion, we make use of the fact that the image $f_S(S)$ of S and the preimage $f_S^{-1}(Y)$ of Y are polyhedral sets. This fact follows from a more general result:

Fact 2.17. Images and preimages of polyhedral sets under affine maps are polyhedral sets.

We refer to Rockafellar [39, Theorem 19.3, p. 174] for a proof. The part for preimages can also be obtained as special case of the result shown in Lemma 2.2 that level set cells $S_{\in R}$ are polyhedral sets: For affine spaces A and B , which we can identify with suitable Euclidean spaces \mathbb{R}^d and \mathbb{R}^r respectively, and an

affine map $f: A \rightarrow B$, we can choose $S = A$ and $f_S = f$ so that for polyhedral sets $R \subseteq B$ the level set cell $S_{\in R}$ agrees with the preimage $f^{-1}(R)$.

We may ask whether the two conditions in the criterion for canonical cells in Corollary 2.16 can be combined into a single one in a fashion similar to the last statement of Lemma 2.15. Clearly, the single criterion of this statement can be rephrased by requiring that $f_S(\text{int}(S))$ intersects the interior of Y or alternatively by requiring that $f_S^{-1}(\text{int}(Y))$ intersects the interior of S . But the more interesting question is, whether the condition that $f_S(S)$ and Y intersect in their interiors or the condition that $f_S^{-1}(Y)$ and S intersect in their interiors can be used as criterion. It turns out that the former condition works, but for the latter condition, one direction of a presumed equivalence fails.

Let us deal with the negative result first. It relies on the observation that interior points of $f_S^{-1}(Y)$ are not necessarily preimages of interior points in Y . A simple example is a constant map f_S whose constant value c is contained in a proper face Z of Y . Then $f_S^{-1}(Y) = f_S^{-1}(c)$ is the whole affine space $\text{aff}(S)$, furthermore $\text{aff}(S)$ and S intersect in their interiors, but $S_{\in Y}$ is not canonical because $S_{\in Z}$ is a representation by smaller cells. On the other hand, if $S_{\in Y}$ is canonical and non-empty, the criterion is fulfilled: By Lemma 2.15, an interior point of S is mapped to an interior point of Y . It suffices to show that preimages of interior points of Y are interior points of $f_S^{-1}(Y)$. To show this, consider the restriction of f_S to the affine hull of $f_S^{-1}(Y)$. The image of this map is a subspace of the affine hull of Y . Let us denote the obtained linear map from $\text{aff}(f_S^{-1}(Y))$ to $\text{aff}(Y)$ by f_Y . For an interior point $y \in Y$ consider in $\text{aff}(Y)$ an open neighbourhood N of y contained in Y . Since f_Y is defined as a restriction of f_S to a set containing all preimages of elements in Y and since N is contained in Y , we have $f_Y^{-1}(N) = f_S^{-1}(N)$. Furthermore this preimage $f_Y^{-1}(N) = f_S^{-1}(N)$ is an open set in $\text{aff}(f_S^{-1}(Y))$ by continuity of f_Y , and it is a subset of $f_S^{-1}(Y)$. This implies that points in $f_S^{-1}(N)$ are interior points of $f_S^{-1}(Y)$, and since preimages of y are contained in $f_S^{-1}(N)$, we have accomplished the proof.

Now we consider the other proposed condition that $f_S(S)$ and Y intersect in their interiors. The difficulty that prevents an immediate prove by Lemma 2.15 lies in the fact that boundary points of S may be mapped to interior points of $f_S(S)$. We remedy this problem by showing that among the preimages of an interior point of $f_S(S)$ there is always an interior point of S . The other part of the proof that interior points of S are mapped to interior points of $f_S(S)$ can be conveniently inferred from the open mapping theorem. We start by stating this theorem and add the implications relevant for us. For further details on the open mapping theorem, we refer to textbooks on functional analysis, for example Conway [15].

Fact 2.18. A Banach space is a normed vector space that is Cauchy complete with respect to the metric induced by the norm [15, Definition III.1.2, p. 65]. The open mapping theorem for Banach spaces states that a surjective continuous linear operator between Banach spaces is open, that is, it maps open sets to open sets [15, Theorem III.12.1, p. 93]. We can easily infer that surjective affine linear maps between finite dimensional affine spaces are open by considering such maps as linear operators on Banach spaces in a natural way: Any affine space can be identified with a vector space by choosing some point as origin. An affine map $f: A \rightarrow B$ between affine spaces A and B turns into a continuous linear operator on vector spaces if the image $f(O)$ of the origin O in A is chosen as

origin of B . All norms on finite dimensional vector spaces are equivalent, the resulting normed vector spaces are Banach spaces, and the induced topology agrees with the standard topology. Hence a map f as above can be viewed as a continuous linear operator between Banach spaces. Enforcing surjectivity of the affine map by considering its image, an affine subspace, as codomain yields the following conclusion: Any affine map between finite dimensional affine spaces maps open sets in its domain to open sets in its image.

Now we can prove that the condition that $f_S(S)$ and Y intersect in their interiors can be used for characterising canonical terms.

Lemma 2.19. *A term $S_{\in Y}$ for non-empty S is canonical if and only if the interiors of $f_S(S)$ and Y intersect. For empty S the term is canonical if and only if Y is empty.*

Proof. The case that S is empty is trivial. For non-empty S , we start with an observation that we need for one of the directions of the equivalence, namely we check that the image $f_S(\text{aff}(S))$ of the map $f_S: \text{aff}(S) \rightarrow \mathbb{R}^r$ agrees with the affine hull $\text{aff}(f_S(S))$ of the image of S under f_S . Clearly, we have $f_S(S) \subseteq f_S(\text{aff}(S))$, which implies immediately $\text{aff}(f_S(S)) \subseteq \text{aff}(f_S(\text{aff}(S)))$. Since the image $f_S(\text{aff}(S))$ is an affine subspace of \mathbb{R}^r it agrees with its own affine hull, thus $\text{aff}(f_S(S)) \subseteq f_S(\text{aff}(S))$. On the other hand, any element z of $\text{aff}(S)$ is an affine combination of elements from S . By linearity of f_S , its image $f_S(z)$ is an affine combination of elements in $f_S(S)$, showing that $f_S(\text{aff}(S)) \subseteq \text{aff}(f_S(S))$.

We continue by showing the direction that $f_S(S)$ and Y intersecting in their interiors is a necessary condition for $S_{\in Y}$ being canonical in case of a non-empty cell S . If $S_{\in Y}$ for non-empty S is canonical then by Lemma 2.15 the cell S contains an interior point x with value $f_S(x)$ in the interior of Y . It suffices to show that interior points of S are mapped to interior points of $f_S(S)$. So let x be an interior point of S . This assumption means that x has an open neighbourhood N in $\text{aff}(S)$ such that N is contained in S . We want to apply our observation in Fact 2.18 based on the open mapping theorem to show that $f_S(N) \subseteq f_S(S)$ is an open neighbourhood of $f_S(x)$ in $\text{aff}(f_S(S))$ because this demonstrates that $f_S(x)$ is an interior point of $f_S(S)$ as desired. By Fact 2.18 the affine map f_S maps the open set N to a set $f_S(N)$ that is open in the image $f_S(\text{aff}(S))$ of f_S . Our preliminary observation that $f_S(\text{aff}(S)) = \text{aff}(f_S(S))$ completes the proof for the necessity of the condition.

Now we prove that $f_S(S)$ and Y intersecting in their interiors is sufficient for $S_{\in Y}$ being canonical. The assumption that the interiors of $f_S(S)$ and Y intersect means that f_S attains on S a value y in the interior of Y and that this value y is also interior to $f_S(S)$. The proof is completed when we show that any interior point of $f_S(S)$ has some preimage in the interior of S , because such a preimage x of the value $y = f_S(x)$ in the interior of Y demonstrates that the criterion in the last statement of Lemma 2.15 is fulfilled. For verifying this claim, we follow a proof suggested in [41]. So let us consider an interior point y of $f_S(S)$. Choose an arbitrary interior point z of S . If $f_S(z)$ agrees with y , we are done. Otherwise consider the line segment between $f_S(z)$ and y . Since y is an interior point, we can extend the line segment at its endpoint y to a slightly longer line segment from $f_S(z)$ to a point y' still lying in $f_S(S)$ such that y lies in the interior of that extended line segment. Let x' be some element of S which is mapped to y' by f_S . By linearity, the line segment ℓ from

z to x' is mapped by f_S bijectively to the line segment from $f_S(z)$ to y' , hence some interior point x of ℓ is mapped to y . Since both endpoints of ℓ lie in S and at least one of them, namely z , lies in the interior of S , all interior points of ℓ including x are interior points of S by convexity. Thus we have found some interior point x of S that is a preimage of y as desired. \square

Criteria for Closed Intervals as Level Ranges

From now on we focus on the typical situation where M is a polytopal complex and the level range is a closed interval. We already discussed the trivial case of the empty set as level range, where $\emptyset_{\in \emptyset}$ is the only canonical term. The other trivial level range besides \emptyset is $R = \mathcal{F}(\mathbb{R})$, where any term $S_{\in \mathbb{R}} = S$ is canonical except for $S = \emptyset$. For the remaining cases of closed intervals, we want to infer simple criteria for terms being canonical using our results for general level set cells. Several of these criteria can be stated in two equivalent variants, where one variant considers only values of f attained at vertices, while the other variant considers values of f attained at arbitrary points of the domain. The equivalence of these variants relies on the following easy preliminary observation:

Fact 2.20. Assume that S is a convex polytope with an affine linear function $f_S: \text{aff}(S) \rightarrow \mathbb{R}^r$. Let J be a convex set in \mathbb{R}^r , not necessarily closed. By convexity, the following equivalence holds:

$$f_S(v) \in J \text{ for all vertices } v \text{ of } S \iff f_S(x) \in J \text{ for all } x \in S$$

This is logically equivalent to the following statement where \bar{J} denotes the complement of an arbitrary convex set.

$$S \text{ contains a vertex } v \text{ with } f_S(v) \in \bar{J} \iff S \text{ contains some } x \text{ with } f_S(x) \in \bar{J}$$

In particular, by choosing for J appropriate intervals with endpoint $h \in \mathbb{R}$, this implies the equivalences listed in Table 2.5, which allow to switch between the whole domain of the cell S and its vertex set when considering certain properties of the values attained by a real valued function f_S .

Real Values Attained in the Interior of a Cell. When we consider a non-empty closed interval I as level range, then the last criterion in Lemma 2.15 states that a term $S \in I$ is canonical if and only if f_S attains an interior value h of I in the interior of S . Therefore we turn to the question when a certain value $h \in \mathbb{R}$ is attained in the interior of a cell S . Recall from Definition 1.8 that the interior of a cell is its underlying subset without the underlying subsets of its proper faces, or equivalently, it is the topological relative interior of the cell viewed as a subset of its affine hull. When considering interval complexes, it should be noted that this definition of interior differs slightly from the notion of the interior of an interval when viewed as subset of \mathbb{R} . The difference lies in the treatment of singleton intervals $I = \{h\}$, whose interior in \mathbb{R} is empty, but is $\{h\}$ itself in $\text{aff}(\{h\}) = \{h\}$. For all other kinds of closed intervals both notions of interior yield the same result. The next lemma differentiates between constant and non-constant maps for the question whether a real value is attained in the interior of a cell or not.

Table 2.5: Some equivalences regarding the values attained by an affine linear map $f_S: \text{aff}(S) \rightarrow \mathbb{R}$ on a convex polytope S and on its vertex set. The value h is an arbitrary real number. The equivalences are instances of the more general observations stated in Fact 2.20.

$$\begin{array}{l}
S \text{ contains a vertex } v \text{ with } f_S(v) < h \iff S \text{ contains some } x \text{ with } f_S(x) < h \\
S \text{ contains a vertex } v \text{ with } f_S(v) > h \iff S \text{ contains some } x \text{ with } f_S(x) > h \\
S \text{ contains a vertex } v \text{ with } f_S(v) \leq h \iff S \text{ contains some } x \text{ with } f_S(x) \leq h \\
S \text{ contains a vertex } v \text{ with } f_S(v) \geq h \iff S \text{ contains some } x \text{ with } f_S(x) \geq h \\
S \text{ contains a vertex } v \text{ with } f_S(v) \neq h \iff S \text{ contains some } x \text{ with } f_S(x) \neq h \\
f_S(v) \geq h \text{ for all vertices } v \text{ of } S \iff f_S(x) \geq h \text{ for all } x \in S \\
f_S(v) \leq h \text{ for all vertices } v \text{ of } S \iff f_S(x) \leq h \text{ for all } x \in S \\
f_S(v) > h \text{ for all vertices } v \text{ of } S \iff f_S(x) > h \text{ for all } x \in S \\
f_S(v) < h \text{ for all vertices } v \text{ of } S \iff f_S(x) < h \text{ for all } x \in S \\
f_S(v) = h \text{ for all vertices } v \text{ of } S \iff f_S(x) = h \text{ for all } x \in S
\end{array}$$

Lemma 2.21. *A linear map $f_S: \text{aff}(S) \rightarrow \mathbb{R}$ attains a value $h \in \mathbb{R}$ in the interior of the polyhedral set S if and only if f_S is constant with value h or S contains elements x^+ and x^- with $f_S(x^+) > h$ and $f_S(x^-) < h$. If S is a convex polytope, x^+ and x^- can be chosen as vertices of S . Note that the second alternative can be expressed in other words by the condition that $f_S^{-1}(h)$ is a hyperplane in the affine hull of S that cuts S .*

Proof. Assume first that f_S is not constant and attains h in a point x in the interior of S . Then $f_S^{-1}(h)$ is a hyperplane in the affine hull of S containing the point x . Consider the line through x perpendicular to that hyperplane. On one side of the hyperplane, the map f_S attains only values strictly greater than h , on the other side the values are strictly smaller. Since x lies in the interior of S , we can choose two points x^+ and x^- on the line with $f_S(x^+) > h$ and $f_S(x^-) < h$ lying sufficiently close to x such that x^+ and x^- are elements of S . If S is a convex polytope, we can apply the appropriate equivalences from Table 2.5 to obtain the desired vertices. This completes the proof for one direction of the equivalence.

For the other direction, the case that f_S is constant is trivial. The other case can be shown by contraposition: We assume that no interior point of S has the f_S -value h and show that this assumption prevents the existence of at least one of the supposed points x^+ and x^- with $f_S(x^+) > h$ and $f_S(x^-) < h$. The assumption means that the interiors of the convex sets S and $f_S^{-1}(h)$ are disjoint. In this situation a variant of the hyperplane separation theorem applies: Two convex sets can be separated by a hyperplane in their common affine hull if and only if their interiors are disjoint [23, Theorem 2.2.2, p. 11]. We can conclude that there is a hyperplane H in $\text{aff}(S)$ separating S and $f_S^{-1}(h)$. But then H must be parallel to $f_S^{-1}(h)$, because otherwise H would cut $f_S^{-1}(h)$ and cannot be separating. Therefore H is the level set hyperplane of f_S for some value h' , so that $H = f_S^{-1}(h')$ divides $\text{aff}(S)$ into $\text{aff}(S)_{f_S \leq h'}$ and $\text{aff}(S)_{f_S \geq h'}$. Since H

is separating S from $f_S^{-1}(h)$, we have two possibilities on which sides of H the two sets lie. Either S lies in $\text{aff}(S)_{f_S \leq h'}$ and $f_S^{-1}(h)$ lies in $\text{aff}(S)_{f_S \geq h'}$, which implies $f(x) \leq h' \leq h$ for all $x \in S$ contradicting the existence of x^+ . Or S lies in $\text{aff}(S)_{f_S \geq h'}$ and $f_S^{-1}(h)$ lies in $\text{aff}(S)_{f_S \leq h'}$, which implies $f(x) \geq h' \geq h$ for all $x \in S$ contradicting the existence of x^- . \square

Rephrasing this criterion for a value to be attained in the interior to a criterion for the contrary yields the following corollary.

Corollary 2.22. *The map $f_S: \text{aff}(S) \rightarrow \mathbb{R}$ does not attain a value $h \in \mathbb{R}$ in the interior of a polyhedral set S if and only if one of the following cases occurs:*

1. S is empty.
2. $f_S(x) \geq h$ for all $x \in S$ and $f_S(x^+) > h$ for some $x^+ \in S$.
3. $f_S(x) \leq h$ for all $x \in S$ and $f_S(x^-) < h$ for some $x^- \in S$.

The latter two alternatives include the case that f_S is constant with a value different from h . If S is a convex polytope, the cases can be reformulated in terms of vertices.

1. S is empty.
2. $f_S(v) \geq h$ for all vertices v of S and $f_S(v^+) > h$ for some vertex v^+ of S .
3. $f_S(v) \leq h$ for all vertices v of S and $f_S(v^-) < h$ for some vertex v^- of S .

Proof. The result can be obtained from the equivalence stated in Lemma 2.21 by negating both sides. The reformulation in terms of vertices is possible by Fact 2.20. \square

For a convex polytope S fulfilling the criterion that it contains vertices with f_S -value strictly greater than h and vertices with strictly smaller value, an interior point x with $f_S(x) = h$ can be constructed from the vertices of S . The construction relies again on a fact that we used already for finding an interior point of a cell suitable for performing an elementary starring or even constructing a derived subdivision. Namely, any point represented by a convex combination with positive coefficients for all points of a finite point set lies in the interior of the convex hull of this point set [23, Exercise 2.3.6, p. 17]. Therefore it suffices to construct a convex combination x of vertices of S with $f_S(x) = h$ such that all coefficients are positive. Let b^+ be the barycentre of all vertices of S with f_S -value strictly greater than h . Analogously b^- denotes the barycentre of all vertices with strictly smaller value. Then $f_S(b^+) > h$ and $f_S(b^-) < h$ and we can choose $\lambda \in (0, 1)$ such that $\lambda f_S(b^+) + (1 - \lambda) f_S(b^-) = h$. Let x be the barycentre of the point $\lambda b^+ + (1 - \lambda) b^-$ and all vertices of S whose f_S -value equals h . Then $f_S(x) = h$ and x lies in the interior of S because its construction yields a representation of x as a convex combination of all vertices of S where all coefficients are positive.

Real Values Attained Only on the Boundary of a Cell. The next result considers the situation that a value h is attained at some point of a cell S , but not at one of its interior points.

Lemma 2.23. *The map $f_S: \text{aff}(S) \rightarrow \mathbb{R}$ does not attain a value $h \in \mathbb{R}$ in the interior of a polyhedral set S but in its faces if and only if $f_S^{-1}(h)$ is a supporting hyperplane of S in the affine hull of S .*

Proof. The left hand side assumption implies that f_S is not constant, hence $f_S^{-1}(h)$ is a hyperplane in the affine hull of S , and it implies that this hyperplane is not disjoint from S . By Lemma 2.21, the hyperplane $f_S^{-1}(h)$ also does not cut S . The only remaining alternative is that $f_S^{-1}(h)$ is a supporting hyperplane of S . For the other direction, if $f_S^{-1}(h)$ is a supporting hyperplane, then its intersection with S is a proper face of S . Hence the value h is attained only at points lying in that face. \square

Note that Lemmata 2.21 and 2.23 yield a trichotomy for considered cells S , non-constant maps $f_S: \text{aff}(S) \rightarrow \mathbb{R}$, and real values h , which can be expressed either in terms of the hyperplane $f_S^{-1}(h)$ or in terms of where the value h is attained: The hyperplane $f_S^{-1}(h)$ either cuts S , or it supports S , or it is disjoint from S ; correspondingly, either the value h is attained in the interior of S , or it is attained in a boundary face, but not in the interior, or it is not attained at all in S . For constant maps f_S , the set $f_S^{-1}(h)$ is not a hyperplane but either the whole space or empty. The case that a value is attained in the boundary but not in the interior does not occur; either h is the constant value of f_S and is attained in the interior of S , or h differs from the constant value of f_S and is not attained at all.

A Criterion for All Kinds of Closed Intervals. As summarised in Section 1.1, smooth Morse theory considers the topological changes in the level sets $M_{\leq h}$ occurring while the value h grows. For an analogous study of level set complexes $M_{\leq h}$ in the piecewise linear case, we might want to ask which terms $S_{\leq h}$ and $S_{=h}$ for varying value of h are canonical. A generalisation of this question to arbitrary level set complexes asks for a given cell S and affine linear map $f_S: \text{aff}(S) \rightarrow \mathbb{R}^r$, for which cells Y the term $S_{\in Y}$ is canonical. Among the criteria for canonical representations developed earlier, the criterion from Lemma 2.19 that $S_{\in Y}$ is canonical if and only if the interior of Y intersects the interior of $f_S(S)$ or both are empty, seems to be particularly well suited for this point of view, when we want to characterise those cells Y that induce canonical terms for S . Applying this criterion for real valued functions f_S and closed intervals I yields the following condition, where we have to keep in mind again that we consider a singleton interval as its own interior.

Lemma 2.24. *For a non-empty polyhedral set S let J be the closed interval of the values attained by the real-valued linear function f_S on S , that is $J = f_S(S)$. If S is bounded then $J = [\min_{x \in S} f_S(x), \max_{x \in S} f_S(x)]$ is the interval bounded by the maximum and minimum of f_S on S and it suffices to take the minimum and maximum value of the vertices. Let I be a closed interval. Then the term $S_{\in I}$ for non-empty S is canonical if and only if the interiors of J and I intersect, with the interiors being taken relative to the respective affine hulls. For empty S the term is canonical if and only if I is empty.*

Proof. The observations regarding the definition of J rely on standard facts: The image of S is a polyhedral set in \mathbb{R} , hence a closed interval. A bounded polyhedral set S is compact, thus the continuous map f_S attains a maximum and a minimum on S . A maximum value h on S is characterised by the conditions that $f_S(x) \leq h$ holds for all $x \in S$ and that S contains some x with $f_S(x) \geq h$. Applying the corresponding equivalences from Table 2.5 yields that h is the maximum value attained on S if and only if it is the maximum value attained on its vertex set $\text{Vrt}(S)$. Similarly the minimum on S agrees with the minimum on $\text{Vrt}(S)$. Therefore it suffices to take the minimum and maximum values at the vertices as endpoints for J . This last observation is in fact nothing other than the famous basic result from linear programming that optimal values are attained at vertices, which means more precisely that if a linear function attains a value on a polyhedral set S that is maximal or minimal for S then there is also a vertex of S at which this value is attained.

The characterisation of the canonical terms is the one from Lemma 2.19 applied to this special case. \square

Singletons and Half-Bounded Intervals. The next result focuses on the level set cells that are defined by intervals with one endpoint h , namely cells of the form $S_{=h}$, $S_{<h}$, and $S_{>h}$. Indicated by parentheses is a variant of the result: When the considered cell S is bounded, that is, a convex polytope, then we can invoke the equivalences from Table 2.5 and replace every occurrence of the word “element” in the statement by the word “vertex”. The result distinguishes six cases depending on the position of h relative to the interval $f_S(S)$. When we think of h as a parameter on the real line that continuously increases from $-\infty$ to $+\infty$ and keep track of which of the cases below applies for a given cell S , then the label number of the applicable case is weakly monotonically increasing. In detail, for non-constant f_S attaining a maximum and a minimum on S , case 1 occurs for $h < \min_{x \in S} f_S(x)$, case 2 for $h = \min_{x \in S} f_S(x)$, case 3 for $\min_{x \in S} f_S(x) < h < \max_{x \in S} f_S(x)$, case 5 for $h = \max_{x \in S} f_S(x)$, and case 6 for $h > \max_{x \in S} f_S(x)$. If maximum or minimum or both are not attained, the condition for case 3 changes whereas cases 1 and 2 and cases 5 and 6 either disappear when the respective optimum appearing in the conditions above is not attained, or apply under the same conditions as above: If f_S attains neither maximum nor minimum, case 3 applies to all $h \in \mathbb{R}$; if f_S attains a maximum but no minimum, cases 1 and 2 disappear, case 3 applies to $h < \max_{x \in S} f_S(x)$, while the conditions for cases 5 and 6 remain unchanged; symmetrically, if f_S attains a minimum but no maximum, the conditions for cases 1 and 2 remain unchanged, case 3 applies to $h > \min_{x \in S} f_S(x)$, and cases 5 and 6 disappear. For constant maps f_S with constant value c , case 1 applies for $h < c$, case 4 for $h = c$ and case 6 for $h > c$.

Lemma 2.25. *For a given non-empty (bounded) cell S and varying value h , we can now list which of the terms $S_{=h}$, $S_{<h}$ and $S_{>h}$ are canonical and which are not.*

1. *If $f_S(x) > h$ for all elements (vertices) $x \in S$, then $S_{>h}$ is canonical and agrees with S , while $S_{=h}$ and $S_{<h}$ are not canonical. They are empty and represented by $\emptyset_{\in \emptyset}$. This case occurs for constant f_S and for non-constant f_S attaining a minimum on S .*

2. If $f_S(x) \geq h$ for all elements (vertices) $x \in S$ and for some elements (vertices) x^+ and x_0 the relations $f_S(x^+) > h$ and $f_S(x_0) = h$ hold, in other words if f_S is non-constant and h its minimum on S , then $S_{\geq h}$ is canonical and agrees with S , while $S_{=h}$ and $S_{\leq h}$ are not canonical. The cells $S_{\leq h}$ and $S_{=h}$ agree and are represented by $T_{=h}$ where T is the face of S defined by the supporting hyperplane $f_S^{-1}(h)$.
3. If S contains elements (vertices) x^+ and x^- fulfilling $f_S(x^+) > h$ and $f_S(x^-) < h$, then $S_{=h}$, $S_{\leq h}$ and $S_{\geq h}$ are canonical and f_S is non-constant.
4. If $f_S(x) = h$ for all elements (vertices) $x \in S$, then f_S is constant and $S_{=h}$ is canonical and agrees with S , while $S_{\leq h}$ and $S_{\geq h}$ are not canonical. They also agree with S and are canonically represented by $S_{=h}$.
5. If $f_S(x) \leq h$ for all elements (vertices) $x \in S$ and for some elements (vertices) x^- and x_0 the relations $f_S(x^-) < h$ and $f_S(x_0) = h$ hold, in other words if f_S is non-constant and h its maximum on S , then $S_{\leq h}$ is canonical and agrees with S , while $S_{=h}$ and $S_{\geq h}$ are not canonical. The cells $S_{\geq h}$ and $S_{=h}$ agree and are represented by $T_{=h}$ where T is the face of S defined by the supporting hyperplane $f_S^{-1}(h)$.
6. If $f_S(x) < h$ for all elements (vertices) $x \in S$, then $S_{\leq h}$ is canonical and agrees with S , while $S_{=h}$ and $S_{\geq h}$ are not canonical. They are empty and represented by $\emptyset_{\in \emptyset}$. This case occurs for constant f_S and for non-constant f_S attaining a maximum on S .

Proof. One easily checks that the six cases are mutually exclusive and that the list covers all possible situations. Since case 5 is symmetric to case 2 and case 6 is symmetric to case 1, we only consider the first four cases. We can express the situation using the notation introduced for the characterisation of canonical terms from Lemma 2.24: On the one hand, we have the closed interval $J = f_S(S)$ of the values attained on S , on the other hand, the interval I is one of $\{h\}$, $(-\infty, h]$, or $[h, \infty)$. Note that f_S is constant if and only if J is a singleton. Now we decide for each case whether J and I intersect in their interior or not. In the former case the corresponding term is canonical by Lemma 2.24, in the latter case the term is not canonical and we describe which term represents the cell canonically instead.

The assumption of the first case means that J has a left endpoint and that h is strictly smaller than this endpoint. Then $\{h\}$ and $(-\infty, h]$ do not intersect J at all, meaning that $S_{=h}$ and $S_{\leq h}$ are empty and represented by $\emptyset_{\in \emptyset}$. For the interval $I = [h, \infty)$ however, I and J intersect in their interior, regardless of whether J is a singleton or not, in fact their intersection is J itself and any interior point of J is also interior to I , hence $S_{\geq h}$ is canonical and agrees with S .

When the second case applies, J has a left endpoint that agrees with h and is not a singleton. Then for sufficiently small ϵ , the value $h + \epsilon$ is interior to both $[h, \infty)$ and J , rendering $S_{\geq h}$ canonical. For the other two intervals $I = \{h\}$ or $I = (-\infty, h]$, the value h is the only value of I that is attained on S and it is not interior to J . Rather, $f_S^{-1}(h)$ is a hyperplane such that S is contained completely in one of its closed half-spaces, namely the half-space defined by the inequality $f_S(x) \geq h$. Furthermore, since the value h is attained on S , the hyperplane $f_S^{-1}(h)$ is a supporting hyperplane for S . The intersection of this supporting hyperplane with S is at the same time a proper face T of S , the

level set cell $S_{=h}$, and the level set cell $S_{\leq h}$. Therefore $T_{=h}$ is the canonical representation for $S_{=h}$ and $S_{\leq h}$. To avoid confusion regarding the variant for polytopal cells S with vertices in place of arbitrary points, let us point out that Table 2.5 does not contain an equivalence stating that S contains a point x_0 with value h if and only if it contains a vertex with value h , because such a statement is obviously wrong. The argument relies instead on the observation which we made already in the proof of Lemma 2.24: Under the assumption that $f_S(x) \geq h$ holds for all $x \in S$, the condition that $f_S(x_0) = h$ holds for some x_0 in S is equivalent to the condition that $f_S(x_0) \leq h$ holds for some x_0 in S , and for this condition an appropriate equivalence in Table 2.5 exists. This is again an instance of the basic fact from linear programming that a minimum is always attained at a vertex.

In the third case, the value h lies in the interior of an interval J that is not a singleton. For sufficiently small ϵ , the values $h - \epsilon$ and $h + \epsilon$ lie still in the interior of J , so that the value h is interior to J and $\{h\}$, the value $h - \epsilon$ is interior to J and $(-\infty, h]$, and the value $h + \epsilon$ is interior to J and $[h, \infty)$. This shows that all three terms $S_{=h}$, $S_{\leq h}$ and $S_{\geq h}$ are canonical in this case.

When the fourth case occurs, J agrees with $\{h\}$, so that the only interior point h of J is interior to $\{h\}$ but not to $(-\infty, h]$ or $[h, \infty)$. Therefore $S_{=h}$ is canonical, and since it agrees with S , $S_{\leq h}$, and $S_{\geq h}$, it is also the canonical representation for $S_{\leq h}$ and $S_{\geq h}$. \square

The next two lemmata are in principle just a reformulation of the previous result. We only rearrange the cases so that we obtain a convenient reference for the following situation: We are given a polyhedral complex M with a real-valued map f linear on cells, inducing for each cell S of M a linear map f_S on the affine hull of S . Furthermore we consider a real value h and ask for which cells S of M the term $S_{=h}$ is canonical. The first lemma lists the criteria for $S_{=h}$ being canonical and afterwards the criteria for $S_{=h}$ being non-canonical. The second lemma provides criteria in the same fashion for the case that we ask for the terms $S_{\leq h}$ or $S_{\geq h}$ instead of $S_{=h}$. Again we indicate in parentheses a variant of the result for bounded polyhedral sets where the word “element” is replaced by the word “vertex” at appropriate places using the equivalences from Table 2.5.

Lemma 2.26. *A term $S_{=h}$ canonically represents the cell $S_{=h}$ for a (bounded) polyhedral set S if and only if one of the following cases occurs:*

1. *The cell S is non-empty and all elements (vertices) x of S have the constant value $f_S(x) = h$.*

While maintaining the requirement to be non-empty, this case can be equivalently described by the condition that S lies completely in $f_S^{-1}(h)$ or by the condition that $f_S^{-1}(h)$ agrees with the whole affine hull of S . In this case the level set cell $S_{=h}$ agrees with the cell S and hence their dimensions $\dim(S_{=h}) = \dim(S)$ agree.

2. *The cell S contains at least one element (vertex) x^+ with $f_S(x^+) > h$ and at least one element vertex x^- with $f_S(x^-) < h$.*

An equivalent criterion for this case is the requirement that $f_S^{-1}(h)$ is a hyperplane that cuts S , that is, for each of the two sides of the hyperplane,

there are elements (vertices) of S lying strictly on that side. In this case we have $\dim(S_{=h}) = \dim(S) - 1$.

On the other hand, $S_{=h}$ is not a canonical representation if and only if one of the following cases occur

1. The cell S is empty.

In this case the canonical representation for the cell $S_{=h}$ is $\emptyset_{\in\emptyset}$.

2. All elements (vertices) x of S have values $f_S(x) \geq h$ and the inequality is strict for at least one element (vertex).
3. All elements (vertices) x of S have values $f_S(x) \leq h$ and the inequality is strict for at least one element (vertex).

In the last two cases the canonical representation for the cell $S_{=h}$ depends on whether the value h is attained on S or not: If the inequalities are strict for all elements (vertices), the intersection $S \cap f_S^{-1}(h)$ is empty, otherwise $f_S^{-1}(h)$ is a supporting hyperplane of S . In the former case $S_{=h}$ is empty and represented by $\emptyset_{\in\emptyset}$, in the latter case $S_{=h} = S \cap f_S^{-1}(h)$ is a proper face T of S with $T = T_{=h}$, hence it is canonically represented by $T_{=h}$.

Arranging the cases differently yields another criterion for $S_{=h}$ being non-canonical: Either S is empty, or the map f_S is constant with a value $c \neq h$, or f_S is non-constant on $\text{aff}(S)$ and the cell S lies completely in one of the closed half-spaces defined by the hyperplane $f_S^{-1}(h)$.

Proof. By the last criterion from Lemma 2.15, the term $S_{=h}$ is canonical if and only if f_S attains the value h in the interior of S . We proved in Lemma 2.21 that this condition is fulfilled if and only if one of the two stated cases occurs. The mentioned alternative criteria are easily verified as equivalent. The remark on the dimension is trivial in the first case. For the second case we use that if a hyperplane H intersects an interior point p of a full-dimensional convex set C , then there is an open ball B around p contained in C , and the fact that $H \cap B$ is an open ball in H shows that $H \cap C$ is a convex set of dimension $\dim(H) = \dim(C) - 1$. Choosing for H the hyperplane $f_S^{-1}(h)$ and for $p \in S = C$ an interior point of S with value $f_S(p) = h$ yields that $S_{=h} = H \cap S$ has dimension $\dim(S) - 1$. From the two cases that ensure a canonical representation, the first agrees with case 4 in Lemma 2.25, the second agrees with case 3 in Lemma 2.25.

On the other hand, the three suggested cases for non-canonical terms are exactly the three cases for h being not attained in the interior of S established in Corollary 2.22. The claim regarding the canonical representation for the cell $S_{=h}$ is trivial for the first case, and can be transferred from Lemma 2.25 for the other two cases, because the second case covers cases 1 and 2 of Lemma 2.25, and the third case covers cases 5 and 6 of Lemma 2.25.

The criterion suggested in the last paragraph can also be read off from the cases in Lemma 2.25: The second alternative corresponds to cases 1 and 6 for constant maps, the third alternative subsumes cases 1, 2, 5, and 6 for non-constant maps. \square

Lemma 2.27. *For a (bounded) polyhedral set S , the term $S_{\leq h}$ is canonical if and only if S contains an element (vertex) x^- such that $f_S(x^-) < h$. In this case the dimension of $S_{\leq h}$ agrees with the dimension of S .*

On the other hand, $S_{\leq h}$ is not canonical if and only if $f_S(x) \geq h$ for all elements (vertices) $x \in S$. If $S_{\leq h}$ is not canonical, then the canonical representation can be $S_{=h}$, $T_{=h}$ for a proper face T of S , or $\emptyset_{\in \emptyset}$. The first possible term $S_{=h}$ is the canonical representation of $S_{\leq h}$ if and only if f_S is constant with value h and S is non-empty. A term $T_{=h}$ is the canonical representation if and only if all elements (vertices) $x \in S$ have values $f_S(x) \geq h$ and S contains elements (vertices) x^+ and x_0 with $f_S(x^+) > h$ and $f_S(x_0) = h$. The appropriate cell T for this case is the face of S defined by the intersection of S with the supporting hyperplane $f_S^{-1}(h)$. The term $\emptyset_{\in \emptyset}$ is the canonical representation if and only if $f_S(x) > h$ for all elements (vertices) $x \in S$.

For the symmetric case of a term $S_{\geq h}$ analogous statements hold: The term $S_{\geq h}$ is canonical if and only if S contains an element (vertex) x^+ such that $f_S(x^+) > h$. The dimension of $S_{\geq h}$ agrees with the dimension of S in this case. Contrariwise, $S_{\geq h}$ is not canonical if and only if $f_S(x) \geq h$ for all elements (vertices) $x \in S$. The cell $S_{\geq h}$ is canonically represented by the term $S_{=h}$ if and only if $S \neq \emptyset$ and f_S is constant with value h ; by a term $T_{=h}$ if and only if all elements (vertices) $x \in S$ have values $f_S(x) \leq h$ and S contains elements (vertices) x^- and x_0 with $f_S(x^-) < h$ and $f_S(x_0) = h$ and if T is the face defined by the supporting hyperplane $f_S^{-1}(h)$; and by the term $\emptyset_{\in \emptyset}$ if and only if $f_S(x) < h$ for all elements (vertices) $x \in S$.

Proof. Since the term $S_{\geq h}$ can be treated in an completely symmetric way, we only consider the term $S_{\leq h}$. For proving the claimed criterion for canonical terms, we apply again the last equivalence from Lemma 2.15, which states in this case that $S_{\leq h}$ is canonical if and only if f_S attains a value in the interior of $(-\infty, h]$ at an interior point of S . But values in the interior of $(-\infty, h]$ are real numbers strictly smaller than h . So we can see immediately that $f_S(x^-) < h$ for some $x^- \in S$ is a necessary condition. For verifying that the condition is also sufficient, we only have to justify why the requirement that x^- is an interior point of S can be dropped. This is due to continuity and to the fact that any neighbourhood of a point in a convex set contains also interior points of the convex set: If x^- is an arbitrary point of S with $f_S(x^-) < h$, then in a sufficiently small neighbourhood of x^- , all values attained by f_S are still less than h , and an interior point of S lying in that neighbourhood is the desired interior point of S with value less than h . Note that the criterion just proven covers the cases 3, 5, and 6 from Lemma 2.25.

The remark on the dimension relies on continuity as well: Consider an interior point x^- of S with $f_S(x^-) < h$. Then all points in a sufficiently small ball around x^- still lie in S and have a value less than h , proving that $S_{\leq h}$ contains a ball of the same dimension as S .

The stated criterion for $S_{\leq h}$ being non-canonical is obviously the negation of the criterion for being canonical. The claimed canonical representations for the cell $S_{\leq h}$ arise from cases 1, 2, and 4 from Lemma 2.25. \square

Level Values Not Attained at Vertices. When we assume that the cell S is polytopal and that the considered value h is not attained at vertices, the situation is simplified considerably. The next lemma shows that if this assumption applies to all cells of a complex M , then the empty cell is the only cell with non-canonical representations in the level set complexes $M_{=h}$, $M_{\leq h}$, and $M_{\geq h}$.

Lemma 2.28. *Assume that M is a polytopal complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells and that the real value h fulfils $f(v) \neq h$ for all vertices v of M . Let S be a cell of M . Then $S_{=h}$, $S_{\leq h}$, and $S_{\geq h}$ are canonical representations if and only if the respective set is non-empty. For $S_{=h}$ this is again equivalent to S containing two vertices v^+ and v^- with $f(v^+) > h$ and $f(v^-) < h$. Furthermore, for $X, Y \in \{(-\infty, h], \{h\}, [h, \infty)\}$ and an additional cell T of M , the cell $T_{\in X}$ is a face of $S_{\in Y}$ if and only if $T_{\in X}$ is empty or T is a face of S and X is a face of Y .*

Proof. For the first claim we show the contraposition that $S_{=h}$, $S_{\leq h}$, and $S_{\geq h}$ are not canonical if and only if they represent the empty cell. Since by assumption $f^{-1}(h)$ does not contain vertices, the cases where $f_S^{-1}(h)$ is a supporting hyperplane of S or contains S disappear for polytopal cells S . This means that the cases 2, 4, and 5 from Lemma 2.25 do not apply, when we ask whether $S_{=h}$, $S_{\leq h}$, or $S_{\geq h}$ are canonical or not. In the remaining cases these terms are non-canonical if and only if they represent the empty cell. When the term $S_{=h}$ is considered, case 3 from Lemma 2.25 is the only remaining case where $S_{=h}$ is canonical. The stated equivalent condition for $S_{=h}$ is exactly the criterion given there for that case. The last assertion follows from Lemma 2.10 because $T_{\in X}$ is either empty as a cell or canonical as a term. \square

Bounded Non-Singleton Intervals. The only remaining type of closed intervals which we did not consider in detail yet consists of the intervals $[a, b]$ with two endpoints. First, we give a criterion characterising the canonical ones among the terms of the form $S_{\in [a, b]}$, and we list which terms under which condition canonically represent a cell $S_{\in [a, b]}$ if the term $S_{\in [a, b]}$ is not canonical, just as we did for $S_{=h}$, $S_{\leq h}$, and $S_{\geq h}$ in Lemmata 2.26 and 2.27. Again a variant for polytopal cells with occurrences of “element” being replaced by “vertex” is indicated by parentheses. Afterwards we state an analogue of Lemma 2.28 for a level set complex $M_{\in [a, b]}$ where M is a polytopal complex and the values a and b are not attained at vertices of M .

Lemma 2.29. *For a (bounded) polyhedral set S and real numbers $a < b$, the term $S_{\in [a, b]}$ is canonical if and only if S contains elements (vertices) x and y with $f_S(x) > a$ and $f_S(y) < b$. In this case the dimensions of S and $S_{\in [a, b]}$ agree.*

If $S_{\in [a, b]}$ is not canonical, then the canonical representation is one of the following:

1. $S_{=a}$ if and only if S is non-empty and f_S is constant with value a .
2. $S_{=b}$ if and only if S is non-empty and f_S is constant with value b .
3. $T_{=a}$ for a proper face T of S if and only if $f_S(x) \leq a$ for all elements (vertices) $x \in S$ and equality and strict inequality are both attained, with T being the face of S defined by the intersection of S with the supporting hyperplane $f_S^{-1}(a)$.
4. $T_{=b}$ for a proper face T of S if and only if $f_S(x) \geq b$ for all elements (vertices) $x \in S$ and equality and strict inequality are both attained, with T being the face of S defined by the intersection of S with the supporting hyperplane $f_S^{-1}(b)$.

5. $\emptyset_{\in \emptyset}$ if and only if $f_S(x) < a$ for all elements (vertices) $x \in S$ or $f_S(x) > b$ for all elements (vertices) $x \in S$.

Proof. The claims can be shown with arguments similar to those used in the proof of Lemma 2.27 for terms of the form $S_{\leq h}$. Applying once again the last equivalence from Lemma 2.15 yields that a term $S_{\in [a,b]}$ is canonical if and only if some interior point of S is mapped to an interior point of $[a, b]$, in other words $a < f_S(x) < b$ for some interior point of S . This shows that the stated criterion is necessary for $S_{\in [a,b]}$ being canonical. Now assume for the other direction that $f_S(x) > a$ and $f_S(y) < b$ hold. As a first step, one easily checks by case distinction depending on which order relations additionally hold between $f_S(x)$, $f_S(y)$, a , and b that the interval between the values $f_S(x)$ and $f_S(y)$ contains a value h in the interior of the interval $[a, b]$. This value h in the interval between $f_S(x)$ and $f_S(y)$ is attained at some point p of S by convexity of S and continuity of f_S . Using again continuity of f_S and the fact that arbitrary small neighbourhoods of points in S contain interior points of S implies that the value of some interior point of S sufficiently close to p still lies in the interior of $[a, b]$. This ensures that $S_{\in [a,b]}$ is canonical by the last criterion from Lemma 2.15. Furthermore an interior point of S with value in the interior of $[a, b]$ has a neighbourhood that is contained completely in $S_{\in [a,b]}$ by continuity, which shows $\dim(S_{\in [a,b]}) = \dim(S)$ whenever the term $S_{\in [a,b]}$ is canonical.

For non-canonical terms $S_{\in [a,b]}$, one easily checks that the five cases in the list are mutually disjoint and cover all remaining alternatives. Furthermore, in each case the claimed term obviously represents the same cell as $S_{\in [a,b]}$ and neither the domain nor the level range used for the term can be replaced by smaller faces, showing that the term is the canonical representation of the cell. \square

Lemma 2.30. *Assume that M is a polytopal complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells and that the real values $a < b$ fulfil $f(v) \notin \{a, b\}$ for all vertices v of M . Let S be a cell of M . For any cell S of M , the term $S_{\in [a,b]}$ is a canonical representation of its cell if and only if $S_{\in [a,b]}$ is non-empty. Furthermore, for $X, Y \in \{\{a\}, [a, b], \{b\}\}$ and cells S and T of M , the level set cell $T_{\in X}$ is a face of $S_{\in Y}$ if and only if $T_{\in X}$ is empty or T is a face of S and X is a face of Y .*

Proof. Under the assumption that S is polytopal and non-empty, all five cases from Lemma 2.29 for non-canonical terms $S_{\in [a,b]}$ require that at least one of the values a and b is attained at a vertex of S . The claim concerning the face relation follows again from Lemma 2.10. \square

Summary in First-Order Logic Style. Note that the Lemmata 2.26, 2.27, and 2.29 provide for each cell $S_{\in Y}$ of a level set complex, with Y being one of $\{h\}$, $(-\infty, h]$, $[h, \infty)$, or $[a, b]$, a criterion for a term $T_{\in X}$ being the canonical representation of $S_{\in Y}$. When the cell S is a convex polytope, then they even provide a criterion in terms of the values that the vertices of S attain. In preparation for a proof in the next section, we collect all these criteria in Table 2.6, expressing them in a way that we can treat them just like formulas of first-order logic. We introduce a quaternary predicate $\phi(S, Y, T, X)$ that denotes the statement that $T_{\in X}$ is the canonical representation of the cell $S_{\in Y}$, where S and T are cells of M , and Y and X are closed intervals. Allowing ourselves another slight abuse of notation, we abbreviate the predicate by $\phi(S_{\in Y}, T_{\in X})$.

Corollary 2.31. *For deciding whether a term $T_{\in X}$ is the canonical representation of the cell represented by $S_{\in Y}$, the previous results provide the criteria listed in Table 2.6. The property that $T_{\in X}$ is the canonical representation of $S_{\in Y}$ is expressed by the predicate $\phi(S_{\in Y}, T_{\in X})$. On the left hand side of the table, we list instances of the predicate, distinguishing cases into certain patterns according to the types of the closed intervals X and Y and their mutual face and order relations and according to the mutual face relations between S and T . Here the term “order relations” refers to the distinction made when Y is an interval $[a, b]$ and X corresponds to one of its endpoints, namely whether the corresponding endpoint is a left or a right endpoint. The listed patterns cover all possible cases where S and T are cells of the considered complex M , and X and Y are closed intervals. Just as Lemmata 2.26, 2.27, and 2.29 were stated in two variants, the table allows two variants of interpretation. The expression $v \in S$ occurring several times on the right hand side in bounded quantifiers and in set builder notation can be interpreted in two meaningful ways: The literal interpretation considers all points of the cell S as the range for the variable v . With this interpretation, the list is correct for all polyhedral complexes M . When we only allow polytopal complexes M , we can alternatively assume that the variable v ranges only over vertices of S , as suggested by the choice of the letter v instead of the letter x . Note that in the expression $\text{conv}\{v \in S \mid f(v) = h\}$, taking the convex hull is only necessary for the latter interpretation; for the former it is unnecessary but still correct.*

Proof. The cases $Y = \emptyset$ and $Y = \mathbb{R}$ are trivial, the remaining criteria are shown in Lemmata 2.26, 2.27, and 2.29. \square

2.3 Combinatorial Equivalences and Collapses

Now that we have a suitable description of the combinatorial structure of a level set complex based on the notion of canonical representation, we can prove for certain pairs of level set complexes that they are combinatorially equivalent. In this section, we collect some results of this kind. They will be useful later for proving certain level sets as PL-homeomorphic because combinatorially equivalent polytopal complexes have PL-homeomorphic domains by Theorem 1.31. In other parts of this section we study collapses that can be performed on level set complexes. Since collapses induce homotopy equivalences, the results yield convenient tools for establishing homotopy equivalences between level sets and other spaces. Throughout the section, we assume that M is a polytopal complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Clearly, the level ranges in this situation are interval complexes, in most cases complexes associated with a single closed interval I .

2.3.1 Natural Combinatorial Equivalences between Level Set Complexes

The first topic addresses natural combinatorial equivalences between level set complexes $M_{\in I_1}$ and $M_{\in I_2}$ that arise when one interval serving as level range is obtained from the other by moving the endpoints slightly as long as they avoid values that are attained at vertices. For example, if no vertex has a value in

Table 2.6: A list of criteria for the predicate $\phi(S_{\in Y}, T_{\in X})$ expressing that $T_{\in X}$ canonically represents $S_{\in Y}$. The first line covers some obviously false cases. For the other lines we assume $T \neq S$ since the case $S = T$ is covered explicitly in separate lines. Predicates $\phi(S_{=c}, T_{=c})$ with $c \in \{a, b, h\}$ referenced on the right hand side as sub-formula can be replaced using the formula for $\phi(S_{=h}, T_{=h})$.

If $T \notin \mathcal{F}(S)$, or $X \notin \mathcal{F}(Y)$, or $T = \emptyset$ and $X \neq \emptyset$, or $T \neq \emptyset$ and $X = \emptyset$:

$$\phi(S_{\in Y}, T_{\in X}) \iff \text{false}$$

Otherwise, for $T \neq S$

$$\begin{aligned} \phi(S_{\in \emptyset}, \emptyset_{\in \emptyset}) &\iff \text{true} \\ \phi(S_{=h}, S_{=h}) &\iff (S \neq \emptyset \wedge \forall v \in S: f(v) = h) \\ &\quad \vee (\exists v \in S: f(v) < h \wedge \exists v \in S: f(v) > h) \\ \phi(S_{=h}, T_{=h}) &\iff ((\forall v \in S: f(v) \leq h \wedge \exists v \in S: f(v) < h) \\ &\quad \vee (\forall v \in S: f(v) \geq h \wedge \exists v \in S: f(v) > h)) \\ &\quad \wedge \exists v \in S: f(v) = h \wedge T = \text{conv}\{v \in S \mid f(v) = h\} \\ \phi(S_{=h}, \emptyset_{\in \emptyset}) &\iff \forall v \in S: f(v) > h \vee \forall v \in S: f(v) < h \\ \phi(S_{\leq h}, S_{\leq h}) &\iff \exists v \in S: f(v) < h \\ \phi(S_{\geq h}, S_{\geq h}) &\iff \exists v \in S: f(v) > h \\ \phi(S_{\leq h}, T_{\leq h}) &\iff \text{false} \\ \phi(S_{\geq h}, T_{\geq h}) &\iff \text{false} \\ \phi(S_{\leq h}, S_{=h}) &\iff S \neq \emptyset \wedge \forall v \in S: f(v) = h \\ \phi(S_{\geq h}, S_{=h}) &\iff S \neq \emptyset \wedge \forall v \in S: f(v) = h \\ \phi(S_{\leq h}, T_{=h}) &\iff \forall v \in S: f(v) \geq h \wedge \phi(S_{=h}, T_{=h}) \\ \phi(S_{\geq h}, T_{=h}) &\iff \forall v \in S: f(v) \leq h \wedge \phi(S_{=h}, T_{=h}) \\ \phi(S_{\leq h}, \emptyset_{\in \emptyset}) &\iff \forall v \in S: f(v) > h \\ \phi(S_{\geq h}, \emptyset_{\in \emptyset}) &\iff \forall v \in S: f(v) < h \\ \phi(S_{\in [a,b]}, S_{\in [a,b]}) &\iff \exists v \in S: f(v) > a \wedge \exists v \in S: f(v) < b \\ \phi(S_{\in [a,b]}, T_{\in [a,b]}) &\iff \text{false} \\ \phi(S_{\in [a,b]}, S_{=a}) &\iff S \neq \emptyset \wedge \forall v \in S: f(v) = a \\ \phi(S_{\in [a,b]}, S_{=b}) &\iff S \neq \emptyset \wedge \forall v \in S: f(v) = b \\ \phi(S_{\in [a,b]}, T_{=a}) &\iff \forall v \in S: f(v) \leq a \wedge \phi(S_{=a}, T_{=a}) \\ \phi(S_{\in [a,b]}, T_{=b}) &\iff \forall v \in S: f(v) \geq b \wedge \phi(S_{=b}, T_{=b}) \\ \phi(S_{\in [a,b]}, \emptyset_{\in \emptyset}) &\iff \forall v \in S: f(v) > b \vee \forall v \in S: f(v) < a \\ \phi(S_{\in \mathbb{R}}, S_{\in \mathbb{R}}) &\iff S \neq \emptyset \\ \phi(S_{\in \mathbb{R}}, T_{\in \mathbb{R}}) &\iff \text{false} \\ \phi(S_{\in \mathbb{R}}, \emptyset_{\in \emptyset}) &\iff S = \emptyset \end{aligned}$$

the interval $[h_1, h_2]$, then $M_{=h_1}$ and $M_{=h_2}$ are combinatorially equivalent by the assignment $S_{=h_1} \mapsto S_{=h_2}$. Since the cells in a level set complex can be identified without double-counting with the canonical terms for level set cells, we could restrict the assignment to canonical terms to avoid ambiguity. But it turns out that this is unnecessary since the assignment by the natural equivalence is independent of the representation of the cell if the two level ranges fulfil the requirements.

Monotonically Combinatorially Equivalent Interval Complexes. In order to describe the sufficient requirement for pairs (I_1, I_2) of intervals that ensures that $M_{\in I_1}$ and $M_{\in I_2}$ are combinatorially equivalent in a natural way, we have to refine our notion of combinatorial equivalence for interval complexes. The notion from Definition 1.26 considers half-bounded intervals $(-\infty, h_1]$ and $[h_2, \infty)$ bounded in different directions as equivalent, but even if $h_1 = h_2$, we cannot expect any combinatorial equivalence between $M_{\leq h_1}$ and $M_{\geq h_2}$. Furthermore it allows for combinatorial equivalences that switch the left and right endpoints of corresponding bounded intervals, that is, $[a_1, b_1]$ is mapped to $[a_2, b_2]$, but $\{a_1\}$ to $\{b_2\}$, and $\{b_1\}$ to $\{a_2\}$. For avoiding these undesirable effects of the general definition, we introduce the notion of an order-respecting combinatorial equivalence for interval complexes.

Definition 2.32 (order-respecting combinatorial equivalence). Two interval complexes \mathcal{I} and \mathcal{J} are called *order-respecting combinatorially equivalent* or *monotonically combinatorially equivalent* if there is a combinatorial equivalence $\Psi: \mathcal{I} \rightarrow \mathcal{J}$ in the sense of Definition 1.26 between them such that the following conditions are fulfilled:

1. The combinatorial equivalence assigns only half-bounded intervals of the same type to each other; in other words, it respects the six kinds of intervals listed in Table 1.1.
2. Its restriction to the vertices respects the natural order on \mathbb{R} as follows: for any x and y with $\{x'\} = \Psi(\{x\})$ and $\{y'\} = \Psi(\{y\})$ the equivalence $x \leq_{\mathbb{R}} y \iff x' \leq_{\mathbb{R}} y'$ holds.

It is not hard to see that the combinatorial equivalence $\Psi: \mathcal{I} \rightarrow \mathcal{J}$ required in the definition is uniquely determined if such an equivalence exists. Hence it is justified to call Ψ *the* order-respecting combinatorial equivalence between the interval complexes.

As usual, we are mainly interested in the case that the considered pair of level ranges consists of the interval complexes $\mathcal{F}(I_1)$ and $\mathcal{F}(I_2)$ associated with closed intervals I_1 and I_2 different from \emptyset and \mathbb{R} . Table 2.7 lists the assignments made by the order-respecting combinatorial equivalence Ψ for such pairs of level ranges.

It remains to formalise the requirement for the intervals I_1 and I_2 that one can be obtained from the other by shifting the endpoints such that the shift does not pass across a value attained at some vertex of M . We suggest the following definition that works for arbitrary interval complexes:

Definition 2.33. Let M be a polytopal complex with a map $f: M \rightarrow \mathbb{R}$ linear on cells. Two monotonically combinatorially equivalent interval complexes \mathcal{I}

Table 2.7: The order-respecting combinatorial equivalence $\Psi: \mathcal{F}(I_1) \rightarrow \mathcal{F}(I_2)$ for the typically considered pairs of closed intervals I_1 and I_2 . For each pair, a table of values listing the assignments made by Ψ is given. The trivial assignment $\Psi: \emptyset \mapsto \emptyset$ is omitted in all lists. For the last row, $a_i \neq b_i$ is assumed.

I_1	I_2	table of values for Ψ		
$\{h_1\}$	$\{h_2\}$	$\{h_1\} \mapsto \{h_2\}$		
$(-\infty, h_1]$	$(-\infty, h_2]$	$\{h_1\} \mapsto \{h_2\}$	$(-\infty, h_1] \mapsto (-\infty, h_2]$	
$[h_1, \infty)$	$[h_2, \infty)$	$\{h_1\} \mapsto \{h_2\}$	$[h_1, \infty) \mapsto [h_2, \infty)$	
$[a_1, b_1]$	$[a_2, b_2]$	$\{a_1\} \mapsto \{a_2\}$	$\{b_1\} \mapsto \{b_2\}$	$[a_1, b_1] \mapsto [a_2, b_2]$

and \mathcal{J} are called *combinatorially equivalent for M and f* , if the order-respecting combinatorial equivalence $\Psi: \mathcal{I} \rightarrow \mathcal{J}$ fulfils the following condition for any pair of corresponding vertices $x \in \text{Vrt}(\mathcal{I})$ and $x' \in \text{Vrt}(\mathcal{J})$, that is $\Psi(\{x\}) = \{x'\}$: If $x <_{\mathbb{R}} x'$, then no vertex of M has an f -value in $[x, x']$, and if $x' <_{\mathbb{R}} x$, then no vertex of M has an f -value in $[x', x]$. We also call two closed intervals I_1 and I_2 combinatorially equivalent for M and f , if their associated complexes $\mathcal{F}(I_1)$ and $\mathcal{F}(I_2)$ are.

Applying this definition to the typically considered pairs of intervals, we obtain the following requirements for I_1 and I_2 being combinatorially equivalent for M and f : For the first three cases in Table 2.7 it is required that no vertex v of M has a value $f(v)$ in the interval $[h_1, h_2]$ or $[h_2, h_1]$, assuming the two values h_1 and h_2 are distinct. For the last case it is required that no vertex v of M has a value $f(v)$ in the interval $[a_1, a_2]$ or $[a_2, a_1]$, unless $a_1 = a_2$, and that no vertex has a value in the interval $[b_1, b_2]$ or $[b_2, b_1]$, unless $b_2 = b_1$.

Transferring Combinatorial Equivalences from Interval Complexes to Level Set Complexes. Now we want to show that if two interval complexes \mathcal{I} and \mathcal{J} are combinatorially equivalent for M and f , then the level set complexes $M_{\in \mathcal{I}}$ and $M_{\in \mathcal{J}}$ are combinatorially equivalent in a natural way, because the equivalence is induced by the order-respecting combinatorial equivalence of \mathcal{I} and \mathcal{J} . We observed already in Lemma 2.10 that the combinatorial structure of a level set complex is uniquely determined by the combinatorial structure of domain and level range and by which terms are canonical. Since M is combinatorially equivalent to itself and, by assumption, \mathcal{I} and \mathcal{J} are monotonically combinatorially equivalent, we can bijectively assign to each pair (S, Y) in the Cartesian product of M and \mathcal{I} a corresponding pair $(S, \Psi(Y))$ in the Cartesian product of M and \mathcal{J} , where $\Psi: \mathcal{I} \rightarrow \mathcal{J}$ denotes the order-respecting combinatorial equivalence. Thus, for any term $S_{\in Y}$ representing a cell in $M_{\in \mathcal{I}}$ there is a corresponding term $S_{\in \Psi(Y)}$ representing a cell in $M_{\in \mathcal{J}}$. Lemma 2.10 states that the face poset of a level set complex is isomorphic to the product order of the posets of domain and level range restricted to the pairs forming canonical terms, with the isomorphism assigning to each such pair the cell represented by the term. Now if for all S and Y the term $S_{\in Y}$ is canonical if and only if $S_{\in \Psi(Y)}$ is canonical, the cited lemma asserts that $M_{\in \mathcal{I}}$ and $M_{\in \mathcal{J}}$ are com-

binatorially equivalent with the combinatorial equivalence being given by the natural bijection assigning to a cell of $M_{\in \mathcal{I}}$ with canonical representation $S_{\in Y}$ the cell of $M_{\in \mathcal{J}}$ with canonical representation $S_{\in \Psi(Y)}$. But we can prove even a bit more for such an assignment: If \mathcal{I} and \mathcal{J} are combinatorially equivalent for M and f , then $T_{\in X}$ is the canonical representation for $S_{\in Y}$ if and only if $T_{\in \Psi(X)}$ is the canonical representation for $S_{\in \Psi(Y)}$, showing that the assignment $S_{\in Y} \mapsto S_{\in \Psi(Y)}$ does not depend on the choice of the representing term for the cell $S_{\in Y}$ and yields a well-defined combinatorial equivalence from $M_{\in \mathcal{I}}$ to $M_{\in \mathcal{J}}$ even if not restricted to canonical terms explicitly.

For proving this claim, recall that we introduced a predicate $\phi(S, Y, T, X)$, abbreviated by $\phi(S_{\in Y}, T_{\in X})$, for the statement that $T_{\in X}$ is the canonical representation of the cell $S_{\in Y}$ and gave a list of suitable criteria in terms of the values attained at vertices in Table 2.6, distinguishing all possible cases where X and Y are closed intervals. It is not hard, but somewhat tedious to check for all cases that if \mathcal{I} and \mathcal{J} are combinatorially equivalent for M and f by Ψ , then $\phi(S, Y, T, X)$ is equivalent to $\phi(S, \Psi(Y), T, \Psi(X))$. We can reduce the amount of work a bit when we treat the criteria just like formulas of first-order logic and decompose them into their atomic formulas.

Lemma 2.34. *Let \mathcal{I} and \mathcal{J} be combinatorially equivalent interval complexes for a polytopal complex M and a map $f: |M| \rightarrow \mathbb{R}$ linear on cells by the order-respecting combinatorial equivalence $\Psi: \mathcal{I} \rightarrow \mathcal{J}$. Then for all cells S and T in M and all closed intervals X and Y in \mathcal{I} , the term $T_{\in X}$ is the canonical representation for $S_{\in Y}$ in $M_{\in \mathcal{I}}$ if and only if $T_{\in \Psi(X)}$ is the canonical representation for $S_{\in \Psi(Y)}$ in $M_{\in \mathcal{J}}$.*

Proof. Expressing the claim in terms of the predicate ϕ , our aim is to prove that $\phi(S, Y, T, X)$ is equivalent to $\phi(S, \Psi(Y), T, \Psi(X))$, or written in short form: $\phi(S_{\in Y}, T_{\in X}) \iff \phi(S_{\in \Psi(Y)}, T_{\in \Psi(X)})$. The cases listed in Table 2.6 cover all kinds of closed intervals that can occur in an interval complex. The order-respecting combinatorial equivalence asserts that $X \sqsubseteq Y$ holds if and only if $\Psi(X) \sqsubseteq \Psi(Y)$ does, and that X is empty if and only if $\Psi(X)$ is. This proves that the obviously false case in the first line of Table 2.6 occurs for the terms $S_{\in Y}$ and $T_{\in X}$ if and only if it occurs for the terms $S_{\in \Psi(Y)}$ and $T_{\in \Psi(X)}$. Similarly, the fact that the order-respecting combinatorial equivalence respects the type of the interval and the order of the endpoints ensures that for each line of Table 2.6, the pattern on the left hand side matches the types of the intervals Y and X and their mutual face and order relations if and only if it matches the type of the intervals $\Psi(Y)$ and $\Psi(X)$ and their mutual face and order relations. Therefore it suffices to check that the truth value of the corresponding right hand sides remains unchanged when X is replaced by $\Psi(X)$ and Y by $\Psi(Y)$.

Since we assume M to be polytopal, we can use the interpretation of the table in which the variable v ranges only over vertices of S . We treat the criteria on the right hand side just like formulas of first-order logic and decompose them into their atomic formulas. This way we obtain, apart from formulas such as $S = \emptyset$ that do not depend on the level range, only formulas of the pattern $f(v) R h$, where R is one of the binary relations $=, <, >, \leq, \geq$ between the value $f(v)$ of some vertex v of M and some value h such that $\{h\}$ is a vertex of the level range, and the formula $T = \text{conv}\{v \in S \mid f(v) = h\}$ for some vertex $\{h\}$ of the level range. We could even reduce the number of patterns by expressing some of the binary relations as a boolean combination in terms of the other ones,

if we like. Now we show first that a replacement of X by $\Psi(X)$ and of Y by $\Psi(Y)$ does not affect the truth value of the atomic formulas. When we perform such a replacement and look at its effect on the atomic formulas of the patterns $f(v) R h$ and $T = \text{conv}\{v \in S \mid f(v) = h\}$, the replacement will only change the value substituted for h from some value h_1 with $\{h_1\} \in \text{Vrt}(\mathcal{I})$ to some value h_2 such that $\{h_2\} = \Psi(\{h_1\}) \in \text{Vrt}(\mathcal{J})$. If $h_1 = h_2$, the atomic formulas $f(v) R h_1$ and $f(v) R h_2$ as well as the formulas $T = \text{conv}\{v \in S \mid f(v) = h_1\}$ and $T = \text{conv}\{v \in S \mid f(v) = h_2\}$ are trivially equivalent. Otherwise, the assumption that \mathcal{I} and \mathcal{J} are combinatorially equivalent for M and f asserts that no vertex of M has an f -value in the interval $[h_1, h_2]$ or $[h_2, h_1]$, whichever of the two values h_1 and h_2 is the smaller one. Under this assertion in turn, the atomic formulas $f(v) R h_1$ and $f(v) R h_2$ are obviously equivalent for any vertex v and any of the five relations R , and furthermore $\text{conv}\{v \in S \mid f(v) = h_1\}$ and $\text{conv}\{v \in S \mid f(v) = h_2\}$ are both empty, meaning that the formulas $T = \text{conv}\{v \in S \mid f(v) = h_1\}$ and $T = \text{conv}\{v \in S \mid f(v) = h_2\}$ are equivalent as well. This proves that the truth value of the atomic formulas is not affected by the replacement. Since the truth value of the whole formula does only depend on the truth values of the atomic formulas, we have established our claim that replacing X and Y by $\Psi(X)$ and $\Psi(Y)$ does not affect the truth values on the right hand sides. Therefore we have $\phi(S_{\in Y}, T_{\in X}) \iff \phi(S_{\in \Psi(Y)}, T_{\Psi(X)})$ as desired. \square

Now we can prove that interval complexes that are combinatorially equivalent for M and f induce level set complexes on M that are combinatorially equivalent in a natural way.

Lemma 2.35. *Let M be a polytopal complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. If \mathcal{I} and \mathcal{J} are combinatorially equivalent level ranges for M and f , then $M_{\in \mathcal{I}}$ and $M_{\in \mathcal{J}}$ are combinatorially equivalent polytopal complexes. The natural equivalence is given by the assignment $S_{\in Y} \mapsto S_{\in \Psi(Y)}$ for the order-respecting combinatorial equivalence Ψ mapping the intervals of \mathcal{I} to the intervals of \mathcal{J} . Under this assignment $S_{\in Y}$ is canonical if and only if $S_{\in \Psi(Y)}$ is canonical. Furthermore the assignment is independent of the representation of the cell $S_{\in Y}$ by a cell S of M and a face Y of \mathcal{I} and hence is well-defined.*

More specific instances of this result for the typically considered intervals are collected in Table 2.8. The trivial assignments made by the corresponding combinatorial equivalence are listed in Table 2.9.

Proof. We first show that the assignment $S_{\in Y} \mapsto S_{\in \Psi(Y)}$ for $Y \in \mathcal{I}$ is well-defined. Assume that $S_{\in Y_1}$ and $S_{\in Y_2}$ represent the same cell and this cell is canonically represented by $T_{\in X}$ in $M_{\in \mathcal{I}}$. By Lemma 2.34, this implies that $T_{\in \Psi(X)}$ is the canonical representation in $M_{\in \mathcal{J}}$ of both $S_{\in \Psi(Y_1)}$ and $S_{\in \Psi(Y_2)}$. Hence, $S_{\in \Psi(Y_1)}$ and $S_{\in \Psi(Y_2)}$ represent the same cell and the assignment $S_{\in Y} \mapsto S_{\in \Psi(Y)}$ yields a well-defined map. For the remainder of the proof, let $\Phi: M_{\in \mathcal{I}} \rightarrow M_{\in \mathcal{J}}$ with $\Phi(S_{\in Y}) = S_{\in \Psi(Y)}$ denote this natural map.

Switching the roles of \mathcal{I} and \mathcal{J} and using the same proof for the inverse map Ψ^{-1} , which is an order-respecting combinatorial equivalence for M and f as well, we obtain that the assignment $S_{\in Z} \mapsto S_{\in \Psi^{-1}(Z)}$ for $Z \in \mathcal{J}$ yields a well-defined map $\Phi': M_{\in \mathcal{J}} \rightarrow M_{\in \mathcal{I}}$. The assignment rules immediately imply the equalities $\Phi'(\Phi(S_{\in Y})) = S_{\in \Psi^{-1}(\Psi(Y))} = S_{\in Y}$ and $\Phi(\Phi'(S_{\in Z})) = S_{\in \Psi(\Psi^{-1}(Z))} = S_{\in Z}$, so that Φ is proven to be bijective with its inverse Φ^{-1} being Φ' .

Table 2.8: For a polytopal complex M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells such that for all vertices v of M the value $f(v)$ does not lie in one of the intervals $[h_1, h_2]$, $[a_1, a_2]$, $[b_1, b_2]$, each of the following pairs of level set complexes is a pair of combinatorially equivalent complexes, with the equivalence denoted by the symbol $\stackrel{\text{comb}}{\equiv}$.

$$\begin{array}{ccc} M_{=h_1} \stackrel{\text{comb}}{\equiv} M_{=h_2} & M_{\leq h_1} \stackrel{\text{comb}}{\equiv} M_{\leq h_2} & M_{\geq h_1} \stackrel{\text{comb}}{\equiv} M_{\geq h_2} \\ M_{\in[a, b_1]} \stackrel{\text{comb}}{\equiv} M_{\in[a, b_2]} & M_{\in[a_1, b]} \stackrel{\text{comb}}{\equiv} M_{\in[a_2, b]} & M_{\in[a_1, b_1]} \stackrel{\text{comb}}{\equiv} M_{\in[a_2, b_2]} \end{array}$$

Table 2.9: Under the assumptions made for Table 2.8, the natural combinatorial equivalences between the level set complexes is described completely by the following assignments.

$$\begin{array}{ccccccc} S_{=h_1} \mapsto S_{=h_2} & S_{\leq h_1} \mapsto S_{\leq h_2} & S_{\geq h_1} \mapsto S_{\geq h_2} & & & & \\ S_{\in[a, b_1]} \mapsto S_{\in[a, b_2]} & S_{\in[a_1, b]} \mapsto S_{\in[a_2, b]} & S_{\in[a_1, b_1]} \mapsto S_{\in[a_2, b_2]} & & & & \\ S_{=a_1} \mapsto S_{=a_2} & S_{=b_1} \mapsto S_{=b_2} & S_{=a} \mapsto S_{=a} & S_{=b} \mapsto S_{=b} & \emptyset_{\in\emptyset} \mapsto \emptyset_{\in\emptyset} & & \end{array}$$

It remains to show that Φ is a combinatorial equivalence in terms of cells. Since $M_{\in\mathcal{I}}$ and $M_{\in\mathcal{J}}$ are both polytopal complexes for polytopal M , it suffices to check that Φ respects the face relation, that is $T_{\in X} \sqsubseteq S_{\in Y}$ holds in $M_{\in\mathcal{I}}$ if and only if $\Phi(T_{\in X}) \sqsubseteq \Phi(S_{\in Y})$ holds in $M_{\in\mathcal{J}}$. For showing the direction going from $M_{\in\mathcal{I}}$ to $M_{\in\mathcal{J}}$, we can assume that $T_{\in X}$ is canonical, since any cell has a canonical representation. This implies by Lemma 2.34 that $T_{\in\Psi(X)}$ is the canonical representation of $\Phi(T_{\in X})$. For the opposite direction we can symmetrically assume that $T_{\in\Psi(X)}$ is the canonical representation for $\Phi(T_{\in X})$ implying that $T_{\in X}$ is canonical as well. Under these assumptions, Lemma 2.10 and the fact that Ψ is a combinatorial equivalence yield the following equivalences completing the proof:

$$\begin{aligned} T_{\in X} \sqsubseteq S_{\in Y} &\iff T \sqsubseteq S \wedge X \sqsubseteq Y \iff T \sqsubseteq S \wedge \Psi(X) \sqsubseteq \Psi(Y) \\ &\iff T_{\in\Psi(X)} \sqsubseteq S_{\in\Psi(Y)} \iff \Phi(T_{\in X}) \sqsubseteq \Phi(S_{\in Y}) \end{aligned}$$

□

2.3.2 More Combinatorially Equivalent Level Set Complexes and First Conclusions

A First Step towards Level Preserving Isotopies. The next result establishes a combinatorial equivalence between the level set complex $M_{\in[a, b]}$ and the product complex $M_{=h} \times [a, b]$ of the level set complex $M_{=h}$ with an interval $[a, b]$ containing h when no value in the interval $[a, b]$ is attained at a vertex of M . This result can be viewed as a first step towards a piecewise linear analogue of Fact 1.3 for smooth Morse theory, that there is a level preserving isotopy between level sets of an interval without critical values. Note that

for any value $t \in [a, b]$ the product complex $M_{=h} \times \{t\}$ can be identified with $M_{=h}$ in the obvious way, and that it can also be identified in a natural way with a level set complex of $M_{=h} \times [a, b]$, namely the complex $(M_{=h} \times [a, b])_{\pi_2=t}$ where π_2 is the projection to the second coordinate of the Cartesian product $|M_{=h}| \times [a, b]$, which is a linear map. For $t = a$ or $t = b$ the complex $M_{=h} \times \{t\}$ is even a subcomplex of $M_{=h} \times [a, b]$ that can be identified with $M_{=h}$. Using this identification, it makes sense to say that a combinatorial equivalence on $M_{=h} \times [a, b]$ restricted to the subcomplex $M_{=h} \times \{a\}$ or $M_{=h} \times \{b\}$ agrees with some combinatorial equivalence defined on $M_{=h}$.

Lemma 2.36. *Assume that M is a polytopal complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells such that $f(v) \notin [a, b]$ for all vertices v of M . Then $M_{\in[a,b]}$ is combinatorially equivalent to $M_{=h} \times [a, b]$ for any $h \in [a, b]$. At the endpoints of the interval $[a, b]$, the combinatorial equivalence restricts to the natural combinatorial equivalences in the sense of Lemma 2.35 between $M_{=a}$ and $M_{=h}$ (identified with $M_{=h} \times \{a\}$) on the one hand and between $M_{=b}$ and $M_{=h}$ (identified with $M_{=h} \times \{b\}$) on the other hand.*

Proof. Under the given assumption that no value in $[a, b]$ is attained at vertices of the polytopal complex M , Lemmata 2.28 and 2.30 yield that for any fixed cell S of M either all of $S_{=a}$, $S_{=b}$, $S_{\in[a,b]}$, and $S_{=h}$ are canonical or all of them are empty. This shows that assignments that map the cell $S_{\in Y}$ in $M_{\in[a,b]}$ to the cell $S_{=h} \times Y$ in $M_{=h} \times [a, b]$ and vice versa are independent of the representation of the cells and inverse to each other. At the endpoints of the interval with the indicated identifications, these assignments obviously agree with the assignments of the natural combinatorial equivalence from Lemma 2.35.

For establishing that the assignments define a combinatorial equivalence between $M_{\in[a,b]}$ and $M_{=h} \times [a, b]$, it remains to check that $T_{\in X}$ is a face of $S_{\in Y}$ in $M_{\in[a,b]}$ if and only if $T_{=h} \times X$ is a face of $S_{=h} \times Y$ in $M_{=h} \times [a, b]$. According to Lemma 2.30 again, the face relation $T_{\in X} \sqsubseteq S_{\in Y}$ holds in $M_{\in[a,b]}$ if and only if $T_{\in X}$ is empty or $T \sqsubseteq S$ and $X \sqsubseteq Y$ hold. By our first observation in this proof, $T_{\in X}$ is empty if and only if $T_{=h}$ is empty or X is empty, which is again equivalent to $T_{=h} \times X$ being the empty cell in $M_{=h} \times [a, b]$. Furthermore, by Lemma 2.28, the relation $T_{=h} \sqsubseteq S_{=h}$ holds if and only if $T_{=h}$ is empty or $T \sqsubseteq S$. We can conclude that $T_{\in X} \sqsubseteq S_{\in Y}$ holds if and only if $T_{=h}$ is empty or X is empty or $T_{=h} \sqsubseteq S_{=h}$ and $X \sqsubseteq Y$ hold. But this last criterion is exactly the characterisation of the face relation $T_{=h} \times X \sqsubseteq S_{=h} \times Y$ in the product complex $M_{=h} \times [a, b]$ as observed in Definition 1.15. \square

Corollary 2.37. *For any value h that is not attained as f -value at vertices of a polytopal complex M , the level set $M_{=h}$ is bicollared in M . If additionally M is a combinatorial manifold and $M_{=h}$ a sub-manifold of M , then $M_{=h}$ is locally flat in M .*

Proof. If $f(v) \neq h$ for all vertices, then for some small enough ϵ no vertex has a value in $[a, b] = [h - \epsilon, h + \epsilon]$. Applying Lemma 2.36 for the intervals $[a, h]$ and $[h, b]$ yields combinatorial equivalences between $M_{=h} \times [a, h]$ and $M_{\in[a,h]}$ on the one hand and $M_{=h} \times [h, b]$ and $M_{\in[h,b]}$ on the other hand. When we identify the common subcomplex $M_{=h} \times \{h\}$ of $M_{=h} \times [a, h]$ and $M_{=h} \times [h, b]$ with $M_{=h}$, both combinatorial equivalences restrict to the natural combinatorial equivalence, which is the identity on $M_{=h}$. Consider the interval complex

$\mathcal{I} = \mathcal{F}([a, h]) \cup \mathcal{F}([h, b])$. Then $M_{\in \mathcal{I}}$ contains $M_{\in [a, h]}$, $M_{\in [h, b]}$, and $M_{=h}$ as subcomplexes and $M_{=h} \times \mathcal{I}$ contains $M_{=h} \times [a, h]$, $M_{=h} \times [h, b]$, and $M_{=h} \times \{h\}$ as subcomplexes. Since the two combinatorial equivalences agree on their common intersection, we can combine them into a single combinatorial equivalence between $M_{\in \mathcal{I}}$ and $M_{=h} \times \mathcal{I}$. We follow the construction of Lemma 1.30 to obtain combinatorially equivalent triangulations of both complexes without introducing new vertices, and transform the combinatorial equivalence between $M_{\in \mathcal{I}}$ and $M_{=h} \times \mathcal{I}$ into a PL homeomorphism $\phi: |M_{=h} \times \mathcal{I}| \rightarrow |M_{\in \mathcal{I}}|$ between their domains by simplex-wise linear interpolation based on the constructed triangulations. According to our observations in Lemma 1.30, the combinatorial equivalence in terms of vertices between the induced triangulations of $M_{=h} \times \{h\}$ and $M_{=h}$ is still the one given by the natural identification of the vertices, so that simplex-wise linear interpolation yields the natural identification of $|M_{=h} \times \{h\}|$ and $|M_{=h}|$, in other words the restriction of ϕ to $|M_{=h} \times \{h\}|$ corresponds to the identity on $|M_{=h}|$. Since the domain of \mathcal{I} is the interval $[a, b]$, we can consider ϕ as a PL embedding of $|M_{=h}| \times [a, b]$ into $|M|$, and as such it fulfils the definition of a bicollar.

For the second claim, consider a point p in $|M_{=h}|$. We have to check two properties: The first is that p is a boundary point of $|M_{=h}|$ if and only if it is a boundary point of $|M|$, which establishes that $(M, M_{=h})$ is a proper manifold pair; and the second is that for some neighbourhood N of p in $|M|$, the pair $(N, N_{=h})$ is homeomorphic to the standard ball pair. We verify both properties in parallel. Note first that any neighbourhood of p in $|M_{\in [a, b]}|$ is also a neighbourhood of p in $|M|$ because $|M_{\in [a, b]}|$ is a neighbourhood of $|M_{=h}|$ by continuity of f . Since $M_{=h}$ is a manifold by assumption, there is a ball neighbourhood N' of p in $|M_{=h}|$. Furthermore, p is a boundary point of $|M_{=h}|$ if and only if it is a boundary point of N' . Now the pair $(N' \times [a, b], N')$ is obviously homeomorphic to the standard ball pair and $N' \times [a, b]$ is a neighbourhood of (p, h) in $|M_{=h} \times [a, b]|$, with (p, h) being a boundary point of $N' \times [a, b]$ if and only if p is a boundary point of N' . The PL homeomorphism ϕ constructed in the first part of the proof maps this pair $(N' \times [a, b], N')$ to some pair (N, N') with $N' = N_{=h}$ where N is some neighbourhood of p in $|M_{\in [a, b]}|$ which is also a neighbourhood of p in $|M|$ as noted above. Furthermore, $\phi(p, h) = p$ is a boundary point of N if and only if (p, h) is a boundary point in $N' \times [a, b]$. Now we can concatenate the single proof steps and establish the two desired properties: Composing the equivalences regarding the property of being a boundary point, we obtain that p is a boundary point of $M_{=h}$ if and only if it is a boundary point of its neighbourhood N in $|M|$, which is again equivalent to p being a boundary point of $|M|$. This proves that $(M, M_{=h})$ is a proper manifold pair. For the second property, we have seen that for the constructed neighbourhood N of p in $|M|$, the pair $(N, N_{=h}) = (N, N')$ is homeomorphic to $(N' \times [a, b], N')$ which is again homeomorphic to the standard ball pair as desired. \square

A Useful Combinatorial Equivalence and Its Implications. For simplicial complexes, there is a useful combinatorial equivalence between certain level sets in the link and the star of a vertex. Namely, for a value b different from the value $f(v)$ at the vertex v , the level set of the star of v at level b is, depending on whether b is greater or less than $f(v)$, equivalent either to the upper level




Figure 2.2: An example for the combinatorial equivalence of $\text{lk}(v)_{\geq b}$ and $\text{st}(v)_{=b}$ for a vertex v of a simplicial complex and a value $b > f(v)$. For an edge $U \in \text{lk}(v)$ and the adjacent triangle $vU \in \text{st}(v)$, the level set cells $U_{\geq b} \in \text{lk}(v)_{\geq b}$ and $(vU)_{=b} \in \text{st}(v)_{=b}$, that correspond to each other in the combinatorial equivalence, are indicated.

set of the link above level b or to the lower level set of the link below level b . An instance of such a pair of combinatorially equivalent level set complexes is depicted in Figure 2.2, including an example for a pair of level set cells that correspond to each other under the combinatorial equivalence defined in the proof below. Actually the proof only requires that the star of the considered vertex v is the cone on its link with apex v . But if each vertex of a polytopal complex fulfils this condition, then an elementary starring at a vertex has no effect on the complex, which means that the complex is in fact simplicial, because triangulating it by starring at each vertex reproduces the original complex. On the other hand, the weaker assumption still has its merits for polytopal complexes in general, since we can perform an elementary starring at some arbitrary point p in its domain, yielding a subdivision with p as a vertex fulfilling $\text{st}(p) = p \text{lk}(p)$. The result plays a crucial role for the construction of level-preserving isotopies in Chapter 4.

Lemma 2.38. *Let M be a polytopal complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Consider a vertex v such that $\text{st}(v)$ is the cone $v \text{lk}(v)$ and some value $b > f(v)$. There is a combinatorial equivalence between the level set complexes $\text{lk}(v)_{\geq b}$ and $\text{st}(v)_{=b}$ that is the identity on their common subcomplex $\text{lk}(v)_{=b}$. Symmetrically, for $b < f(v)$, there is a combinatorial equivalence between the level set complexes $\text{lk}(v)_{\leq b}$ and $\text{st}(v)_{=b}$ that is the identity on their common subcomplex $\text{lk}(v)_{=b}$.*

Proof. We only give the proof for the case $b > f(v)$. Our aim is to define a poset isomorphism $\phi^{\text{comb}}: \text{lk}(v)_{\geq b} \rightarrow \text{st}(v)_{=b}$ that extends the identity on $\text{lk}(v)_{=b}$. Clearly, any cell of $\text{lk}(v)_{\geq b}$ that does not already belong to $\text{lk}(v)_{=b}$ can be represented by a term $U_{\geq b}$ for some cell U of $\text{lk}(v)$ and its canonical representation

is necessarily of that form. Similarly, since we assume $\text{st}(v) = v\text{lk}(v)$, any cell of $\text{st}(v)_{=b}$ that does not belong to $\text{lk}(v)_{=b}$ is canonically represented by a term $(vW)_{=b}$ for some cell $W \in \text{lk}(v)$. Singling out the canonical ones among those terms $U_{\geq b}$ and $(vW)_{=b}$ with the help of Lemmata 2.26 and 2.27 yields the following bijective correspondences: A term $U_{\geq b}$ represents canonically an i -dimensional cell of $\text{lk}(v)_{\geq b}$ not belonging to $\text{lk}(v)_{=b}$ if and only if U is an i -dimensional cell of $\text{lk}(v)$ that contains at least one vertex with a value greater than b ; a term $(vW)_{=b}$ represents canonically an i -dimensional cell of $\text{st}(v)_{=b}$ not belonging to $\text{lk}(v)_{=b}$ if and only if W is an i -dimensional cell of $\text{lk}(v)$ that contains at least one vertex with a value greater than b . Therefore the assignment $\phi^{\text{comb}}(U_{\geq b}) = (vU)_{=b}$ defines a dimension-respecting bijection between $\text{lk}(v)_{\geq b} \setminus \text{lk}(v)_{=b}$ and $\text{st}(v)_{=b} \setminus \text{lk}(v)_{=b}$ when restricted to cells U of $\text{lk}(v)$ that contain at least one vertex with a value greater than b .

For verifying that ϕ^{comb} , defined as identity on $\text{lk}(v)_{=b}$ and by the above assignment on $\text{lk}(v)_{\geq b} \setminus \text{lk}(v)_{=b}$, is a poset isomorphism, it remains to check that it respects the face relation, that is $\Sigma \sqsubseteq S$ holds in $\text{lk}(v)_{\geq b}$ if and only if $\phi^{\text{comb}}(\Sigma) \sqsubseteq \phi^{\text{comb}}(S)$ holds in $\text{st}(v)_{=b}$. This is trivial if Σ and S are both in $\text{lk}(v)_{=b}$, or if $\Sigma = \emptyset$, which is equivalent to $\phi^{\text{comb}}(\Sigma) = \emptyset$. So let S be canonically represented by $U_{\geq b}$ and assume first that Σ is a non-empty face of S . By Lemma 2.1, we have $\Sigma = F_{\geq b}$ or $\Sigma = F_{=b}$ for some $F \sqsubseteq U$. Furthermore we can assume without loss of generality that this representation is canonical, because any canonical representation of a non-empty face of $U_{\geq b}$ is of one of the two forms. In the former case, $\phi^{\text{comb}}(\Sigma) = (vF)_{=b} \sqsubseteq (vU)_{=b} = \phi^{\text{comb}}(S)$ holds. In the latter case we obtain $\phi^{\text{comb}}(\Sigma) = F_{=b} \sqsubseteq U_{=b} \sqsubseteq (vU)_{=b} = \phi^{\text{comb}}(S)$. It remains to show the other direction where we assume that $\phi^{\text{comb}}(\Sigma)$ is a face of $\phi^{\text{comb}}(S) = (vU)_{=b}$. Non-empty faces of $(vU)_{=b}$ are canonically represented either by a term $(vF)_{=b}$ or by a term $F_{=b}$ for some face F of U . If $\phi^{\text{comb}}(\Sigma)$ is canonically represented by $(vF)_{=b}$, we obtain the desired face relation $\Sigma = F_{\geq b} \sqsubseteq U_{\geq b} = S$. Otherwise, if $\phi^{\text{comb}}(\Sigma)$ is canonically represented by $F_{=b}$, we can conclude $\Sigma = F_{=b} \sqsubseteq U_{\geq b} = S$, which completes the proof. \square

Applying the fact that combinatorial equivalences induce PL homeomorphisms to the combinatorial equivalence constructed in the previous lemma allows to transfer topological properties of $\text{lk}(v)_{\geq b}$ to $\text{st}(v)_{=b}$. The next lemma relies on that observation.

Lemma 2.39. *Let M be a d -dimensional combinatorial manifold with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Consider a point $x \in |M|$ that is contained in the interior of a cell where f is non-constant, which means in particular that x must not be a vertex. Then the level set $|M|_{=f(x)}$ is a $(d-1)$ -dimensional manifold in some neighbourhood of x .*

Proof. Choose a vertex v of the cell that contains x in its interior such that $f(x) \neq f(v)$. We can assume $f(x) > f(v)$ without loss of generality, otherwise a completely symmetric argument applies. By Lemma 2.38, the level set $\text{st}(v)_{=f(x)}$ is combinatorially equivalent to $\text{lk}(v)_{\geq f(x)}$ with $\text{lk}(v)_{=f(x)}$ being fixed. Thus, there is a PL homeomorphism between $|\text{st}(v)|_{=f(x)}$ and $|\text{lk}(v)|_{\geq f(x)}$ which is the identity on $|\text{lk}(v)|_{=f(x)}$. By our choice of the vertex v , the point x is contained in $|\text{st}(v)|$, but not in $|\text{lk}(v)|$, implying that it belongs to $|\text{st}(v)|_{=f(x)}$ but not to $|\text{lk}(v)|_{=f(x)}$. Therefore the PL homeomorphism maps x to some point y in $|\text{lk}(v)|_{>f(x)}$. Now $|\text{lk}(v)|$ is a PL $(d-1)$ -manifold, in fact a ball or sphere,

since we assume M to be a d -manifold. Hence its open subset $|\text{lk}(v)|_{>f(x)}$ is a PL $(d-1)$ -manifold as well, and also its corresponding set under the PL homeomorphism, namely $(|\text{st}(v)| \setminus |\text{lk}(v)|)_{=f(x)}$. The general fact for any vertex v of a combinatorial manifold M that $|\text{st}(v)| \setminus |\text{lk}(v)|$ is an open neighbourhood in $|M|$ of any of its points, shows that $|M|_{=f(x)}$ is a $(d-1)$ -manifold in the neighbourhood $|\text{st}(v)| \setminus |\text{lk}(v)|$ of x as desired, namely the manifold $(|\text{st}(v)| \setminus |\text{lk}(v)|)_{=f(x)}$. Furthermore x is mapped to a boundary point of $\text{lk}(v)$ if and only if x itself is a boundary point. \square

Note that if any two vertices forming an edge of M have different f -values, then the assumption on the point x in the Lemma is fulfilled for any point that is not a vertex. Recall that we defined a map linear on cells to be in general position if all its values at vertices are distinct, so that this remark applies in particular to maps in general position. In addition, we can invoke Corollary 2.37 for any level set $|M|_{=h}$ such that the value h is not attained at vertices.

Corollary 2.40. *Let M be a d -dimensional combinatorial manifold with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells.*

1. *If f is in general position (it suffices to assume that any vertices u, v such that uv is an edge of M have distinct f -values), then any point in $|M|_{=h}$ that is not a vertex of M has some neighbourhood in $|M|$ in which $|M|_{=h}$ is a $(d-1)$ -manifold.*
2. *Even without general position assumption on f , it is true for any value h not attained at vertices, that $|M|_{=h}$ is a $(d-1)$ -dimensional locally flat sub-manifold of $|M|$.*

Proof. Under the given assumptions, any point that is not a vertex is contained in the interior of a cell where f is non-constant and Lemma 2.39 yields the first claim. Since being a manifold is a local property, a set that is a manifold in the neighbourhood of any of its points, is also globally a manifold. If the value h is not attained at vertices, then no point x in $|M|_{=h}$ is contained in a cell where f is constant, because the constant value for the cell would be h and it would be attained at all vertices of the cell. This observation holds even if f is not in general position. Therefore $|M|_{=h}$ is a $(d-1)$ -dimensional manifold for any value h not attained at vertices, because it is a manifold in the neighbourhood of any of its points by Lemma 2.39. It follows from Corollary 2.37 that $|M|_{=h}$ is locally flat in $|M|$. \square

2.3.3 Collapsing Lower Level Set Complexes

Now we turn to the investigation of collapses that can be performed on level set complexes.

Collapsing Canonically Represented Pairs. For lower and upper level set complexes $M_{\leq h}$ or $M_{\geq h}$, the cells S of M that are neither completely contained in nor disjoint from the level set contribute two cells to the complex, namely $S_{=h}$ and the lower part $S_{\leq h}$ or the upper part $S_{\geq h}$ respectively. We show that this pair of cells can be collapsed, which may require prior collapsing of co-faces. Note that even if M is simplicial, the level set complexes $M_{\leq h}$ or $M_{\geq h}$

are usually not simplicial. Therefore the results rely on the remark made in Subsection 1.2.10, that collapses can be defined without change on polytopal complexes as well. Below, we formulate all results only for lower level set complexes. The analogous statements for upper level set complexes are also true and can be proven in the same way. We start with a proof for maximal cells S .

Lemma 2.41. *Let S be a maximal cell of a polytopal complex M such that the terms $S_{\leq h}$ and $S_{=h}$ are canonical. Then an elementary collapse of the cells $S_{\leq h}$ and $S_{=h}$ can be performed in the complex $M_{\leq h}$.*

Proof. For verifying that two cells $C \sqsubseteq D$ in a complex can be removed by an elementary collapse, it suffices to show that C and D are distinct and that the only co-faces of the smaller cell C are C itself and the other cell D . The assumption that $S_{\leq h}$ is canonical guarantees that $S_{\leq h}$ and $S_{=h}$ are distinct cells. Since $S_{=h}$ is canonical, the co-faces of $S_{=h}$ are of the form $T_{\leq h}$ or of the form $T_{=h}$ where T is a co-face of S by Lemma 2.9. This implies $T = S$ because S is assumed to be maximal in M . Hence $S_{=h}$ and $S_{\leq h}$ are the only co-faces of $S_{=h}$ as desired. \square

For a cell S with the property that the level set cells $S_{\leq h}$ and $S_{=h}$ are canonical, this property is inherited by all the co-faces of S . Therefore we can collapse $S_{\leq h}$ and $S_{=h}$ after having collapsed inductively all cells $T_{\leq h}$ and $T_{=h}$ for T a co-face of S . This result can be obtained as a corollary of the following more general strategy for performing collapses on level set complexes:

Lemma 2.42. *Let K be a subcomplex of a polytopal complex M such that for each cell $S \in M \setminus K$ either the terms $S_{\leq h}$ and $S_{=h}$ are both canonical or the level set cell $S_{\leq h}$ is contained in $|K|$. Then $M_{\leq h}$ collapses to $K_{\leq h}$. More specifically, for each cell $S \in M \setminus K$ fulfilling the first alternative, the cells $S_{\leq h}$ and $S_{=h}$ can be eventually collapsed after having collapsed co-faces of the same type, whereas for each cell $S \in M \setminus K$ that fulfils the second but not the first alternative there is a cell $T \in K$ such that $S_{\leq h} = S_{=h} = T_{=h} = T$ holds. Performing all collapses of the cells $S_{\leq h}$ and $S_{=h}$ for S fulfilling the first alternative in a suitable order yields a collapse from $M_{\leq h}$ to $K_{\leq h}$.*

Proof. We start with showing the claimed property for cells in $S \in M \setminus K$ that fulfil the second but not the first alternative. Note that among the cases in Lemma 2.25, only Case 3 has the property that the terms $S_{\leq h}$ and $S_{=h}$ are both canonical. This case occurs when f attains on S values strictly greater than h as well as values strictly smaller than h , a property that is obviously inherited by all co-faces of S . Hence, if the first alternative does not apply for S , then f is bounded on S by h , either from below or from above. We show that the assumption that a value strictly smaller than h is attained on S leads to contradiction with the other assumptions, excluding Cases 5 and 6 from Lemma 2.25: If a value strictly smaller than h were attained on S , then h would be an upper bound, meaning $S_{\leq h} = S$. Since we assume that the second alternative applies, we would obtain that S is contained in $|K|$, which means that it belongs to the subcomplex K , contradicting the assumption $S \in M \setminus K$. The only remaining possibility is that h is a lower bound for f on S , which implies $S_{\leq h} = S_{=h}$, and we also know that these two level set cells agree with a face T of S , namely with \emptyset in Case 1 of Lemma 2.25, where all f -values on S are strictly greater than h , with S itself in Case 4, where f is constant on S with

value h , and with the proper face T of S defined by the supporting hyperplane $f_S^{-1}(h)$ in Case 2, where h is the minimum attained by some non-constant f on S . In each case, the equality $T = T_{=h}$ is obvious from $T = S_{=h}$, and the face $T = S_{\leq h}$ is a cell of K , because we assume that $S_{\leq h}$ is contained in $|K|$.

Let (S^1, S^2, \dots, S^k) be an enumeration of the cells in $M \setminus K$ such that whenever S^i is a co-face of S^j then $i \leq j$. We claim that the sequence of pairs of level set cells $((S_{\leq h}^1, S_{=h}^1), (S_{\leq h}^2, S_{=h}^2), \dots, (S_{\leq h}^k, S_{=h}^k))$ describes a sequence of valid elementary collapses or vacuous operations from the starting complex $M_{\leq h}$, with a pair describing an elementary collapse if $S_{\leq h}^j$ and $S_{=h}^j$ are both canonical, and a vacuous operation if $S_{\leq h}^j = S_{=h}^j = T_{=h}^j = T^j$ for some $T^j \in K$. In other words, we let $K^0 = M$ and $K^j = K^{j-1} \setminus S^j$ for $1 \leq j \leq k$ and show that $K_{\leq h}^{j-1}$ collapses to or agrees with $K_{\leq h}^j$ for $1 \leq j \leq k$. This sequence gives a collapse from $K_{\leq h}^0 = M_{\leq h}$ to $K_{\leq h}^k = K_{\leq h}$.

We consider the complex $K_{\leq h}^{j-1}$ and the pair $(S_{\leq h}^j, S_{=h}^j)$. If $S_{\leq h}^j = S_{=h}^j = T_{=h}^j = T^j$ for some $T^j \in K$, then clearly $K_{\leq h}^{j-1}$ and $K_{\leq h}^j$ agree because $T^j \in K^j$. Otherwise we want to collapse the pair $(S_{\leq h}^j, S_{=h}^j)$. Since K is a subcomplex, all co-faces of S^j are contained in $M \setminus K$ and by our assumption on the sequence (S^1, S^2, \dots, S^k) , each co-face of S^j is of the form S^i for some $i \leq j$. Hence S^j is a maximal cell of $K^{j-1} = M \setminus \{S^1, S^2, \dots, S^{j-1}\}$. By Lemma 2.41, an elementary collapse of $S_{\leq h}^j$ and $S_{=h}^j$ can be performed in $K_{\leq h}^{j-1}$, meaning $K_{\leq h}^{j-1}$ collapses to $K_{\leq h}^j$ as desired. \square

Corollary 2.43. *Let S be a cell of a polytopal complex M such that $S_{\leq h}$ and $S_{=h}$ are canonical. Then there is a collapse of $M_{\leq h}$ which removes all cells $T_{\leq h}$ and $T_{=h}$ for T a co-face of S .*

Proof. Let K be the subcomplex of M obtained by removing S and all its co-faces from M . It follows immediately from our characterisations of canonically represented cells that the terms $T_{\leq h}$ and $T_{=h}$ for T a co-face of S are also canonical: The cell S contains elements x^+ and x^- with $f(x^+) > h$ and $f(x^-) < h$, and then so does T . Hence the cells from $M \setminus K$ fulfil the requirements of Lemma 2.42, yielding a collapse of $M_{\leq h}$ to $K_{\leq h}$, which is the complex obtained from $M_{\leq h}$ by removing all cells $T_{\leq h}$ and $T_{=h}$ for T a co-face of S . \square

Partial Collapse to a Certain Level. In another variant of the previous results, instead of collapsing a whole cell $S_{\leq h}$ we only collapse the part of the cell that lies above a certain level. Formally speaking, we consider the subdivision $M_{\leq a} \cup M_{\geq a}$ of the complex M induced by the level set $M_{=a}$ where $a < h$ and the cell $S_{\geq a} \in M_{\leq a} \cup M_{\geq a}$, and apply the collapse to the pair consisting of $(S_{\geq a})_{\leq h} = S_{\in[a, h]}$ and $(S_{\geq a})_{=h} = S_{=h}$. We call such an operation a collapse of $S_{\leq h}$ to $S_{\leq a}$ or a collapse to level a , because in the complex obtained from $M_{\leq h}$ by subdividing and collapsing, the cell $S_{\leq h}$ is first subdivided into $S_{\leq a}$ and $S_{\in[a, h]}$, and then the latter part $S_{\in[a, h]}$ is removed, leaving only the former part $S_{\leq a}$.

Corollary 2.44. *Assume $a < h$ and consider the subdivision $M_{\leq a} \cup M_{\geq a}$ of a polytopal complex M induced by the level set $M_{=a}$. Then the following statements regarding collapses in $(M_{\leq a} \cup M_{\geq a})_{\leq h} = M_{\leq a} \cup M_{\in[a, h]}$ hold:*

Let S be a maximal cell of a polytopal complex M such that $S_{\leq h}$ and $S_{=h}$ are canonical. Then on $M_{\leq a} \cup M_{\in[a,h]}$ an elementary collapse removing the cells $(S_{\geq a})_{\leq h} = S_{\in[a,h]}$ and $(S_{\geq a})_{=h} = S_{=h}$ can be performed.

Let K be a subcomplex of M such that for each cell $S \in M \setminus K$ the terms $S_{\leq h}$ and $S_{=h}$ are canonical or the level set cell $S_{\leq h}$ is contained in $|K|$ or in $|M|_{\leq a}$. Then $(M_{\leq a} \cup M_{\geq a})_{\leq h} = M_{\leq a} \cup M_{\in[a,h]}$ collapses to $M_{\leq a} \cup (K_{\geq a})_{\leq h} = M_{\leq a} \cup K_{\in[a,h]}$.

Proof. We show that the first claim can be obtained from Lemma 2.41 and the second from Lemma 2.42 by applying them to $M_{\leq a} \cup M_{\geq a}$ in place of M and $M_{\leq a} \cup K_{\geq a}$ in place of K . Note that we can consider $M_{\leq a} \cup M_{\geq a}$ as a single level set complex $M_{\in \mathcal{I}}$ for the interval complex $\mathcal{I} = \mathcal{F}((-\infty, a]) \cup \mathcal{F}([a, \infty))$. At first we observe that $S_{\leq h}$ and $S_{=h}$ being both canonical implies that $(S_{\geq a})_{\leq h}$ and $(S_{\geq a})_{=h}$ are both canonical: As used already in the proof of Lemma 2.41, the terms $S_{\leq h}$ and $S_{=h}$ being both canonical is equivalent to S containing elements x^+ and x^- with $f(x^+) > h$ and $f(x^-) < h$. This condition implies obviously that the upper part $S_{\geq a}$ of the subdivided cell S contains elements y^+ and y^- with $f(y^+) > h$ and $f(y^-) < h$, meaning that the terms $(S_{\geq a})_{\leq h}$ and $(S_{\geq a})_{=h}$ are canonical representations for their cells in $(M_{\geq a})_{\leq h}$. Now if S is a maximal cell in M with $S_{\leq h}$ and $S_{=h}$ being canonical, then we can also prove that $S_{\geq a}$ is maximal in $M_{\in \mathcal{I}}$ as follows: The term $S_{\geq a}$ is canonical because S contains an element x^+ with $f(x^+) > h$. Therefore Lemma 2.9 shows that any co-face of $S_{\geq a}$ in $M_{\in \mathcal{I}}$ is of the form $T_{\in Y}$ where T is a co-face of S in M and Y a co-face of $(-\infty, a]$ in \mathcal{I} . Since S is maximal in M and $(-\infty, a]$ is maximal in \mathcal{I} , the cell $S_{\geq a}$ is also maximal in $M_{\in \mathcal{I}}$. This finishes already the proof for the first claim, because the maximal cell $S_{\geq a}$ in $M_{\leq a} \cup M_{\geq a}$ fulfils the requirements of Lemma 2.41, and applying it yields that an elementary collapse of $(S_{\geq a})_{\leq h} = S_{\in[a,h]}$ and $(S_{\geq a})_{=h} = S_{=h}$ can be performed in $(M_{\leq a} \cup M_{\geq a})_{\leq h}$.

For the second claim, it remains to check that under the given assumptions for any cell C in $(M_{\leq a} \cup M_{\geq a}) \setminus (M_{\leq a} \cup K_{\geq a})$ either $C_{\leq h}$ and $C_{=h}$ are both canonical or $C_{\leq h}$ is contained in $|M|_{\leq a} \cup |K|_{\geq a}$. Such a cell C from the given set is necessarily representable by a term of the form $S_{\geq a}$ for some $S \in M \setminus K$. For such cells S in turn, we assume that either $S_{\leq h}$ and $S_{=h}$ are both canonical or that $S_{\leq h}$ is contained in $|K|$ or $|M|_{\leq a}$. We have already seen that the former assumption implies that $C_{\leq h} = (S_{\geq a})_{\leq h}$ and $C_{=h} = (S_{\geq a})_{=h}$ are canonical. If the latter assumption applies and $S_{\leq h}$ is contained in $|K|$ or $|M|_{\leq a}$, then $C_{\leq h} = (S_{\geq a})_{\leq h} = (S_{\leq h})_{\geq a}$ is contained in $|K|_{\geq a}$ or $|M|_{\leq a}$ as desired. Thus, the subcomplex $M_{\leq a} \cup K_{\geq a}$ of $M_{\leq a} \cup M_{\geq a}$ fulfils the requirements of Lemma 2.42, showing that $M_{\leq a} \cup M_{\in[a,h]}$ collapses to $M_{\leq a} \cup K_{\in[a,h]}$. \square

Performing Collapses to All Applicable Cells. To conclude this chapter, we examine which cells remain when we perform the collapses to all applicable cells. We consider both variants, the complete collapse of all pairs $S_{\leq h}$ and $S_{=h}$ in $M_{\leq h}$ such that both terms are canonical according to Corollary 2.43 and the collapse to level a of all cells $S_{\in[a,h]}$ and $S_{=h}$ in $M_{\leq a} \cup M_{\in[a,h]}$ such that $S_{\leq h}$ and $S_{=h}$ are canonical according to Corollary 2.44. A small example illustrating both variants is depicted in Figure 2.3. In the first variant, it turns out that we obtain not only a subcomplex of $M_{\leq h}$ but also a subcomplex of M . It consists of those cells of M that have only vertices whose f -values are at most h , in other words, it is the subcomplex induced by such vertices. This complex can also be




Figure 2.3: Instances of the collapses of level set complexes described in Lemma 2.46. The light and dark shaded areas together form the original level set complex, on the left hand side $M_{\leq h}$, on the right hand side its subdivision $M_{\leq a} \cup M_{\in[a, h]}$. The dark shaded areas alone represent the resulting complex after the collapse, namely $\text{scp}_{\leq h}(M)$ on the left side and $M_{\leq a} \cup (\text{scp}_{\leq h}(M))_{\in[a, h]}$ on the right side.

characterised as the largest subcomplex of M contained in $|M|_{\leq h}$. We denote such a subcomplex of M consisting of all cells of M lying completely in $|M|_{\leq h}$ for some level h by $\text{scp}_{\leq h}(M)$, so that the result can be stated symbolically as $M_{\leq h} \searrow \text{scp}_{\leq h}(M)$.

This first version of the result has been remarked already previously in the literature: It is mentioned in passing by Brehm and Kühnel [9, p. 468] and stated by Kühnel [30, Lemma 7.2 (ii), p. 97], with a focus on the case that M is a combinatorial manifold and $h = f(v)$ for some vertex v . Immediate consequences of this result are observed by Morozov [37, p. 24], who notes that there is a deformation retraction from $M_{\leq h}$ to $\text{scp}_{\leq h}(M)$, and by Bauer [6, Proposition 2.25, p. 31], who points out the homotopy equivalence of the two complexes. Deformation retractions similar to those induced by the collapses in the second variant, which retract $M_{\in[a, a+\epsilon]}$ to $M_{=a}$ when $f^{-1}(a, a+\epsilon]$ contains no vertices and $M_{\in[f(v)-\epsilon, f(v)]}$ to $M_{=f(v)-\epsilon} \cup (v(\text{scp}_{\leq f(v)}(\text{lk}(v))) \cap M_{\in[f(v)-\epsilon, f(v)]})$ when v is the only vertex attaining a value in $[f(v)-\epsilon, f(v)]$, are presented by Bestvina [7, Propositions 2.4 and 2.7].

Definition 2.45. For a polytopal complex M with a map f linear on cells and a real value h we denote by $\text{scp}_{\leq h}(M)$ the subcomplex of M consisting of all cells of M lying completely in $|M|_{\leq h}$, in other words the subcomplex induced by the vertices with f -value at most h . We can define an analogous subcomplex $\text{scp}_{\geq h}(M)$, which is induced by the vertices with f -value at least h .

Lemma 2.46. *Let M be a polytopal complex with a map f linear on cells. Then for any real value h , the level set complex $M_{\leq h}$ collapses to the subcomplex $\text{scp}_{\leq h}(M)$ of M .*

$$M_{\leq h} \searrow \text{scp}_{\leq h}(M)$$

Furthermore, for any value $a \leq h$, the complex $M_{\leq a} \cup M_{\in[a, h]}$ collapses to $M_{\leq a} \cup (\text{scp}_{\leq h}(M))_{\in[a, h]}$. In particular, if $f^{-1}(a, a+\epsilon]$ does not contain vertices for some $\epsilon > 0$, then $M_{\leq a} \cup M_{\in[a, a+\epsilon]}$ collapses to $M_{\leq a}$.

Proof. We start by showing that the subcomplex $K = \text{scp}_{\leq h}(M)$ fulfils the assumptions of the complex K in Lemma 2.42 and Corollary 2.44. A cell S outside of K necessarily contains a vertex with f -value strictly greater than h . Examining again the cases from Lemma 2.25, we see that if S also contains a vertex with f -value strictly less than h then $S_{\leq h}$ and $S_{=h}$ are both canonical. Otherwise the equalities $S_{\leq h} = S_{=h} = T_{=h} = T$ hold for some cell $T \in K$, meaning that $S_{\leq h}$ is contained in $|K|$. This proves that K has the desired properties.

Applying Lemma 2.42 yields that $M_{\leq h}$ collapses to $K_{\leq h}$. But since K contains only cells of M where f is bounded from above by h , the complexes $K_{\leq h}$ and K agree. This shows the first claim. For the second claim, we apply Corollary 2.44 and obtain that $M_{\leq a} \cup M_{\in[a, h]}$ collapses to $M_{\leq a} \cup (\text{scp}_{\leq h}(M))_{\in[a, h]}$.

For the last claim, it suffices to show that $(\text{scp}_{\leq a+\epsilon}(M))_{\in[a, a+\epsilon]}$ is a subcomplex of $M_{\leq a}$. Note that with the given assumption, all cells of M lying completely in $|M|_{\leq a+\epsilon}$ also lie completely in $|M|_{\leq a}$ and hence the subcomplex $\text{scp}_{\leq a+\epsilon}(M)$ agrees with $\text{scp}_{\leq a}(M)$. This gives $(\text{scp}_{\leq a+\epsilon}(M))_{\in[a, a+\epsilon]} = (\text{scp}_{\leq a}(M))_{\in[a, a+\epsilon]} = (\text{scp}_{\leq a}(M))_{=a}$, with the last term obviously being a subcomplex of $M_{\leq a}$. \square

Chapter 3

Piecewise Linear Morse Functions

The object of study in classical smooth Morse theory are smooth manifolds with a smooth map defined on them that fulfils a certain property, namely that all its critical points are non-degenerate. Such maps are called Morse functions. In this chapter, we study piecewise linear analogues of the definitions related to Morse functions, such as definitions for regular and critical points and the distinction of the latter kind into degenerate and non-degenerate ones. A short comparison of some previously studied notions of criticality for piecewise linear maps is also part of the chapter. Furthermore we draw some connections between Morse functions on simplicial complexes in the piecewise linear sense and discrete Morse functions on simplicial complexes.

The usual definition of critical points for smooth functions as points where the derivative vanishes has no direct analogue for piecewise linear maps. But the Morse lemma in the variant described in Fact 1.1 provides characterisations of regular and non-degenerate critical points that can be adapted from the smooth category to other categories of topology. It states that a non-degenerate critical point p of index λ of a smooth function f is characterised by the existence of a neighbourhood with local coordinates (x_1, \dots, x_d) such that $p = (0, \dots, 0)$ and $f(x_1, \dots, x_d) = f(p) - x_1^2 - \dots - x_\lambda^2 + x_{\lambda+1}^2 + \dots + x_d^2$ in this coordinate system. Similarly, regular points p are characterised by a neighbourhood with local coordinates (x_1, \dots, x_d) such that $p = (0, \dots, 0)$ and $f(x_1, \dots, x_d) = f(p) + x_1$ in this coordinate system. The remaining points are degenerate critical and a smooth function is a Morse function if all points are either regular or non-degenerate critical.

In a paper of Morse [38], analogous characterisations are used to define “topologically ordinary points”, “topologically critical points (associated with a canonical mapping of index λ)”, and “topologically non-degenerate functions” on general topological manifolds. The only change is the type of the local coordinate system. In classical Morse theory, the local coordinates are given by diffeomorphisms, whereas the local coordinates in the general topology setting of [38] can be given by arbitrary homeomorphisms.

This suggests to define analogous objects in the piecewise linear category. Clearly, the local coordinate system in the characterisations would be given by

PL homeomorphisms. But we also have to change the “standard” functions for non-degenerate critical points, because a sum of squared coordinates or its negatives is not a PL function. For our study, we separate the characterisation of regular and non-degenerate critical points into two parts that can be treated independently.

For the first part we introduce in Definition 3.1 the notion of local equivalence, that allows to express the statement that some function agrees locally with some other function when expressed in suitable PL local coordinates in terms of this equivalence relation. The class of objects that is partitioned into equivalence classes by this relation is formed by triples (M, x, f) consisting of a complex M , a point x in its domain, and a map f linear on cells of the complex. For a given complex M with a map f linear on cells, the points in the domain can be classified into regular points and different kinds of non-degenerate and degenerate critical points according to the equivalence class of (M, x, f) . We establish several equivalent characterisations for local equivalence, which in turn gives several equivalent characterisations for regular and critical points. The list in Theorem 3.11 gives a synopsis on the different characterisations of local equivalence.

The second part consists only of identifying the “standard” functions that exhibit the desired local behaviour at regular or non-degenerate critical points. More specifically, certain triples are used as reference for defining regular or the different kinds of non-degenerate critical points of f on M as those points x such that (M, x, f) is locally equivalent to the respective reference triple. For example, a suitable reference triple for regular points in d -dimensional manifolds could consist of \mathbb{R}^d , that is the complex $\{\emptyset, \mathbb{R}^d\}$ to be more precise, its origin as considered point, and the linear map assigning to each point its first coordinate x_1 . The reference triples and which kind of points they define are listed in Definition 3.12. For the resulting notion of PL Morse functions (Definition 3.13) we show in Theorem 3.24 that any combinatorial manifold has a PL Morse function that is linear on cells of its derived subdivision. The construction used in the proof is a special case of a more general construction described in Lemma 3.26, that produces from a given discrete Morse function on a combinatorial manifold a PL Morse function linear on cells of the derived of that manifold, whose critical points correspond to critical cells of the original function.

The first section of this chapter is devoted to the study of local equivalence and the different ways to characterise this relation. The definition of regular and critical points and the comparison with other notions of criticality is presented in the second section. Additional properties and characterisations of regular points will be considered in a chapter of its own following the present one.

3.1 Local Equivalence

The criterion for non-degenerate critical points given by the Morse lemma and its variant for regular points have in common that they require the existence of a local coordinate system such that the representation of the considered function in these coordinates agrees with a certain function. It seems useful to capture this common point by introducing a notion that considers functions as equivalent at points of their domain, when they can be represented by the same

$$\begin{array}{ccc}
N_x & \xrightarrow{\phi} & N'_{x'} \\
\downarrow f & & \downarrow f' \\
\mathbb{R} & \xrightarrow{h \mapsto h+c} & \mathbb{R}
\end{array}$$

Figure 3.1: Commuting diagram for Definition 3.1 illustrating local equivalence of f at x and f' at x' .

function in suitable local coordinate systems for both points. We can even go further and avoid the detour via coordinate systems by using a direct homeomorphism between suitable neighbourhoods of the points. This way the notion of local equivalence also applies to points that do not have a neighbourhood homeomorphic to some Euclidean space or half-space. Since in Morse theory, the classification of points into regular and critical points does not change when the function is shifted by adding a constant, we consider functions also as equivalent when they differ only by a constant. These considerations lead to a notion that defines two functions at two points in polyhedral complexes as equivalent if the points have PL-homeomorphic neighbourhoods such that the function values of corresponding points in these neighbourhoods differ by the same constant throughout the neighbourhoods. We use PL-invariance of links and stars to show alternative characterisations of local equivalence of two maps at two points. In particular we show that it is enough to find a PL-homeomorphism of the links of the points that respects the subdivision into upper level link and lower level link, with $(\text{lk}(p))_{\geq f(p)}$ being the upper level link and $(\text{lk}(p))_{\leq f(p)}$ being the lower level link of a point p with respect to a map f . This result is stated in Corollary 3.10.

A notion of equivalence based on the above ideas for defining regular and critical points in the piecewise linear category has been suggested previously by Kosinski in [29], where such a notion is even defined for maps that take values in higher dimensional Euclidean spaces. Although the literal definitions of Kosinski's notion of equivalence and of the local equivalence studied below differ in some details, they turn out to be equivalent in all cases where both notions can be applied. We postpone the comparison of the two notions to the next section, when the results of the current section are available and make the comparison easier.

3.1.1 Definition and First Observations

Definition 3.1 (local equivalence). Let M and M' be polyhedral complexes with maps $f: |M| \rightarrow \mathbb{R}$ and $f': |M'| \rightarrow \mathbb{R}$ linear on cells. The function f at $x \in |M|$ is called *locally equivalent* to f' at $x' \in |M'|$ if there are neighbourhoods N_x of x in M and $N'_{x'}$ of x' in M' , a PL homeomorphism $\phi: N_x \rightarrow N'_{x'}$, mapping x to x' , and a constant c , such that for any $y \in N_x$, the difference between the value at the point itself and the value at its corresponding point $\phi(y)$ agrees with the constant c , that is, $f'(\phi(y)) - f(y) = c$. In other words, restricted to N_x , the maps $f' \circ \phi$ and $f + c$ agree, so that the diagram in Figure 3.1 commutes. Note that the value of the constant c is necessarily $f'(x') - f(x)$, and that the requirement can also be expressed by the condition that the map

$$\begin{array}{ccc}
 N_x & \xrightarrow{\phi} & N'_{x'} \\
 & \searrow g & \downarrow g' \\
 & & \mathbb{R}
 \end{array}$$

Figure 3.2: Commuting diagram for Remark 3.2 illustrating local equivalence of f at x and f' at x' by the shifted maps $g = f - f(x)$ and $g' = f' - f'(x')$.

$(f' \circ \phi) - f$ is constant throughout N_x , again with the constant being necessarily $c = f'(x') - f(x)$.

Variants of the Definition. The definition can be rephrased in various ways by shifting the maps f and f' or by making additional assumptions on the neighbourhoods that are possible because the PL homeomorphism ϕ can be restricted to smaller neighbourhoods once a suitable pair of neighbourhoods is established. We collect some easy observations of this kind. The possibility to switch freely between more and less restrictive assumptions facilitate the proofs of other characterisations later on.

Remark 3.2. In Definition 3.1, instead of considering the values of the maps f and f' itself, we can equivalently consider values that are shifted by some constants. In other words, f and f' are locally equivalent at x and x' if and only if the maps $f + s$ and $f' + s'$ defined by $(f + s)(y) = f(y) + s$ for $y \in |M|$ and $s \in \mathbb{R}$ and by $(f' + s')(y') = f'(y') + s'$ for $y' \in |M'|$ and $s' \in \mathbb{R}$ are locally equivalent at these two points. For a fixed pair of points x and x' , a particularly useful choice for shifting the maps is the consideration of the maps $g = f - f(x)$ and $g' = f' - f'(x')$. One advantage of this choice is that the function values at the two points agree and even that the common value is zero: $g(x) = g'(x') = 0$. Another advantage is that it leads to the simplified commutative diagram from Figure 3.2 where $g' \circ \phi = g$ holds when the maps are restricted to the neighbourhoods N_x and $N'_{x'}$:

Proof. The maps $f' \circ \phi$ and $f + c$ agree on N_x if and only if the maps $(f' + s') \circ \phi$ and $(f + s) + (c - s + s')$ agree on N_x . \square

The above remark shows that when we ask for local equivalence at a fixed pair of points x and x' , we can always assume $f(x) = f'(x') = 0$ without loss of generality. Furthermore, the requirement that the difference $f'(\phi(y)) - f(y)$ between the function values at corresponding points y and $\phi(y)$ is constantly $f'(x') - f(x)$ throughout the neighbourhood is equivalent to the requirement that the values of the shifted maps $f - f(x)$ and $f' - f'(x')$ agree at corresponding points y and $\phi(y)$ throughout the neighbourhood: $f(y) - f(x) = f'(\phi(y)) - f'(x')$.

Remark 3.3. For the neighbourhoods N_x and $N'_{x'}$ in Definition 3.1 certain additional properties can be assumed without loss of generality. In the following, we focus on the neighbourhood N_x , but all assumptions can be made analogously for $N'_{x'}$, even if not stated explicitly. First of all, the assumptions described in Subsection 1.2.8 on pseudo-radial projection can be applied: We can assume

that N_x is a cone neighbourhood xL_x with a polytopal complex L_x as link complex of x . Furthermore it can be assumed that the restriction of the map f to the neighbourhood $|xL_x|$ is linear on cells of the complex xL_x . A similar assumption regarding the PL homeomorphism ϕ is possible: Without loss of generality the map ϕ is a cellular bijection between cone neighbourhoods given by complexes xL_x and $x'L_{x'}$.

Proof. Once a pair of PL-homeomorphic neighbourhoods N_x and $N_{x'}$, with PL homeomorphism $\phi: N_x \rightarrow N_{x'}$, and the properties required in Definition 3.1 is established, we can restrict ϕ to any smaller neighbourhood $\tilde{N}_x \subseteq N_x$ and consider \tilde{N}_x and its image $\tilde{N}_{x'} = \phi(\tilde{N}_x) \subseteq N_{x'}$. Then the pair of neighbourhoods \tilde{N}_x and $\tilde{N}_{x'}$ obviously fulfils the required properties as well, with the restriction $\tilde{\phi}$ of ϕ being the PL homeomorphism. Therefore we start with some neighbourhoods N_0 of x and N'_0 of x' fulfilling the requirements in Definition 3.1 with some PL homeomorphism $\phi_0: N_0 \rightarrow N'_0$, and construct iteratively smaller PL-homeomorphic neighbourhoods $N_i \subseteq N_{i-1}$ of x and $N'_i \subseteq N'_{i-1}$ of x' with the PL homeomorphism $\phi_i: N_i \rightarrow N'_i$ being a restriction of ϕ_0 , until we obtain a pair of neighbourhoods N_k and N'_k that fulfils all suggested additional properties.

As a first step, we choose for N_1 a compact cone neighbourhood $x\tilde{L}_1$ contained in $\text{st}_M(x) \cap N_0$, for example an ϵ -neighbourhood for sufficiently small ϵ . Since N_1 is compact, its image $N'_1 = \phi_0(N_1) = \phi_1(N_1)$ is also compact, and we can find polytopal complexes K_0 and K'_0 whose domains are the considered neighbourhoods $N_1 = |K_0|$ and $N'_1 = |K'_0|$. We subdivide both complexes by taking the intersection complexes $K_1 = K_0 \cap M$ and $K'_1 = K'_0 \cap M'$. Recall from Theorem 1.31 and its proof, that we can find subdivisions K_2 of K_1 and K'_2 of K'_1 such that $\phi_1: |K_2| \rightarrow |K'_2|$ is a cellular bijection with respect to the subdivisions K_2 and K'_2 inducing a combinatorial equivalence between K_2 and K'_2 . This means that a restriction of ϕ_1 yields a cellular bijection between the combinatorially equivalent subcomplexes $\text{st}_{K_2}(x)$ and $\text{st}_{K'_2}(x')$. Since $\phi_1(x) = x'$, this map remains a cellular bijection when we subdivide $\text{st}_{K_2}(x)$ and $\text{st}_{K'_2}(x')$ by elementary starrings at x and x' into $x \text{lk}_{K_2}(x)$ and $x' \text{lk}_{K'_2}(x')$.

We show that $N_2 = |x \text{lk}_{K_2}(x)|$ and $N'_2 = |x' \text{lk}_{K'_2}(x')|$ have all desired properties. By construction, N_2 and N'_2 are cone neighbourhoods with polytopal complexes $\text{lk}_{K_2}(x)$ and $\text{lk}_{K'_2}(x')$ as link complexes. The step where we took the intersection complexes with M and M' to obtain K_1 and K'_1 ensures that each cell of K_1 is contained in a cell of M and that each cell of K'_1 is contained in a cell of M' . Now $x \text{lk}_{K_2}(x)$ and $x' \text{lk}_{K'_2}(x')$ are subdivisions of the subcomplexes $\text{st}_{K_2}(x)$ and $\text{st}_{K'_2}(x')$ of the subdivisions K_2 and K'_2 of K_1 and K'_1 respectively. This implies that each cell of $x \text{lk}_{K_2}(x)$ is contained in a cell of M and that each cell of $x' \text{lk}_{K'_2}(x')$ is contained in a cell of M' , so that f is linear on cells of $x \text{lk}_{K_2}(x)$ and f' is linear on cells of $x' \text{lk}_{K'_2}(x')$. We noted already above that the restriction ϕ_2 of ϕ_1 is a cellular bijection between $x \text{lk}_{K_2}(x)$ and $x' \text{lk}_{K'_2}(x')$. \square

Characterisation Using Link Complexes. When we restrict the PL homeomorphism $\phi_2: N_2 \rightarrow N'_2$ constructed in the preceding proof to the link complexes $|\text{lk}_{K_2}(x)|$ and $|\text{lk}_{K'_2}(x')|$ and apply the cone construction with the apices x and x' to the restriction, we obtain ϕ_2 itself again. Since a PL homeomorphism between link complexes always induces a PL homeomorphism between the corresponding stars, we could ask whether finding a PL homeomorphism

$\phi_{\text{lk}}: |L_x| \rightarrow |L'_{x'}|$ between link complexes L_x for x and $L'_{x'}$ for x' such that $(f' \circ \phi_{\text{lk}}) - f$ has the constant value $f'(x') - f(x)$ throughout L_x is sufficient for a local equivalence of f at x and f' at x' . Although the answer is negative in general, a simple additional condition on the induced cone neighbourhoods xL_x and $x'L'_{x'}$ turns the suggested condition into a characterisation of local equivalence. We have to ensure that the values of the maps on the links and at the points x and x' themselves capture the local behaviour of the maps near x and x' completely. This is the case if f and f' are linear on all the line segments connecting x or x' with points in the links. Since this additional property is inherited by smaller neighbourhoods $|x\tilde{L}_x| \subseteq |xL_x|$ and $|x'\tilde{L}'_{x'}| \subseteq |x'L'_{x'}|$ when fulfilled by the larger ones, we could say informally that the characterisation works for sufficiently small link complexes.

Lemma 3.4. *Let M and M' be polyhedral complexes with maps $f: |M| \rightarrow \mathbb{R}$ and $f': |M'| \rightarrow \mathbb{R}$ linear on cells. Then f at $x \in |M|$ is locally equivalent to f' at $x' \in |M'|$ if and only if there are link complexes L_x for x and $L'_{x'}$ for x' with a PL homeomorphism $\phi_{\text{lk}}: |L_x| \rightarrow |L'_{x'}|$ such that the following properties hold:*

- *The link complex L_x witnesses that f is conical at x in the sense that the equality $f(\lambda x + (1 - \lambda)y) = \lambda f(x) + (1 - \lambda)f(y)$ for $y \in |L_x|$ and $\lambda \in [0, 1]$ characterises the values of f in the cone neighbourhood xL_x .*
- *The link complex $L'_{x'}$ witnesses that f' is conical at x' in the sense that the equality $f'(\lambda x' + (1 - \lambda)y') = \lambda f'(x') + (1 - \lambda)f'(y')$ for $y' \in |L'_{x'}|$ and $\lambda \in [0, 1]$ characterises the values of f' in the cone neighbourhood $x'L'_{x'}$.*
- *The map $(f' \circ \phi_{\text{lk}}) - f$ has the constant value $f'(x') - f(x)$ throughout $|L_x|$.*

For the first property, it is sufficient to choose a link complex L_x whose domain is contained in $|\text{st}_M(x)|$. Analogously, choosing $L'_{x'}$ with $|L'_{x'}| \subseteq |\text{st}_{M'}(x')|$ guarantees the second property.

Proof. For showing that the condition is necessary for local equivalence, we can take the link complexes of the cone neighbourhoods constructed in the previous proof for Remark 3.3. Since f and f' are linear on cells of these cone neighbourhoods, the first two properties are fulfilled. Restricting the PL homeomorphism ϕ_2 from the previous proof to the links yields a PL homeomorphism ϕ_{lk} with the desired properties.

For proving sufficiency, extend ϕ_{lk} by the cone construction to a PL homeomorphism ϕ between the cone neighbourhoods xL_x and $x'L'_{x'}$. Since by assumption f and f' are also given by the cone construction on the respective cone neighbourhoods, the map $(f' \circ \phi) - f$ has not only the constant value $f'(x') - f(x)$ on $|L_x|$ and at x but also throughout $|xL_x|$ by interpolation, because we can compute the value for any $z = \lambda x + (1 - \lambda)y \in |xL_x|$ with $y \in |L_x|$ as follows:

$$\begin{aligned}
((f' \circ \phi) - f)(z) &= f'(\phi(\lambda x + (1 - \lambda)y)) - f(\lambda x + (1 - \lambda)y) \\
&= f'(\lambda x' + (1 - \lambda)\phi_{\text{lk}}(y)) - (\lambda f(x) + (1 - \lambda)f(y)) \\
&= \lambda(f'(x') - f(x)) + (1 - \lambda)(f'(\phi_{\text{lk}}(y)) - f(y)) \\
&= \lambda(f'(x') - f(x)) + (1 - \lambda)(f' \circ \phi_{\text{lk}} - f)(y) \\
&= \lambda(f'(x') - f(x)) + (1 - \lambda)(f'(x') - f(x)) \\
&= f'(x') - f(x)
\end{aligned}$$

Since f is assumed to be linear on cells of M , it is linear on each line segment from x to a point in $|\text{st}_M(x)|$. This proves that f being conical at x is witnessed by any link of x contained in $|\text{st}_M(x)|$. Analogously, f' being conical at x' is witnessed by any link of x' contained in $|\text{st}_{M'}(x')|$ \square

Characterisation in Terms of Combinatorially Equivalent Subdivisions. We can use the fact that the existence of a PL homeomorphism is equivalent to the existence of combinatorially equivalent subdivisions to obtain another characterisation of local equivalence. This characterisation would require combinatorially equivalent link complexes for x and x' , but in general we would still have to check whether the induced PL homeomorphism ϕ_{lk} fulfils $(f' \circ \phi_{\text{lk}}) - f = f'(x') - f(x)$ at every point. When we ensure that f and f' can be computed by interpolation between the vertices of the considered link complexes and the points x and x' itself, then checking the values at vertices suffices.

Remark 3.5. Let M and M' be polyhedral complexes with maps $f: |M| \rightarrow \mathbb{R}$ and $f': |M'| \rightarrow \mathbb{R}$ linear on cells. Then f at $x \in |M|$ is locally equivalent to f' at $x' \in |M'|$ if and only if there are combinatorially equivalent link complexes L_x for x and $L'_{x'}$ for x' with the following three properties: The link complexes witness that f and f' are conical at x and x' respectively; the map f is linear on cells of L_x , and f' is linear on cells of $L'_{x'}$; for any pair of vertices $v \in \text{Vrt}(L_x)$ and $v' \in \text{Vrt}(L'_{x'})$ that correspond under the combinatorial equivalence, the difference of the function values is constantly $f'(v') - f(v) = f'(x') - f(x)$.

Proof. For necessity, we can again use the link complexes constructed in the proof for Remark 3.3. They fulfil all properties by construction.

For showing sufficiency, we can assume that L_x and $L'_{x'}$ are simplicial by subdividing them according to Lemma 1.30 without introducing new vertices. Such a subdivision does not affect the three properties. The PL homeomorphism $\phi_{\text{lk}}: |L_x| \rightarrow |L'_{x'}|$ induced by simplex-wise linear interpolation fulfils $(f' \circ \phi_{\text{lk}})(v) - f(v) = f'(x') - f(x)$ initially at every vertex v of L_x by assumption. Since ϕ_{lk} is simplicial, f linear on cells of L_x , and f' linear on cells of $L'_{x'}$, this property is inherited from the vertices by all points in $|L_x|$ by simplex-wise linear interpolation. This shows that $(f' \circ \phi_{\text{lk}}) - f = f'(x') - f(x)$ holds throughout $|L_x|$ and applying Lemma 3.4 yields local equivalence of f at x and f' at x' . \square

Transitivity of the Local Equivalence Relation. We still have to show that our definition of local equivalence induces an equivalence relation on the class of all triples (M, x, f) where $x \in |M|$ is a point in the domain of a polyhedral complex M with a function $f: |M| \rightarrow \mathbb{R}$ linear on cells. Reflexivity and symmetry are obvious. For transitivity, we can use the observation that the neighbourhoods in the definition can be replaced by smaller neighbourhoods.

Lemma 3.6. *The local equivalence relation is transitive and hence an equivalence relation on the class of all triples (M, x, f) where $x \in |M|$ is a point in the domain of a polyhedral complex M with a function $f: |M| \rightarrow \mathbb{R}$ linear on cells.*

Proof. Assume that f' at $x' \in |M'|$ is locally equivalent to f at $x \in |M|$ by the neighbourhood homeomorphism $\phi: N' \rightarrow N_1$ commuting with $g = f - f(x)$ and

$g' = f' - f'(x')$ and that f at $x \in |M|$ is locally equivalent to f'' at $x'' \in |M''|$ by $\psi: N_2 \rightarrow N''$ commuting with $g = f - f(x)$ and $g'' = f'' - f''(x'')$. Take a neighbourhood N of x contained in both neighbourhoods N_1 and N_2 and replace the neighbourhoods of x' and x'' by the corresponding images of N under the homeomorphism, namely $\tilde{N}' = \phi^{-1}(N)$ and $\tilde{N}'' = \psi(N)$. Then $\psi \circ \phi: \tilde{N}' \rightarrow \tilde{N}''$ gives the desired PL-homeomorphism between neighbourhoods of x' and x'' commuting with g' and g'' . \square

3.1.2 Characterisations Based on Signs Instead of Values

Now we want to show that the precise function values on the links are not really important for checking local equivalence.

Sufficiency of Matching Signs. So far, we have required for a point y in a link of x and its corresponding point $y' = \phi_{\text{lk}}(y)$ in a link of x' , that the values $g(y) = f(y) - f(x)$ and $g'(y') = f'(y') - f'(x')$ agree exactly. But it turns out that it suffices that the signs of $g(y)$ and $g'(y')$ agree to guarantee local equivalence. Rephrased in terms of f and f' , this means that it suffices that the order relations of the function values of f in comparison to the value $f(x)$ and the order relations of f' in comparison to the value $f'(x')$ agree at corresponding points of the links. In the case that the complexes M and M' are polytopal, an additional advantage of this weaker requirement is that we can always assume that the considered link complexes are subdivisions of $\text{lk}_M(x)$ and $\text{lk}_{M'}(x')$. We start by proving sufficiency of such a criterion.

Lemma 3.7. *Consider two polyhedral complexes M and M' with two associated maps $f: |M| \rightarrow \mathbb{R}$ and $f': |M'| \rightarrow \mathbb{R}$ linear on cells. Assume that there are combinatorially equivalent link complexes L for x and L' for x' such that f is linear on cells of xL and f' linear on cells of $x'L'$, for example some combinatorially equivalent subdivisions of $\text{lk}_M(x)$ and $\text{lk}_{M'}(x')$ in polytopal complexes M and M' . Furthermore assume that for each vertex v of L and its corresponding vertex v' of L' under the combinatorial equivalence, the value $f(v)$ has the same order relation in comparison to $f(x)$ as the value $f'(v')$ in comparison to $f'(x')$, that is: $f(v) \geq f(x) \leftrightarrow f'(v') \geq f'(x')$ and $f(v) \leq f(x) \leftrightarrow f'(v') \leq f'(x')$. This assumption can be restated more conveniently in terms of the shifted maps $g = f - f(x)$ and $g' = f' - f'(x')$: Assume that the shifted maps g and g' have the same sign at corresponding vertices of L and L' . Then f at x is locally equivalent to f' at x' .*

Proof. We can assume that L and L' are simplicial by subdividing them according to Lemma 1.30 without introducing new vertices. Recall from Subsection 1.2.8 on pseudo-radial projection that moving the vertices of L and L' on rays emanating from x and x' in direction towards x and x' does not affect the combinatorial equivalence between the complexes. Choose $\epsilon > 0$ smaller than any non-zero absolute value $|g(v)|$ and $|g'(v')|$ of the shifted functions at the vertices v of L and v' of L' . Construct a link L_1 for x from L by keeping all vertices v with $g(v) = 0$ and replacing all other vertices v by the vertex v_1 such that v_1 lies on the ray through v emanating from x and $|g(v_1)| = \epsilon$. By our choice of ϵ , the replacement of v by v_1 corresponds to a movement in direction of x on the ray. Therefore L_1 and L are combinatorially equivalent links




Figure 3.3: Schematic example for the construction of the link complex L_1 for x from the link complex L for x in the proof of Lemma 3.7. At the vertices of L , the symbols $+$, $-$, or 0 indicate that the shifted function g has positive or negative sign, or is zero at the respective vertex. At the corresponding vertices of L_1 the value of g is $+\epsilon$, $-\epsilon$, or zero. The link complexes L and L_1 are combinatorially equivalent by pseudo-radial projection in a similar fashion as the link complexes in Figure 1.4, because L_1 is obtained from L by replacing each vertex by a vertex lying on the same ray emanating from x .

for x by pseudo-radial projection according to Fact 1.32. See Figure 3.3 for a schematic example of such a construction of L_1 from L . Analogously, construct a link L'_1 for x' from L' by keeping all vertices v' with $g'(v') = 0$ and replacing all other vertices v' by the vertex v'_1 such that v'_1 lies on the ray through v' emanating from x' and $|g'(v'_1)| = \epsilon$. Again L'_1 is combinatorially equivalent to L' by pseudo-radial projection. We show that L_1 and L'_1 are combinatorially equivalent link complexes fulfilling the criteria in Remark 3.5.

The combinatorial equivalence between L and L' induces a combinatorial equivalence between L_1 and L'_1 in the obvious way. By construction, each cell S_1 of L_1 is contained in a cell xS , where S is the cell of L which corresponds to S_1 by pseudo-radial projection. Since f is linear on xS by assumption, it is linear on $xS_1 \subseteq xS$ as well. This guarantees that f is linear on cells of the link complex L_1 and furthermore that L_1 witnesses that f is conical at x . Similarly, each cell of L'_1 is contained in a cell of $x'L'$, so that L'_1 witnesses that f' is conical at x' and f' is linear on cells of L'_1 .

The last required property in Remark 3.5 that $f'(v'_1) - f(v_1) = f'(x') - f(x)$ holds at corresponding vertices v_1 of L_1 and v'_1 of L'_1 can be rephrased in terms of the shifted maps g and g' by requiring that $g(v_1) = g'(v'_1)$ holds at corresponding vertices. A classification of the vertices of L and L' into three classes according to the sign of g and g' , induces a partition into three classes for the vertices of

L_1 and L'_1 : For corresponding vertices v of L and v' of L' with $g(v) = g'(v') = 0$, we chose $v_1 = v$ in L_1 and $v'_1 = v'$ in L'_1 , so that for such pairs of corresponding vertices the equality $g(v_1) = g'(v'_1)$ still holds. For corresponding vertices v of L and v' of L' with $g(v) > 0$ and $g'(v') > 0$, we chose the corresponding vertices v_1 of L_1 and v'_1 of L'_1 in such a way that $g(v_1) = \epsilon = g'(v'_1)$ holds. And finally, for corresponding vertices v of L and v' of L' with $g(v) < 0$ and $g'(v') < 0$, we chose the corresponding vertices v_1 of L_1 and v'_1 of L'_1 in such a way that $g(v_1) = -\epsilon = g'(v'_1)$ holds. Since we assumed that the signs of g and g' agree at corresponding vertices of L and L' , all possible cases are covered and the proof is finished. \square

The Other Direction. It remains to show that a criterion as in the previous lemma is also necessary for local equivalence.

Lemma 3.8. *If (M, x, f) and (M', x', f') are locally equivalent, then any pair of link complexes L for x and L' for x' such that $|L| \subseteq |\text{st}_M(x)|$ and $|L'| \subseteq |\text{st}_{M'}(x')|$, in particular for polytopal complexes M and M' the pair $L = \text{lk}_M(x)$ and $L' = \text{lk}_{M'}(x')$, has combinatorially equivalent subdivisions such that corresponding vertices have the same order relation in comparison to the values at x and x' respectively.*

Proof. Apply the characterisation of local equivalence from Lemma 3.4 to the pair of triples (M, x, f) and (M', x', f') . This yields link complexes L_x for x and $L'_{x'}$ for x' witnessing f and f' as conical at x and x' respectively and a PL homeomorphism $\phi_{\text{lk}}: |L_x| \rightarrow |L'_{x'}|$ with $g' \circ \phi_{\text{lk}} = g$ on $|L_x|$, where $g = f - f(x)$ and $g' = f' - f'(x')$ as usual. Then we use pseudo-radial projection to find combinatorially equivalent subdivisions \tilde{L}_x of L_x and \tilde{L} of L as described in Subsection 1.2.8: We take the polyhedral fans L_x^+ and L^+ and construct the intersection complexes $\tilde{L}_x = L_x^+ \cap L^+ \cap L_x$ and $\tilde{L} = L_x^+ \cap L^+ \cap L$. By construction, the link complexes \tilde{L}_x and \tilde{L} induce the same polyhedral fan $L_x^+ \cap L^+$. An analogous construction yields combinatorially equivalent subdivisions $\tilde{L}'_{x'}$ of $L'_{x'}$ and \tilde{L}' of L' .

As a next step, we apply the fact that PL-homeomorphic complexes can be subdivided in such a way that the PL homeomorphism is a simplicial bijection with respect to these subdivisions. We obtain subdivisions \bar{L}_x of \tilde{L}_x and $\bar{L}'_{x'}$ of $\tilde{L}'_{x'}$ such that $\phi_{\text{lk}}: |\bar{L}_x| \rightarrow |\bar{L}'_{x'}|$ is simplicial. These subdivisions and the simplicial map ϕ_{lk} can be lifted by the pseudo-radial projections between \tilde{L}_x and \tilde{L} on the one hand and between $\tilde{L}'_{x'}$ and \tilde{L}' on the other hand to subdivisions \bar{L} of \tilde{L} and \bar{L}' of \tilde{L}' and a simplicial bijection $\psi: |\bar{L}| \rightarrow |\bar{L}'|$ as follows: The polyhedral fan \bar{L}_x^+ is a subdivision of the polyhedral fan $\tilde{L}_x^+ = \tilde{L}^+$. Therefore pseudo-radial projection of the simplicial link complex \bar{L}_x to \tilde{L} yields a simplicial subdivision of \tilde{L} , namely $\bar{L} = \bar{L}_x^+ \cap \tilde{L}$. With this construction, we can observe that \bar{L} and \bar{L}_x are combinatorially equivalent by pseudo-radial projection. In particular, a vertex \bar{v} of \bar{L} corresponds to a vertex \bar{v}_x of \bar{L}_x under this equivalence if and only if the extended cones \bar{v}^+ and \bar{v}_x^+ define the same ray emanating from x . The subdivision \bar{L}' of \tilde{L}' is obtained analogously by pseudo-radial projection of $\bar{L}'_{x'}$ to \tilde{L}' . The simplicial map $\psi: |\bar{L}| \rightarrow |\bar{L}'|$ is merely the composite of the three simplicial bijections inherent to our construction: first, the simplicial bijection induced by pseudo-radial projection from $|\bar{L}|$ and $|\bar{L}_x|$, then the simplicial bijection ϕ_{lk} from $|\bar{L}_x|$ to $|\bar{L}'_{x'}|$, and finally the simplicial bijection induced by pseudo-radial projection from $|\bar{L}'_{x'}|$ to $|\bar{L}'|$.

We show that \bar{L} and \bar{L}' are combinatorially equivalent subdivisions of L and L' with the desired property that the maps g and g' have the same sign at corresponding vertices of \bar{L} and \bar{L}' . The combinatorial equivalence is given by the simplicial bijection ψ . A vertex \bar{v} of \bar{L} corresponds to a vertex \bar{v}' of \bar{L}' under ψ if and only if their counterparts by pseudo-radial projection, that is, the vertices \bar{v}_x in \bar{L}_x and $\bar{v}'_{x'}$ in $\bar{L}'_{x'}$, with the property that \bar{v}_x and \bar{v} lie on the same ray emanating from x and that $\bar{v}'_{x'}$ and \bar{v}' lie on the same ray emanating from x' , correspond to each other via ϕ_{lk} . So let \bar{v} and \bar{v}' such a pair of corresponding vertices in \bar{L} and \bar{L}' with counterparts \bar{v}_x in \bar{L}_x and $\bar{v}'_{x'}$ in $\bar{L}'_{x'}$. Since $g' \circ \phi_{\text{lk}}$ agrees with g on $|\bar{L}_x| = |L_x|$, the function values $g(\bar{v}_x)$ and $g'(\bar{v}'_{x'})$ at the counterparts agree. Therefore it suffices to show that $g(\bar{v}_x)$ has the same sign as $g(\bar{v})$, and that $g'(\bar{v}'_{x'})$ has the same sign as $g'(\bar{v}')$. We only show the former claim, the latter follows from a completely analogous proof.

The vertices \bar{v} and \bar{v}_x lie on the same ray emanating from x , hence one of the line segments $x\bar{v}$ and $x\bar{v}_x$ is contained in the other. The map f is linear on both line segments: It is linear on $x\bar{v}$ because v is contained in $|L|$ which in turn is assumed to be contained in $|\text{st}_M(x)|$; it is linear on $x\bar{v}_x$ because L_x witnesses that f is conical at x . Clearly, the shifted map g is also linear on both line segments, and its value at x is zero. Since a linear map on a line segment is either constant or uniformly strictly increasing or decreasing in a given direction, we can conclude that when we follow the ray emanating from x in direction of v and v_x , pass through the point with smaller distance to x of the two points, and continue until we reach the point with larger distance, then the values of the map g along the path will either be constantly zero, implying $g(v) = g(v_x) = 0$, or they will be strictly increasing, implying that $g(v)$ and $g(v_x)$ are both positive, or they will be strictly decreasing, implying that $g(v)$ and $g(v_x)$ are both negative. This completes the proof. \square

From Matching Signs at Vertices to Matching Signs at All Points. In the previous two lemmata, the criterion asks only for matching signs of g and g' at corresponding vertices of combinatorially equivalent subdivisions. When we switch from combinatorial equivalences to PL homeomorphisms, we could ask for matching signs at all corresponding points. Such a characterisation of local equivalence is possible and stated in the next lemma, but a PL homeomorphism with that property cannot be obtained immediately from a combinatorial equivalence with matching signs only by interpolation. For example, two edges that each have one endpoint with negative value and one endpoint with positive value are combinatorially equivalent with matching signs at the vertices. But the intermediate points on the edges where the functions attain the value zero are not necessarily corresponding under the induced simplicial map between the two edges. For obtaining a PL homeomorphism between the edges that respects the sign at all intermediate points, we have to subdivide both edges into a non-negative and a non-positive part at the point where zero is attained and then interpolate on the non-negative and non-positive parts individually.

Lemma 3.9. *Let M and M' be polyhedral complexes with maps $f: |M| \rightarrow \mathbb{R}$ and $f': |M'| \rightarrow \mathbb{R}$ linear on cells. Assume that L is a link complex for $x \in |M|$ whose domain is contained in $|\text{st}_M(x)|$, and analogously that L' is a link complex for $x' \in |M'|$ whose domain is contained in $|\text{st}_{M'}(x')|$. Then f at x is locally equivalent to f' at x' if and only if there is a PL homeomorphism*

$\phi_{\text{lk}}: |L| \rightarrow |L'|$ such that ϕ_{lk} maps $|L_{\geq f(x)}|$ to $\phi_{\text{lk}}(|L_{\geq f(x)}|) = |L'_{\geq f'(x')}|$ and $|L_{\leq f(x)}|$ to $\phi_{\text{lk}}(|L_{\leq f(x)}|) = |L'_{\leq f'(x')}|$, in other words if and only if there is a PL homeomorphism between L and L' such that $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding points in $|L|$ and $|L'|$.

Proof. We assume first local equivalence and construct a PL homeomorphism. Subdivide L into $\tilde{L} = L_{\leq f(x)} \cup L_{\geq f(x)}$ and L' into $\tilde{L}' = L'_{\leq f'(x')} \cup L'_{\geq f'(x')}$. By Lemma 3.8, the complexes \tilde{L} and \tilde{L}' have combinatorially equivalent subdivisions \bar{L} and \bar{L}' such that $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding vertices. We can assume that \bar{L} and \bar{L}' are simplicial using Lemma 1.30. Then a PL homeomorphism $\phi_{\text{lk}}: |\bar{L}| \rightarrow |\bar{L}'|$ is induced by simplex-wise linear interpolation of the combinatorial equivalence in terms of vertices.

By construction, \bar{L} contains $\bar{L}_{=f(x)}$, $\bar{L}_{\leq f(x)}$ and $\bar{L}_{\geq f(x)}$ as subcomplexes, and analogously, \bar{L}' contains $\bar{L}'_{=f'(x')}$, $\bar{L}'_{\leq f'(x')}$ and $\bar{L}'_{\geq f'(x')}$ as subcomplexes. Since $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding vertices of \bar{L} and \bar{L}' and on each simplex the respective map attains either only non-negative or only non-positive values, the signs also agree at the interpolated corresponding points as follows: The sign of $f(y) - f(x)$ for $y \in |\bar{L}|$ is zero if and only if y is contained in a simplex of $\bar{L}_{=f(x)}$; a simplex belongs to $\bar{L}_{=f(x)}$ if and only if its corresponding simplex under the combinatorial equivalence belongs to $\bar{L}'_{=f'(x')}$, meaning that y is contained in a simplex of $\bar{L}_{=f(x)}$ if and only if $\phi_{\text{lk}}(y)$ is contained in a simplex of $\bar{L}'_{=f'(x')}$; but $\phi_{\text{lk}}(y)$ is contained in a simplex of $\bar{L}'_{=f'(x')}$ if and only if the sign of $f'(\phi_{\text{lk}}(y)) - f'(x')$ is zero. For a positive sign, we can argue that $f(y) - f(x)$ for $y \in |\bar{L}|$ is positive if and only if y is contained in the interior of a simplex of $\bar{L}_{\geq f(x)}$ that does not belong to $\bar{L}_{=f(x)}$; a simplex of \bar{L} belongs to $\bar{L}_{\geq f(x)}$ but not to $\bar{L}_{=f(x)}$ if and only if the corresponding simplex in \bar{L}' belongs to $\bar{L}'_{\geq f'(x')}$ but not to $\bar{L}'_{=f'(x')}$, meaning that y is contained in the interior of a simplex of $\bar{L}_{\geq f(x)} \setminus \bar{L}_{=f(x)}$ if and only if $\phi_{\text{lk}}(y)$ is contained in a simplex of $\bar{L}'_{\geq f'(x')} \setminus \bar{L}'_{=f'(x')}$; but $\phi_{\text{lk}}(y)$ is contained in the interior of a simplex of $\bar{L}'_{\geq f'(x')} \setminus \bar{L}'_{=f'(x')}$ if and only if $f'(\phi_{\text{lk}}(y)) - f'(x')$ has positive sign. A symmetric argument using $\bar{L}_{\leq f(x)}$ and $\bar{L}'_{\leq f'(x')}$ yields the desired equivalence for negative signs at corresponding points.

Now we show that local equivalence is implied by a PL homeomorphism $\phi_{\text{lk}}: |L| \rightarrow |L'|$ respecting the signs at corresponding points. When we subdivide L into \tilde{L} by taking its intersection complex with $x \text{lk}_M(x)$, the map f is linear on cells of $x\tilde{L}$. For an analogous subdivision \tilde{L}' of L' , the map f' is linear on cells of $x\tilde{L}'$. Regarding ϕ_{lk} as PL homeomorphism from $|\tilde{L}|$ to $|\tilde{L}'|$, we can find combinatorially equivalent simplicial subdivisions \bar{L} of \tilde{L} and \bar{L}' of \tilde{L}' such that the restriction of ϕ_{lk} to the vertex set of \bar{L} forms the combinatorial equivalence in terms of vertices between the two subdivisions. Then the link complexes \bar{L} and \bar{L}' fulfil all conditions required in Lemma 3.7 for ensuring local equivalence, namely that the signs of $f - f(x)$ and $f' - f'(x')$ agree at corresponding vertices, that f is linear on cells of $x\bar{L}$, and that f' is linear on cells of $x'\bar{L}'$. \square

When the considered complexes are polytopal, the previous lemma can be applied in particular to $L = \text{lk}_M(x)$ and $L' = \text{lk}_{M'}(x')$.

Corollary 3.10. *For polytopal complexes M and M' with maps $f: |M| \rightarrow \mathbb{R}$ and $f': |M'| \rightarrow \mathbb{R}$ linear on cells, f at $x \in |M|$ is locally equivalent to f' at $x' \in |M'|$ if and only if there is a PL homeomorphism $\phi_{\text{lk}}: |\text{lk}_M(x)| \rightarrow |\text{lk}_{M'}(x')|$ between the links of x and x' mapping the lower level link $|\text{lk}_M(x)|_{\leq f(x)}$ of x to the lower level link $|\text{lk}_{M'}(x')|_{\leq f'(x')}$ of x' and the upper level link $|\text{lk}_M(x)|_{\geq f(x)}$ of x to the upper level link $|\text{lk}_{M'}(x')|_{\geq f'(x')}$ of x' .*

Proof. Apply Lemma 3.9 to $L = \text{lk}_M(x)$ and $L' = \text{lk}_{M'}(x')$. \square

The PL homeomorphism ϕ_{lk} in Lemma 3.9 can be extended by the cone construction to a PL homeomorphism $\phi: |xL| \rightarrow |x'L'|$ between the corresponding cone neighbourhoods, and this homeomorphism obviously respects the signs of $f - f(x)$ and $f' - f'(x')$ at every point of the cone neighbourhood if ϕ_{lk} respects the signs on the links.

Summary. The next theorem summarises the criteria for local equivalence we collected in the preceding results. Note that some of the criteria establish a certain kind of trade-off between exactly matching values and matching signs on the one hand and the mere existence of suitable corresponding neighbourhoods and a freedom to choose a pair of sufficiently small neighbourhoods on the other hand: For the criteria where the PL homeomorphism is only required to respect the signs, such a homeomorphism exists between any pair of sufficiently small cone neighbourhoods with polytopal link complexes if the maps are locally equivalent at the points. For the criteria that require exactly matching values at corresponding points, we can only ensure for sufficiently small neighbourhoods that a matching neighbourhood at the other point exists if the maps are locally equivalent.

Theorem 3.11. *The following conditions equivalently characterise the property that (M, x, f) and (M', x', f') are locally equivalent triples of a polytopal complex, a point in its domain, and a map linear on its cells. All but the last two criteria also work for polyhedral complexes.*

1. *There is a PL homeomorphism between some cone neighbourhoods of x and x' mapping x and x' to each other such that $f - f(x)$ and $f' - f'(x')$ commute with the PL homeomorphism.*
2. *There are combinatorially equivalent cone neighbourhoods of x and x' such that x corresponds to x' and the maps $f - f(x)$ and $f' - f'(x')$ are linear on cells of the cone neighbourhoods and agree at corresponding vertices.*
3. *There are PL-homeomorphic link complexes for x and x' contained in the stars of x and x' such that $f - f(x)$ and $f' - f'(x')$ commute with the PL homeomorphism.*
4. *There are combinatorially equivalent link complexes for x and x' contained in the stars of x and x' such that $f - f(x)$ and $f' - f'(x')$ are linear on cells of the link complexes and agree at corresponding vertices.*
5. *There are PL-homeomorphic link complexes for x and x' contained in the stars of x and x' such that $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding points.*

6. *There are combinatorially equivalent link complexes for x and x' contained in the stars of x and x' such that $f - f(x)$ and $f' - f'(x')$ are linear on cells of the link complexes and have the same sign at corresponding vertices.*
7. *Any pair of link complexes for x and x' that are contained in the stars of x and x' is PL-homeomorphic in such a way that $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding points.*
8. *Any pair of link complexes for x and x' that are contained in the stars of x and x' has combinatorially equivalent subdivisions such that $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding vertices.*
9. *There is a PL homeomorphism between $\text{lk}(x)$ and $\text{lk}(x')$ mapping upper level link to upper level link and lower level link to lower level link, that is, $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding points.*
10. *The links $\text{lk}(x)$ and $\text{lk}(x')$ have combinatorially equivalent subdivisions such that $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding vertices.*

3.2 Regular and Critical Points

In smooth Morse theory, a Morse function is defined as a smooth map whose critical points are all non-degenerate. For an analogous definition of piecewise linear Morse functions (Definition 3.13), we suggest in Definition 3.12 a classification of the points in the domain of a piecewise linear function into regular points, non-degenerate critical points, and degenerate critical points based on the local equivalence class of the map at the point. We show in Lemma 3.22 that a map linear on cells of a combinatorial manifold and in general position has only critical points that are vertices of the combinatorial manifold. Furthermore, any combinatorial manifold has a piecewise linear Morse function in the sense of the proposed definition as stated in Theorem 3.24: For any derived subdivision of the combinatorial manifold, a PL Morse function linear on cells of the subdivision can be constructed.

We also compare our definition of regular and critical points with previously suggested and studied notions. The definitions given by Kosinski in [29] and by Eells and Kuiper in [17] are basically equivalent to ours. Other variants amount to a weakening of the requirements for what we called local equivalence from a topological condition to a condition based only on homology. The definition by Brehm and Kühnel [9] considers the homology with coefficients in a field of the lower level link $(\text{lk}(v))_{\leq f(v)}$ and classifies the vertices by the Betti numbers. Banchoff's approach in [4] amounts to classify the vertices by the Euler characteristic of the lower level link.

The discrete Morse theory developed by Forman [18, 19] differs fundamentally from the other approaches. The functions considered are not PL functions on the domain, but functions that assign numbers to cells. Nevertheless some links between discrete Morse theory and PL Morse theory can be found: Our proof for the existence of PL Morse functions on the derived of any combinatorial manifold exhibits some similarities with the existence proof for discrete Morse functions that establishes such a function with all cells being critical. In fact, it can be viewed as a special case of a more general method described in

Lemma 3.26 for constructing PL Morse functions from discrete Morse functions preserving criticality in a certain sense.

3.2.1 Reference Triples

Recall that regular points and non-degenerate critical points in smooth Morse theory can be characterised as points where a suitable local coordinate system allows to represent the map f by the standard example of its kind. Regular points p have a p -centred local coordinate system (X_1, \dots, X_d) in terms of which the map is represented by $f(p) + X_1$, and non-degenerate critical points p have a p -centred local coordinate system (X_1, \dots, X_d) in terms of which the map is represented by $f(p) - X_1^2 - \dots - X_\lambda^2 + X_{\lambda+1}^2 + \dots + X_d^2$. At least for regular points, this characterisation can easily be transferred to the piecewise linear setting. For smooth Morse theory, the local coordinate system is understood to be a smooth coordinate system, and the piecewise linear analogue is to require a piecewise linear local coordinate system instead. This suggests the following characterisation for a PL-ordinary or PL-regular point p of a piecewise linear map f on the domain of a polyhedral complex M : The point is PL-regular if and only if there is a p -centred piecewise linear local coordinate system (X_1, \dots, X_d) such that f is represented by the term $f(p) + X_1$ in these coordinates.

We can rephrase this condition using the notion of local equivalence introduced in the previous section. The existence of a p -centred piecewise linear local coordinate system (X_1, \dots, X_d) such that f is represented by a map $f'(X_1, \dots, X_d)$ in terms of these coordinates is equivalent to (M, p, f) being locally equivalent to some triple $(U, 0, f')$ with U being some suitable polyhedral complex whose domain is a neighbourhood of 0 in \mathbb{R}^d where the map f' is defined. Since we can neglect the shifting by the constant $f(p)$ when using local equivalence for the characterisation, we obtain that p is a PL-regular point for f if and only if (M, p, f) is locally equivalent to $(\mathbb{R}^d, 0, f')$ with $f'(X_1, \dots, X_d) = X_1$. In this fashion, we can classify points in the domain of a polyhedral complex with a real-valued map linear on cells according to the local equivalence class. Certain triples (M, x, f) are treated as reference to represent their local equivalence class. The reference triples used here consist of a suitable subdivision of a Euclidean space or half-space, the origin as considered point, and a standard map whose level sets near the origin behave as desired.

For the piecewise linear analogue of the non-degenerate critical points, we cannot adopt the standard reference maps used in the smooth case because a sum of signed squares of the coordinates is not a piecewise linear map. Replacements for these maps should be chosen in such a way that the level sets have the same topological features near the point as their smooth counterparts. A natural candidate for a piecewise linear analogue of the square of a coordinate is its absolute value. Hence we could replace a sum of signed squares of coordinates $-X_1^2 - \dots - X_\lambda^2 + X_{\lambda+1}^2 + \dots + X_d^2$ by a sum of correspondingly signed absolute values of the coordinates and define a PL non-degenerate critical point of index λ as a point at which f is locally equivalent to the map $-|X_1| - \dots - |X_\lambda| + |X_{\lambda+1}| + \dots + |X_d|$ at the origin of \mathbb{R}^d .

In a paper on the construction of contour trees by Chiang, Lenz, Lu, and Rote [13], also points on the boundary of a 2- or 3-dimensional combinatorial manifold with a real-valued map linear on cells are taken into account for the classification into regular and critical points. Inspired by this classification, we

also suggest standard functions for non-degenerate boundary points. However, apart from some basic observations on such points, we do not investigate PL Morse theory for manifolds with boundary any further.

Definition 3.12 (regular and critical points). The following reference maps are used as standard examples for defining regular and non-degenerate critical points:

$$\begin{aligned}\pi_j^d(X_1, \dots, X_d) &= X_j \\ f_\lambda^d(X_1, \dots, X_d) &= -|X_1| - \dots - |X_\lambda| + |X_{\lambda+1}| + \dots + |X_d| \\ g_\lambda^d(X_1, \dots, X_d) &= -|X_1| - \dots - |X_\lambda| + |X_{\lambda+1}| + \dots + |X_{d-1}| - |X_d|\end{aligned}$$

The maps π_j^d and f_λ^d occur in reference triples for both interior and boundary points, whereas the maps g_λ^d are only used for boundary points. For the globally linear coordinate maps π_j^d we use the Euclidean space \mathbb{R}^d or the Euclidean half-space $\mathbb{R}_+^d = \mathbb{R}_{\pi_1^d \geq 0}^d$ without further subdivision, that is, the induced complexes $\mathcal{F}(\mathbb{R}^d) = \{\emptyset, \mathbb{R}^d\}$ or $\mathcal{F}(\mathbb{R}_+^d) = \{\emptyset, \mathbb{R}^{d-1}, \mathbb{R}_+^d\}$ where \mathbb{R}^{d-1} is identified with the hyperplane $\mathbb{R}_{\pi_1^d=0}^d$ bounding the half-space. The maps f_λ^d and g_λ^d are linear on each orthant of \mathbb{R}^d . Hence we subdivide the Euclidean space or half-space into orthants by the hyperplanes $X_j = 0$ for each j , yielding the subdivisions Ω^d of \mathbb{R}^d and Ω_+^d of \mathbb{R}_+^d . The maps f_λ^d are considered on both subdivisions Ω^d and Ω_+^d ; the maps g_λ^d are only considered on the subdivision Ω_+^d of the half-space \mathbb{R}_+^d . These subdivisions can be defined formally as intersection complexes of level set complexes with respect to the coordinate functions and the interval complex $\mathcal{I} = \{\emptyset, (-\infty, 0], \{0\}, [0, \infty)\}$ as follows:

$$\Omega^d = \bigcap_{j=1}^d (\mathcal{F}(\mathbb{R}^d))_{\pi_j^d \in \mathcal{I}} \quad \Omega_+^d = \bigcap_{j=1}^{d-1} (\mathcal{F}(\mathbb{R}_+^d))_{\pi_j^d \in \mathcal{I}}$$

In each case, we use the origin 0 of \mathbb{R}^d as reference point.

Let M be a polyhedral complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. A point $x \in |M|$ is called a *PL-regular* or *PL-ordinary* (interior) point of f (in dimension $d > 0$) if (M, x, f) is locally equivalent to $(\mathbb{R}^d, 0, \pi_1^d)$; it is called a *PL-regular boundary point* of f (in dimension $d > 1$) if (M, x, f) is locally equivalent to $(\mathbb{R}_+^d, 0, \pi_1^d)$. All points that are not PL-regular points of f are called *PL-critical* points of f . Certain types of PL-critical points listed below are called *PL non-degenerate critical points*, the remaining critical points are *PL degenerate critical points*: A point $x \in |M|$ is called a *PL non-degenerate critical* (interior) point of f (in dimension $d \geq 0$) of index λ if (M, x, f) is locally equivalent to $(\Omega^d, 0, f_\lambda^d)$ where the index λ is an integer in the range $0 \leq \lambda \leq d$; it is called a *PL non-degenerate boundary-critical point* of f (in dimension $d > 0$) of index λ for $0 \leq \lambda < d$ if it is locally equivalent to $(\Omega_+^d, 0, f_\lambda^d)$ or $(\Omega_+^d, 0, g_\lambda^d)$; when we want to distinguish the two types of boundary-critical points, we call points with f_λ^d as reference map of lower type, and points with g_λ^d as reference map of upper type. We usually omit the prefix ‘‘PL’’ and call the points regular, non-degenerate critical, and so on. A value $h \in \mathbb{R}$ is called a *critical value* if its preimage $f^{-1}(h)$ contains a critical point, otherwise it is called a *regular value*.

Definition 3.13 (PL Morse function). A map $f: |M| \rightarrow \mathbb{R}$ linear on cells of a polyhedral complex M is called a *PL Morse function*, if for some fixed dimension d all points in $|M|$ are PL-regular or PL non-degenerate critical points in dimension d .

Remark 3.14. Note that if M has a map $f: |M| \rightarrow \mathbb{R}$ that is a PL Morse function, then $|M|$ is necessarily a PL manifold.

Proof. Since f is at each point locally equivalent to one of the standard maps at the origin of some d -dimensional Euclidean space or half-space for some fixed dimension d , each point has a neighbourhood PL-homeomorphic to a neighbourhood of the origin in \mathbb{R}^d or \mathbb{R}_+^d . This means that $|M|$ is a PL d -manifold, and the interior points of $|M|$ are exactly the regular interior points and non-degenerate critical interior points of f , and the boundary points of $|M|$ are exactly the regular boundary points and non-degenerate boundary-critical points of f . \square

Unambiguity of the Classification. As a first step, we check that the classification of the points in Definition 3.12 is unambiguous in the sense that for each point $x \in |M|$ and each map $f: |M| \rightarrow \mathbb{R}$ linear on cells of M at most one of the reference triples used in Definition 3.12 is locally equivalent to (M, x, f) . We consider the standard cross-polytope as link complex for the origin and show that no PL homeomorphism can map the lower level link with respect to one reference map to the lower level link with respect to another reference map. In fact two such lower level links are not even homotopy equivalent since they have different Betti numbers. For distinguishing the cases for boundary points we do not only consider the link as a whole but also its restriction to the boundary. The occurring lower level links are homotopy equivalent to either a single point, or the empty set, or some λ -sphere. We recall some facts that we need for the proof.

Fact 3.15. Any cone is collapsible, that is, for any complex K , the cone aK collapses to the apex a [21, Example 2, p. 49]. Any maximal cell S of K is a free face of aS , and collapsing such a pair yields the cone $a(K \setminus \{S\})$, hence induction on the number of cells of K proves the claim.

Fact 3.16. A single point, the empty set, and the λ -spheres for $\lambda \geq 0$ each have different sets of Betti numbers. Therefore any two spaces that are each homotopy equivalent to one of the above spaces but not to the same one are neither homotopy equivalent nor PL-homeomorphic to each other.

For a point and the empty set the homology groups and the Betti numbers can be easily computed using simplicial homology with integer coefficients and the complex $\{\emptyset\}$ for the empty set and a complex $\mathcal{F}(v) = \{\emptyset, \{v\}\}$ with some vertex v for a point: For the complex $\{\emptyset\}$, all chain groups $C_n(\{\emptyset\})$, boundary maps δ_n , and consequently all homology groups $H_n(\{\emptyset\})$ are zero, yielding $\beta_n(\{\emptyset\}) = 0$ for all n . For $\mathcal{F}(v)$, the only non-zero chain group is $C_0(\mathcal{F}(v))$ which is isomorphic to \mathbb{Z} , and all boundary maps are zero, yielding $H_0(\mathcal{F}(v))$ isomorphic to \mathbb{Z} as only non-zero homology group and $\beta_0(\mathcal{F}(v)) = 1$ as only non-zero Betti number.

For the computation of the homology groups of the λ -spheres S^λ , we refer to Hatcher [25, Corollary 2.14, p. 114] stating for the reduced homology groups that $\tilde{H}_\lambda(S^\lambda)$ is isomorphic to \mathbb{Z} and that all other groups $\tilde{H}_n(S^\lambda)$ for $n \neq \lambda$ are zero. For transferring this result to unreduced homology, recall from Subsection 1.2.10

Table 3.1: The first five Betti numbers β_0 to β_4 for a point, the empty set, and the λ -spheres for $\lambda \leq 4$.

	β_0	β_1	β_2	β_3	β_4	\dots
point	1	0	0	0	0	\dots
empty set	0	0	0	0	0	\dots
0-sphere	2	0	0	0	0	\dots
1-sphere	1	1	0	0	0	\dots
2-sphere	1	0	1	0	0	\dots
3-sphere	1	0	0	1	0	\dots
4-sphere	1	0	0	0	1	\dots
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots

the fact for non-empty spaces X , that the unreduced homology group $H_0(X)$ is isomorphic to the direct sum of \mathbb{Z} with the reduced homology group $\tilde{H}_0(X)$ and that for all other $n \neq 0$, the unreduced homology group $H_n(X)$ is isomorphic to the reduced homology group $\tilde{H}_n(X)$ [25, p. 110]. This gives the Betti numbers $\beta_0(S^0) = 2$, $\beta_0(S^\lambda) = 1$ for $\lambda > 0$, $\beta_\lambda(S^\lambda) = 1$ for $\lambda > 0$, and $\beta_n(S^\lambda) = 0$ for $n \neq 0$ and $n \neq \lambda$. The Betti numbers β_n with $0 \leq n \leq 4$ for a point, the empty set, and S^λ with $\lambda \leq 4$ are listed in Table 3.1.

Lemma 3.17. *Any two distinct reference triples used in Definition 3.12 are not locally equivalent, namely the triples $(\mathbb{R}^d, 0, \pi_1^d)$ for regular interior points, the triples $(\mathbb{R}_+^d, 0, \pi_1^d)$ for regular boundary points, the triples $(\Omega^d, 0, f_\lambda^d)$ for non-degenerate critical interior points, and the triples $(\Omega_+^d, 0, f_\lambda^d)$ and $(\Omega_+^d, 0, g_\lambda^d)$ for non-degenerate boundary-critical points.*

Proof. Any link complex of the origin in \mathbb{R}^d is a $(d-1)$ -sphere, and any link complex of the origin in \mathbb{R}_+^d is a $(d-1)$ -ball. Since local equivalence implies that the link complexes of the considered points are PL-homeomorphic, two reference triples could only be locally equivalent if the dimensions match and either both triples classify interior points or both triples classify boundary points. Note that all reference maps have the value 0 at the origin. Furthermore for simplifying the proof, we can replace the reference triple $(\mathbb{R}^d, 0, \pi_1^d)$ by the locally equivalent triple $(\Omega^d, 0, \pi_1^d)$ and the reference triple $(\mathbb{R}_+^d, 0, \pi_1^d)$ by the locally equivalent triple $(\Omega_+^d, 0, \pi_1^d)$. The local equivalence of these two pairs of triples is due to the fact that for any subdivision K of a complex M , the triples (M, x, f) and (K, x, f) are locally equivalent as witnessed by the identity $\text{id}: |M| \rightarrow |K|$ between the neighbourhoods $|M| = |K|$ of x . Let us deal with triples classifying interior points first.

Interior Points. As link complex for the origin in \mathbb{R}^d , we choose the boundary of the standard cross-polytope, in other words the unit sphere S_1^{d-1} with respect to the 1-norm around the origin. It can be considered as a simplicial complex in a natural way: For each orthant or \mathbb{R}^d , the intersection with S_1^{d-1} is a simplex spanned by the end points of the unit vectors spanning the orthant, and these simplices are the maximal simplices of a simplicial complex triangulat-

ing S_1^{d-1} . The vertices of this triangulation are the points whose coordinates are those of the standard base vectors or their negatives, that is, all but one of the coordinates are zero and the exceptional value is 1 or -1 . We denote by e_i the vertex with coordinate $X_i = 1$ and by $-e_i$ the vertex with coordinate $X_i = -1$.

So let f and g be two distinct functions among the reference maps π_1^d and f_λ^d . If $(\Omega^d, 0, f)$ were locally equivalent to $(\Omega^d, 0, g)$, Lemma 3.9 would imply that there is a PL homeomorphism $\phi_{\text{lk}}: S_1^{d-1} \rightarrow S_1^{d-1}$ mapping $(S_1^{d-1})_{f \leq 0}$ homeomorphically to $(S_1^{d-1})_{g \leq 0}$. We will show that $(S_1^{d-1})_{\pi_1^d \leq 0}$ is a collapsible PL ball, hence homotopy equivalent to a point, $(S_1^{d-1})_{f_0^d \leq 0}$ is empty, and $(S_1^{d-1})_{f_\lambda^d \leq 0}$ is homotopy equivalent to a $(\lambda - 1)$ -sphere for $1 \leq \lambda \leq d$, rendering the existence of ϕ_{lk} impossible because of Fact 3.16.

The values of the standard maps at the vertices of S_1^{d-1} are easy to compute. The map π_1^d is zero at all but two vertices of S_1^{d-1} , with $\pi_1^d(e_1) = 1$ and $\pi_1^d(-e_1) = -1$ being the exceptions. Since no simplex of S_1^{d-1} contains both e_1 and $-e_1$, the lower level complex $(S_1^{d-1})_{\pi_1^d \leq 0}$ is the subcomplex of S_1^{d-1} consisting of those simplices that do not contain e_1 , in other words the simplices lying in the half-space $X_1 \leq 0$. Similarly, $(S_1^{d-1})_{\pi_1^d = 0}$ is the $(d - 2)$ -dimensional unit sphere with respect to the 1-norm in the subspace $X_1 = 0$ and the subcomplex of S_1^{d-1} consisting of the simplices lying in that subspace. Moreover, $(S_1^{d-1})_{\pi_1^d \leq 0}$ is the cone on $(S_1^{d-1})_{\pi_1^d = 0}$ with apex $-e_1$. Therefore $(S_1^{d-1})_{\pi_1^d \leq 0}$ is a $(d - 1)$ -ball, and it collapses to the apex $-e_1$ by Fact 3.15.

The map f_0^d has the value 1 at all vertices of S_1^{d-1} , hence $(S_1^{d-1})_{f_0^d \leq 0}$ is empty. For maps f_λ^d with $\lambda > 0$, we apply Lemma 2.46 stating for polytopal complexes M that $M_{\leq h}$ collapses to $\text{scp}_{\leq h}(M)$. Accordingly, $(S_1^{d-1})_{f_\lambda^d \leq 0}$ collapses to and is homotopy equivalent to the subcomplex $\text{scp}_{f_\lambda^d \leq 0}(S_1^{d-1})$ induced by the vertices with f_λ^d -value at most 0. At the vertices e_i and $-e_i$ with $i \leq \lambda$, we have $f_\lambda^d(e_i) = f_\lambda^d(-e_i) = -1 < 0$, and at the remaining vertices e_i and $-e_i$ with $i > \lambda$, the values are $f_\lambda^d(e_i) = f_\lambda^d(-e_i) = 1 > 0$. Therefore $\text{scp}_{f_\lambda^d \leq 0}(S_1^{d-1})$ agrees with the $(\lambda - 1)$ -dimensional unit sphere with respect to the 1-norm in the subspace spanned by the first λ standard base vectors.

Boundary Points. Now we turn to the reference triples for boundary points. The basic proof strategy is the same as for interior points. As link complex for the origin in \mathbb{R}_+^d we choose the part of S_1^{d-1} lying in \mathbb{R}_+^d . It is the subcomplex of S_1^{d-1} induced by all vertices except $-e_d$ and a cone with apex e_d on the intersection of S_1^{d-1} and the subspace $X_d = 0$ as base. Identifying the subspace $X_d = 0$ with \mathbb{R}^{d-1} , the base can be identified with the standard $(d - 2)$ -dimensional unit sphere S_1^{d-2} in that subspace, illustrating that $S_1^{d-1} \cap \mathbb{R}_+^d = e_d S_1^{d-2}$ is a $(d - 1)$ -ball with boundary sphere S_1^{d-2} . Again the assumption that the triples $(\Omega_+^d, 0, f)$ and $(\Omega_+^d, 0, g)$ are locally equivalent for some distinct maps f and g among the standard reference maps, which are π_1^d , f_λ^d , and g_λ^d in this case, would imply by Lemma 3.9 that there is a PL homeomorphism $\phi_{\text{lk}}: e_d S_1^{d-2} \rightarrow e_d S_1^{d-2}$ such that the signs of $f(y)$ and $g(\phi(y))$ agree at each point $y \in e_d S_1^{d-2}$. This homeomorphism necessarily restricts to a PL homeomorphism from the boundary sphere S_1^{d-2} to itself respecting the signs of f and g in the above sense. But applying the characterisation of local equivalence in Lemma 3.9 in the other direction shows that the restricted homeomorphism between the boundaries witnesses that the restrictions of f and g to the subspace \mathbb{R}^{d-1} are locally equivalent at the origin. Obviously, the restriction

of π_1^d to \mathbb{R}^{d-1} is π_1^{d-1} , and the restriction of both f_λ^d and g_λ^d to \mathbb{R}^{d-1} is f_λ^{d-1} . Now the first part of the proof for interior points in \mathbb{R}^{d-1} already excludes the existence of ϕ_{lk} for most of the possible pairs f and g . The only pairs that are not covered yet consist of a map f_λ^d and a map g_λ^d with the same index λ . For such pairs we establish a contradiction by showing that $\text{scp}_{f_\lambda^d \leq 0}(e_d S_1^{d-2})$ is empty, that $\text{scp}_{f_\lambda^d \leq 0}(e_d S_1^{d-2})$ is a $(\lambda - 1)$ -sphere for $\lambda > 0$ and that $\text{scp}_{g_\lambda^d \leq 0}(e_d S_1^{d-2})$ collapses to a point. Since $f_\lambda^d(e_d) = 1 > 0$, we can conclude immediately that $\text{scp}_{f_\lambda^d \leq 0}(e_d S_1^{d-2})$ is the subcomplex $\text{scp}_{f_\lambda^{d-1} \leq 0}(S_1^{d-2})$, for which we have shown in the first part of the proof that it is empty for $\lambda = 0$ and a $(\lambda - 1)$ -sphere for $\lambda > 0$. For g_λ^d , the value at e_d is $g_\lambda^d(e_d) = -1 < 0$, hence $\text{scp}_{g_\lambda^d \leq 0}(e_d S_1^{d-2})$ is the cone with apex e_d and base $\text{scp}_{g_\lambda^{d-1} \leq 0}(S_1^{d-2})$, which collapses to e_d by Fact 3.15. \square

3.2.2 Constructing Link Complexes

Identifying regular and critical points by means of Definition 3.12 becomes more convenient when we have some tools available to construct suitable link complexes of the origin in \mathbb{R}^d or \mathbb{R}_+^d for which a combinatorial equivalence with some link complex of the considered point can be found easily. For meeting the requirements of Lemma 3.7 which assert local equivalence, the standard map in question has to be linear on cells of the cone on the link complex with apex at the origin. This condition is vacuous for the standard map π_1^d since it is globally linear, for the other standard maps it suffices that each cell of the constructed link complex is contained in one of the orthants. The results in the following provide some useful tools for constructing link complexes of the origin.

Lemma 3.18. *If L is a combinatorial $(d - 1)$ -sphere embedded in \mathbb{R}^d such that its domain $|L|$ and the singleton $\{0\}$ consisting of the origin are independent sets in the sense of Definition 1.11, in other words, such that the cone $0L$ with base L and apex 0 is a valid simplicial complex in \mathbb{R}^d , then L is a link complex of the origin in \mathbb{R}^d . For an analogous result in the half-space \mathbb{R}_+^d , we assume that L is a combinatorial $(d - 1)$ -ball embedded in \mathbb{R}_+^d with domain independent of $\{0\}$ such that the boundary sphere of the ball L agrees with the intersection of L with the bounding hyperplane \mathbb{R}^{d-1} of \mathbb{R}^d .*

Proof. Since $0L$ is a cone by assumption, we only have to check that $|0L|$ is a neighbourhood of the origin in \mathbb{R}^d or \mathbb{R}_+^d . We start with the proof for the whole Euclidean space \mathbb{R}^d . By the assumption that L is a combinatorial sphere, there is a PL homeomorphism $\phi: |S_1^{d-1}| \rightarrow |L|$ from the $(d - 1)$ -dimensional 1-norm unit sphere S_1^{d-1} in \mathbb{R}^d to L . Performing the cone construction to ϕ gives a PL homeomorphism $\phi_0: |0S_1^{d-1}| \rightarrow |0L|$ between the two cones with the origin as apex. On this homeomorphism ϕ_0 , we want to apply the invariance of domain theorem [40, p. 33]. It states that the image $\psi(U)$ of a continuous injection $\psi: U \rightarrow \mathbb{R}^d$ defined on an open set $U \subseteq \mathbb{R}^d$ is open in \mathbb{R}^d . Choosing for U the open set $|0S_1^{d-1}| \setminus |S_1^{d-1}|$ and the restriction of ϕ_0 to this subset as continuous injection ψ , the invariance of domain theorem implies that the image $\phi_0(|0S_1^{d-1}| \setminus |S_1^{d-1}|) = |0L| \setminus |L|$, which contains the origin, is an open set in \mathbb{R}^d . Therefore the superset $|0L|$ of this image is a neighbourhood of the origin in \mathbb{R}^d .

The proof for the half-space can be based on the proof for the whole space. Under the given assumptions, L and its mirror image L_- obtained by reflection

at the bounding hyperplane \mathbb{R}^{d-1} together form a combinatorial $(d-1)$ -sphere $L \cup L_-$ such that $\{0\}$ and $|L \cup L_-|$ are independent sets. Then $L \cup L_-$ is a link complex for the origin in \mathbb{R}^d by the first part, and consequently, its intersection with the half-space $L = (L \cup L_-) \cap \mathbb{R}_+^d$ is a link complex for the origin in \mathbb{R}_+^d . \square

Joining Link Complexes. The next result shows how to construct link complexes for the origin from lower dimensional link complexes for the origin in complementary subspaces. We consider \mathbb{R}^d as decomposed into the subspaces \mathbb{R}^k and \mathbb{R}^{d-k} , with \mathbb{R}^k being identified with the subspace of vectors whose last $d-k$ coordinates X_{k+1}, \dots, X_d are zero and with \mathbb{R}^{d-k} being identified with the subspace of vectors whose first k coordinates X_1, \dots, X_k are zero. When considering the origin in \mathbb{R}_+^d , the second subspace \mathbb{R}^{d-k} is replaced by the half-space \mathbb{R}_+^{d-k} identified with the set of vectors whose first k coordinates X_1, \dots, X_k are zero and whose last coordinate X_d is non-negative.

Lemma 3.19. *Let L_1 be a link complex for the origin in \mathbb{R}^k and let L_2 be a link complex for the origin in \mathbb{R}^{d-k} for some k in the range $1 \leq k \leq d-1$. Then the domains of $L_1 \times \{0\}^{d-k}$ and $\{0\}^k \times L_2$ are independent sets in \mathbb{R}^d and the join of the two complexes is a link complex for the origin in \mathbb{R}^d . If L_2 is a link complex for the origin in \mathbb{R}_+^{d-k} instead, then the domains of $L_1 \times \{0\}^{d-k}$ and $\{0\}^k \times L_2$ are independent sets as well and their join is a link complex for the origin in \mathbb{R}_+^d . When we consider the empty set as link complex for the origin in \mathbb{R}^0 , we can also allow the extreme cases $k=0$ and $k=d$, where $L_1 = L_1 \emptyset$ or $L_2 = \emptyset L_2$ are already link complexes for the origin in \mathbb{R}^d or \mathbb{R}_+^d themselves.*

Proof. The cases $k=0$ and $k=d$ are trivial. We first prove that $|L_1 \times \{0\}^{d-k}|$ and $|\{0\}^k \times L_2|$ are independent sets. Assume that $\lambda(\{p_1\} \times \{0\}^{d-k}) + (1-\lambda)(\{0\}^k \times \{p_2\}) = \mu(\{q_1\} \times \{0\}^{d-k}) + (1-\mu)(\{0\}^k \times \{q_2\})$ for some points p_1 and q_1 in $|L_1|$ and p_2 and q_2 in $|L_2|$ with parameters λ and μ in $[0, 1]$. We have to show $\lambda = \mu$ and $p_1 = q_1$ unless $\lambda = \mu = 0$ and $p_2 = q_2$ unless $\lambda = \mu = 1$. Considering the first k and the last $d-k$ coordinates separately, the assumption implies $\lambda p_1 = \mu q_1$ and $(1-\lambda)p_2 = (1-\mu)q_2$. Since a link complex for the origin has the property that any ray emanating from the origin intersects the link at exactly one point, we can conclude that the equality $\lambda p_1 = \mu q_1$ for p_1 and q_1 in $|L_1|$ implies $\lambda = \mu$ and $p_1 = q_1$ unless $\lambda = \mu = 0$, and that $(1-\lambda)p_2 = (1-\mu)q_2$ for p_2 and q_2 in $|L_2|$ implies $\lambda = \mu$ and $p_2 = q_2$ unless $\lambda = \mu = 1$ as desired.

Let us denote the join $(L_1 \times \{0\}^{d-k})(\{0\}^k \times L_2)$ by L . For verifying that $0L$ is a valid cone in \mathbb{R}^d , we have to check that any ray emanating from the origin intersects $|L|$ at most once. Consider a ray generated by the non-zero vector (x_1, \dots, x_d) . Assuming $\mu(x_1, \dots, x_d) = p$ for some $\mu \geq 0$ and a point $p = \lambda(\{p_1\} \times \{0\}^{d-k}) + (1-\lambda)(\{0\}^k \times \{p_2\})$ with λ , p_1 , and p_2 as above implies $\mu(x_1, \dots, x_k) = \lambda p_1$ and $\mu(x_{k+1}, \dots, x_d) = (1-\lambda)p_2$. The case $\mu = 0$ is impossible since at least one of λp_1 or $(1-\lambda)p_2$ is non-zero. The case $\lambda = 0$ happens if and only if $x_1 = \dots = x_k = 0$ holds, and then μ and p_2 are uniquely determined by the equality $\mu(x_{k+1}, \dots, x_d) = p_2$ because L_2 is a link complex for the origin in \mathbb{R}^{d-k} . Symmetrically, $\lambda = 1$ happens if and only if $x_{k+1} = \dots = x_d = 0$ holds, and then μ and p_1 are determined by $\mu(x_1, \dots, x_k) = p_1$. Otherwise the equalities $(\mu/\lambda)(x_1, \dots, x_k) = p_1$ and $(\mu/(1-\lambda))(x_{k+1}, \dots, x_d) = p_2$ determine uniquely p_1 , p_2 , and the quotients μ/λ and $\mu/(1-\lambda)$. Since both quotients are positive, their reciprocals λ/μ and

$(1 - \lambda)/\mu$, their sum $\lambda/\mu + (1 - \lambda)/\mu = 1/\mu$, and the reciprocal μ of the sum are also positive and uniquely determined. In particular, the uniqueness of μ shows that the ray in direction (x_1, \dots, x_d) intersects $|L|$ in at most one point.

For completing the proof, we can either observe that the equalities above do not only reduce the uniqueness problem of the solutions μ , λ , p_1 , and p_2 to the complementary subspaces but also their existence problem. Therefore solutions for the equalities above exist and can be computed from the solutions in the complementary subspaces for any non-zero vector (x_1, \dots, x_d) in \mathbb{R}^d (with $x_d \geq 0$ in the case that L_2 is a link complex for \mathbb{R}_+^{d-k}), which shows that $|L|$ is a link for the origin in \mathbb{R}^d or \mathbb{R}_+^d respectively. Or we can apply Lemma 3.18: We just proved that $|L|$ and $\{0\}$ are independent sets. Furthermore L_1 is a $(k - 1)$ -sphere and if L_2 is a link in \mathbb{R}^d , it is a $(d - k - 1)$ -sphere, or if L_2 is a link in \mathbb{R}_+^d , it is a $(d - k - 1)$ -ball. By Fact 1.38 the join $L = (L_1 \times \{0\}^{d-k})(\{0\}^k \times L_2)$ is a $(d - 1)$ -sphere in the former case, and a $(d - 1)$ -ball in the latter case. Moreover in the latter case, the boundary of the $(d - 1)$ -ball is the join of $L_1 \times \{0\}^{d-k}$ and $\{0\}^k \times \text{bd}(L_2)$, and since L_2 is a link in \mathbb{R}_+^{d-k} its boundary coincides with its intersection with the hyperplane defined by the last coordinate being zero. Hence $L_1 \times \{0\}^{d-k}$ and $\{0\}^k \times \text{bd}(L_2)$ both lie in the bounding hyperplane \mathbb{R}^{d-1} and their join is at the same time the boundary sphere of L and its intersection with the bounding hyperplane \mathbb{R}^{d-1} . Therefore all requirements of Lemma 3.18 are fulfilled implying that L is a link complex for the origin. \square

Perturbing Link Complexes. For some of the following proofs, we need a result stating that a slight perturbation of the positions of the vertices of a simplicial complex embedded in \mathbb{R}^n preserves the property of being a simplicial complex and even its combinatorial structure, provided the perturbation is small enough. More precisely, for any subset of vertices spanning a simplex of the original complex, the perturbed set of vertices still spans a simplex and the resulting family of perturbed simplices still forms a simplicial complex combinatorially equivalent to the original one with the combinatorial equivalence assigning each vertex to its perturbed counterpart. The following Fact formalises such a claim and includes a sketch for a proof.

Fact 3.20. Let K be a simplicial complex embedded in \mathbb{R}^n . An ϵ -perturbation of the positions of the vertices is a map $P_0: \text{Vrt}(K) \rightarrow \mathbb{R}^n$ such that for each vertex $v \in \text{Vrt}(K)$ the distance between v and its perturbed position $P_0(v)$ is at most ϵ . Such an ϵ -perturbation induces a map P on K , that assigns to each simplex $S \in K$ a convex polytope $P(S) = \text{conv}\{P_0(v) \mid v \in \text{Vrt}(S)\}$ which is the convex hull of the perturbed vertices of S . The image of P , that is the collection $\{P(S) \mid S \in K\}$ of all these convex polytopes is denoted by $P(K)$ and called an ϵ -perturbation of K . There is an $\epsilon > 0$, depending on K , such that for any ϵ -perturbation P_0 of the positions of the vertices of K , the induced ϵ -perturbation $P(K)$ of K is a simplicial complex combinatorially equivalent to K , with $P: K \rightarrow P(K)$ being the combinatorial equivalence in terms of cells and $P_0: \text{Vrt}(K) \rightarrow \text{Vrt}(P(K))$ being the combinatorial equivalence in terms of vertices.

A proof for this claim can be based on a continuity argument: Let the sequence (v_1, \dots, v_N) be an enumeration of the vertices of K . It is possible to find a finite family of continuous functions $f_j: (\mathbb{R}^n)^N \rightarrow \mathbb{R}$ in terms of the coordinates of an N -tuple of points in \mathbb{R}^n , such that $f_j(v_1, \dots, v_N) > 0$ holds for each

j and whenever $f_j(u_1, \dots, u_N) > 0$ holds for each j , the N -tuple (u_1, \dots, u_N) of points $u_i \in \mathbb{R}^n$ generate a simplicial complex combinatorially equivalent to K , with $\phi^{\text{Vrt}}(v_i) = u_i$ defining a combinatorial equivalence in terms of vertices. The continuity of the f_j ensures that there is an ϵ such that if each u_i has distance at most ϵ from v_i , then $f_j(u_1, \dots, u_N) > 0$ still holds for each j . The maps f_j are chosen in such a way, that they witness by the computation of suitable determinants, that certain subsets of the points are affinely independent or that certain points lie on different sides of some hyperplane containing some other points. Such tests are sufficient to witness that a certain collection of convex hulls of subsets of the points fulfils the defining properties of a simplicial complex: Affine independence of a subset of points witnesses that they span a simplex. The fact that two simplices intersect only in a common face can be witnessed by a separating hyperplane containing the common vertices of the simplices, such that the non-common vertices of one simplex and the non-common vertices of the other simplex lie strictly on opposite sides of the separating hyperplane.

We illustrate the construction of the maps f_j by some schematic examples. The continuity of the constructed maps will follow immediately from the fact that the map assigning to an n -tuple of vectors its determinant is continuous. Let v_{i_0}, \dots, v_{i_k} span a k -simplex of K . This means that the points are affinely independent and consequently that the vectors $v_{i_1} - v_{i_0}, \dots, v_{i_k} - v_{i_0}$ are linearly independent. If necessary, namely if $k < n$, we can extend this family of vectors to a base of \mathbb{R}^n by some suitable vectors w_{k+1}, \dots, w_n . The fact that these vectors form a base is then witnessed by their determinant being non-zero. If the determinant is positive, we add the map that computes the determinant of the vectors $v_{i_1} - v_{i_0}, \dots, v_{i_k} - v_{i_0}, w_{k+1}, \dots, w_n$ as a function of the coordinates of the points v_{i_0}, \dots, v_{i_k} to the family of maps f_j we want to construct, otherwise we add the negative of this map to the family. When we compute the value of the map added to the family with the points v_{i_0}, \dots, v_{i_k} replaced by some points u_{i_0}, \dots, u_{i_k} and the obtained value is still positive, this fact witnesses that the points u_{i_0}, \dots, u_{i_k} are affinely independent and span a k -simplex.

Now assume that $v_{i_{-k'}}, \dots, v_{i_{-1}}, v_{i_0}, \dots, v_{i_k}$ spans a simplex of K , and that $v_{i_0}, \dots, v_{i_k}, v_{i_{k+1}}, \dots, v_{i_{k''}}$ spans another simplex of K , with v_{i_0}, \dots, v_{i_k} being exactly the common vertices of the two simplices. The fact that the two simplices intersect only in their common face spanned by v_{i_0}, \dots, v_{i_k} is witnessed by the existence of a hyperplane H containing v_{i_0}, \dots, v_{i_k} such that the points $v_{i_{-k'}}, \dots, v_{i_{-1}}$ lie strictly on one side, and the points $v_{i_{k+1}}, \dots, v_{i_{k''}}$ lie strictly on the other side. This property can be checked by computing $k' + k'' - k$ determinants of certain vectors as follows: The first $n - 1$ vectors span the underlying vector space of H and the last vector points from the hyperplane to one of the other vertices $v_{i_{-k'}}, \dots, v_{i_{-1}}$ or $v_{i_{k+1}}, \dots, v_{i_{k''}}$. To determine these vectors we choose points $p_0 = v_{i_0}, \dots, p_k = v_{i_k}$ and possibly some additional points p_{k+1}, \dots, p_{n-1} in H such that p_0, \dots, p_{n-1} are affinely independent and their affine hull is the separating hyperplane H . Note that $k = -1$ is possible when the simplices are disjoint. Then we take the vectors $w_i = p_i - p_0$ for $1 \leq i \leq n - 1$ and compute for each choice $w_n = v_{i_{-k'}} - p_0, \dots, w_n = v_{i_{-1}} - p_0, w_n = v_{i_{k+1}} - p_0, \dots, w_n = v_{i_{k''}} - p_0$ the determinant of the vectors w_1, \dots, w_n . The results for the former choices with $w_n = v_{i_{-k'}} - p_0, \dots, w_n = v_{i_{-1}} - p_0$ have the same sign which is opposite to the common sign of the results for the latter choices with $w_n = v_{i_{k+1}} - p_0, \dots, w_n = v_{i_{k''}} - p_0$. We can assume without loss

of generality that the former choices yield negative results and the latter choices yield positive results. Then we add maps computing the negative determinant of the vectors w_1, \dots, w_n for the choices $w_n = v_{i_{-k'}} - p_0, \dots, w_n = v_{i_{-1}} - p_0$ and maps computing the determinant of the vectors w_1, \dots, w_n for the choices $w_n = v_{i_{k+1}} - p_0, \dots, w_n = v_{i_{k''}} - p_0$ to the family of maps f_j , expressing each of these maps as a function of the coordinates of the points $v_{i_{-k'}}, \dots, v_{i_{k''}}$. When we compute the values of these maps with the points $v_{i_{-k'}}, \dots, v_{i_{k''}}$ replaced by some points $u_{i_{-k'}}, \dots, u_{i_{k''}}$ and the obtained values are all positive, we can conclude that the convex hulls of the points $u_{i_{-k'}}, \dots, u_{i_{-1}}, u_{i_0}, \dots, u_{i_k}$ on the one hand and the points $u_{i_0}, \dots, u_{i_k}, u_{i_{k+1}}, \dots, u_{i_{k''}}$ on the other hand intersect only in a common face spanned by u_{i_0}, \dots, u_{i_k} , because the computed determinants witness that the points $u_{i_{-k'}}, \dots, u_{i_{-1}}$ lie on one side, and the points $u_{i_{k+1}}, \dots, u_{i_{k''}}$ on the other side of a hyperplane spanned by the points $u_{i_0}, \dots, u_{i_k}, p_{k+1}, \dots, p_{n-1}$.

As a first application of this fact, we show that for sufficiently small ϵ , an ϵ -perturbation of a link complex of the origin in \mathbb{R}^d is still a link complex of the origin in \mathbb{R}^d . For an analogous conclusion concerning link complexes in \mathbb{R}_+^d , we have to assume that the perturbation moves the points on the boundary of the link complex only within the bounding hyperplane \mathbb{R}^{d-1} .

Lemma 3.21. *Let L be a link complex for the origin in \mathbb{R}^d . There exists an $\epsilon > 0$, such that any ϵ -perturbation $P(L)$ of L is a link complex for the origin in \mathbb{R}^d as well and combinatorially equivalent to L .*

For any link complex L of the origin in \mathbb{R}_+^d exists an ϵ such that for any ϵ -perturbation P_0 of the positions of the vertices of L that maps points in the bounding hyperplane \mathbb{R}^{d-1} to points lying again in \mathbb{R}^{d-1} , the induced perturbation $P(L)$ is a link complex of the origin in \mathbb{R}_+^d combinatorially equivalent to L .

Proof. For a link complex L of the origin in \mathbb{R}^d , any ϵ -perturbation $P(L)$ for sufficiently small ϵ is combinatorially equivalent to L by Fact 3.20 and hence a $(d-1)$ -sphere. Extending the perturbation of L to a perturbation of $0L$ by keeping the vertex at the origin fixed and applying again Fact 3.20 shows that $0(P(L))$ is a valid cone complex in \mathbb{R}^d . This asserts that the requirements of Lemma 3.18 are fulfilled and $P(L)$ is a link complex for the origin in \mathbb{R}^d .

For a link complex L of the origin in \mathbb{R}_+^d the analogous applications of Lemma 3.18 show that $P(L)$ is a $(d-1)$ -ball combinatorially equivalent to L and $0(P(L))$ is a valid cone in \mathbb{R}_+^d . We can assume that ϵ is so small that no interior vertex of L is moved to the bounding hyperplane \mathbb{R}^{d-1} . Together with the assumption that boundary vertices of L are moved within \mathbb{R}^{d-1} , this implies the requirement from Lemma 3.18 for links in \mathbb{R}_+^d that the boundary of $P(L)$ is exactly the intersection of $P(L)$ with \mathbb{R}^{d-1} . Therefore Lemma 3.18 yields that $P(L)$ is a link complex for the origin in \mathbb{R}_+^d as desired. \square

3.2.3 Restricting Criticality to Vertices

Now we turn to the claim that critical points of a map in general position on a combinatorial manifold can only occur at vertices. We show that any point lying in the interior of a simplex of a combinatorial manifold where the considered map linear on cells is non-constant on that simplex is a regular point. This

implies the claim because vertices are the only simplices on which a map in general position is constant.

Lemma 3.22. *Assume that M is a combinatorial d -manifold $f: |M| \rightarrow \mathbb{R}$ is a map linear on cells of M . If a point q lies in the interior of a cell where f is non-constant, then q is a regular point of f . This claim holds likewise for interior and boundary points, with only minor differences in the proof; the necessary changes for boundary points are indicated by parentheses.*

Proof. Let Q be the k -cell containing q in its interior and f_Q the non-constant linear map on $\text{aff}(Q)$ agreeing with f on Q . Recall from Table 1.3 that $\text{lk}(q)$ is the join of the boundary $\text{bd}(Q)$ and the link $\text{lk}(Q)$ of Q . The boundary $\text{bd}(Q)$ is a $(k-1)$ -sphere formed by the boundary complex of a k -simplex. Since M is a manifold, the link $\text{lk}(Q)$ is a $(d-k-1)$ -sphere (-disk) by Fact 1.37 and $\text{lk}(q)$ is a $(d-1)$ -sphere (-ball). We show how to construct a cone neighbourhood of the origin that is combinatorially equivalent to a subdivision of $\text{lk}(q)$ such that $f - f(q)$ and π_1^d have the same sign at corresponding vertices. Since π_1^d is linear on $\mathbb{R}_{(+)}^d$ as a whole, this implies by Lemma 3.7 that (M, q, f) and $(\mathbb{R}_{(+)}^d, 0, \pi_1^d)$ are locally equivalent and consequently q is a regular point.

Map $\text{aff}(Q)$ by a suitable bijective affine transformation in such a way to the k -dimensional subspace \mathbb{R}^k of \mathbb{R}^d where all coordinates X_{k+1} up to X_d are zero, that q is mapped to the origin and the half-space $\text{aff}(Q)_{f_Q \leq f(q)}$ is mapped to the half-space $X_1 \leq 0$ of \mathbb{R}^k . The image of the simplex Q under this map is a k -simplex B in \mathbb{R}^k that contains the origin in its relative interior. Hence $\text{bd}(B)$ is a $(k-1)$ -sphere, $\{0\}$ and $|\text{bd}(B)|$ are independent sets because $0 \in \text{bd}(B)$ can be obtained by an elementary starring of B at the origin, and $\text{bd}(B)$ is a link complex for the origin in \mathbb{R}^k by Lemma 3.18. Furthermore the signs of $f - f(q)$ at the vertices of Q and the signs of π_1^d at the corresponding vertices of B agree. Therefore the affine transformation induces a combinatorial equivalence between $\text{bd}(Q)$ and $\text{bd}(B)$ such that $f - f(q)$ and π_1^d have the same sign at corresponding vertices.

Consider the 1-norm unit $(d-k-1)$ -(half)-sphere in the $(d-k)$ -dimensional subspace (sub-half-space) where all coordinates X_1 up to X_k are zero (and $X_d \geq 0$). Let S be a subdivision of this (half-)sphere that is combinatorially equivalent to a subdivision S' of $\text{lk}(Q)$. Clearly S is a link complex for the origin in the subspace (sub-half-space). By Lemma 3.19, the join of S and $\text{bd}(B)$ is a link complex for the origin. It is combinatorially equivalent to the subdivision $S' \text{bd}(Q)$ of $\text{lk}(q) = \text{lk}(Q) \text{bd}(Q)$. The only property we have not fulfilled yet are the matching signs of $f - f(q)$ and π_1^d at corresponding vertices of S' and S . To achieve this, we choose, according to Lemma 3.21, a sufficiently small $\epsilon > 0$ such that any ϵ -perturbation of $S \text{bd}(B)$ (leaving vertices in \mathbb{R}^{d-1} within \mathbb{R}^{d-1}) is again a link complex for the origin combinatorially equivalent to $S \text{bd}(B)$ and apply the lemma to a specific ϵ -perturbation P_0 for the positions of the vertices of $S \text{bd}(B)$ defined as follows: We shift the vertices of S that correspond to vertices of S' with positive/negative $f - f(q)$ -value slightly in positive/negative X_1 -direction by an ϵ amount, the other vertices remain unchanged. The resulting perturbed complex $P(S \text{bd} B)$ is still a link complex for the origin in $\mathbb{R}_{(+)}^d$ combinatorially equivalent to $S' \text{bd}(Q)$ and π_1^d and $f - f(q)$ have the same sign at corresponding vertices by construction. \square

The fact that points in the interior of a cell where f is non-constant are

regular provides another proof for the result from Corollary 2.40 that $|M|_{=h}$ is a $(d-1)$ -manifold for any value h not attained at vertices of M as follows: Let $x \in |M|_{=h}$. Then x lies in the interior of a cell where f is non-constant, because otherwise the value $h = f(x)$ would be attained at the vertices of the cell containing x in its interior. By Lemma 3.22, the point x is regular, and the characterisation of local equivalence via PL-homeomorphic neighbourhoods such that the shifted maps agree at corresponding points yields for (M, x, f) and $(\mathbb{R}^d, 0, \pi_1^d)$ a local coordinate system centred at x such that the shifted map $f - h$ is represented by π_1^d in these coordinates. Therefore $|M|_{=h}$ in some neighbourhood of x corresponds to the subspace $\pi_1^d = 0$ in some neighbourhood of the origin in the local coordinates. This shows that in a neighbourhood of x , the level set $|M|_{=h}$ is a $(d-1)$ -sub-manifold embedded in $|M|$. Since this applies to any point $x \in |M|_{=h}$, the level set $|M|_{=h}$ as a whole is a $(d-1)$ -sub-manifold embedded in $|M|$.

3.2.4 Existence of PL Morse Functions and Their Construction from Discrete Morse Functions

For the proof of the next result that any combinatorial manifold has a PL Morse function on its derived, we need a simple observation concerning the links of vertices in a derived subdivision.

Lemma 3.23. *Consider a vertex v_S in a derived subdivision $\text{sdv}^{\text{drv}}(M)$ of a simplicial complex M which is the vertex chosen in the interior of a cell S of M to construct the derived. Then the link $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ in the derived is combinatorially equivalent to the join of the derived $\text{sdv}^{\text{drv}}(\text{lk}_M(S))$ of the link of S in M and the derived $\text{sdv}^{\text{drv}}(\text{bd}(S))$ of the boundary of S . The combinatorial equivalence in terms of vertices is given by the following assignment: A vertex v_T in $\text{sdv}^{\text{drv}}(\text{bd}(S))$ with $T \in \text{bd } S$ is mapped to v_T in $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ and a vertex v_U from $\text{sdv}^{\text{drv}}(\text{lk}_M(S))$ with $U \in \text{lk}_M(S)$ is mapped to the vertex v_{SU} from $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$.*

Proof. Recall from Definition 1.23 that a simplex in a derived subdivision of M is spanned by vertices v_{S_0}, \dots, v_{S_k} such that the cells S_0, \dots, S_k form a chain in the face poset of M . Therefore a simplex in $\text{sdv}^{\text{drv}}(M)$ containing the vertex v_S is spanned by v_S , some vertices v_{T_1}, \dots, v_{T_m} , and some vertices v_{W_1}, \dots, v_{W_n} with $T_1 \sqsubset_M \dots \sqsubset_M T_m \sqsubset_M S \sqsubset_M W_1 \sqsubset_M \dots \sqsubset_M W_n$. Consequently, simplices in $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ are the join of a simplex spanned by some of the vertices v_{T_1}, \dots, v_{T_m} as above and a simplex spanned by some of the vertices v_{W_1}, \dots, v_{W_n} as above. Since we assume $T_1 \sqsubset_M \dots \sqsubset_M T_m \sqsubset_M S$, the cells T_i lie in the boundary of S and a simplex spanned by vertices v_{T_1}, \dots, v_{T_m} is a simplex in the derived $\text{sdv}^{\text{drv}}(\text{bd}(S))$ of the boundary of S . On the other hand, the co-faces W_j of S can be represented as a join of S and some simplices $U_j \in \text{lk}_M(S)$ with $U_1 \sqsubset_M \dots \sqsubset_M U_n$. The corresponding vertices v_{U_1}, \dots, v_{U_n} span a simplex in the derived $\text{sdv}^{\text{drv}}(\text{lk}_M(S))$ of the link of S . With these observations, we can easily check that the vertices $v_{T_1}, \dots, v_{T_m}, v_{SU_1}, \dots, v_{SU_n}$ span a simplex in $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ if and only if the vertices v_{T_1}, \dots, v_{T_m} span a simplex in $\text{sdv}^{\text{drv}}(\text{bd}(S))$ and the vertices v_{U_1}, \dots, v_{U_n} span a simplex in $\text{sdv}^{\text{drv}}(\text{lk}_M(S))$. This proves that the assignments of v_T to itself for $T \in \text{bd}(S)$ and of v_U to v_{SU} for $U \in \text{lk}_M(S)$ define a combinatorial equivalence in terms of

vertices between the join of $\text{sdv}^{\text{drv}}(\text{lk}_M(S))$ and $\text{sdv}^{\text{drv}}(\text{bd}(S))$. Note that the fact $\text{lk}_M(v_S) = \text{bd}(S) \text{lk}_M(S)$ from Table 1.3 asserts that $|\text{sdv}^{\text{drv}}(\text{lk}_M(S))| = |\text{lk}_M(S)|$ and $|\text{sdv}^{\text{drv}}(\text{bd}(S))| = |\text{bd}(S)|$ are independent sets in the ambient space of M . \square

Existence Proof. The statement of the following theorem, that any combinatorially manifold has a PL Morse function, was also observed by Eells and Kuiper in [17], where PL Morse functions are called C^{omb} -non-degenerate functions. Their proof is based on the same idea as the proof given below.

Theorem 3.24. *For any derived subdivision $\text{sdv}^{\text{drv}}(M)$ of a combinatorial d -manifold M , there is a PL Morse function on $|M|$ linear on cells of $\text{sdv}^{\text{drv}}(M)$.*

Proof. A derived subdivision of M contains for any cell $S \in M$ a vertex v_S in the interior of S . For a map $f_0: \text{Vrt}(\text{sdv}^{\text{drv}}(M)) \rightarrow \mathbb{R}$, let f be the map linear on cells of $\text{sdv}^{\text{drv}}(M)$ obtained by simplex-wise linear interpolation between the values of f_0 . We show that any such map f obtained from a map f_0 that fulfils $f_0(v_T) < f_0(v_S)$ whenever T is a proper face of S in M is a PL Morse function. A possible choice for f_0 fulfilling the requirement is $f_0(v_S) = \dim(S)$. This is also the map used in the existence proof of Eells and Kuiper in [17]. When a map in general position is desired, we can modify this example by adding small individual values ϵ_S at each vertex v_S with all ϵ_S being distinct positive numbers less than 1, yielding distinct values $f_0(v_S) = \dim(S) + \epsilon_S$. So let f be the simplex-wise linear interpolation of a map f_0 with $f_0(v_T) < f_0(v_S)$ whenever T is a proper face of S . We claim that all non-vertices are regular and all vertices v_S are non-degenerate critical with index $\lambda = \dim(S)$.

Non-Vertices. If $x \in |\text{sdv}^{\text{drv}}(M)|$ is not a vertex, then the cell that contains x in its interior is a simplex formed by vertices v_{S_i} where the cells S_i when indexed in suitable order form a strictly increasing sequence in the face poset of M . Consequently, the values $f(v_{S_i})$ arranged in the same order form a strictly increasing sequence of real numbers and are distinct. Lemma 3.22 shows that x is regular.

Proof Idea for Vertices. For a vertex v_S , denote by ϕ^{comb} the combinatorial equivalence in terms of cells $\phi^{\text{comb}}: \text{sdv}^{\text{drv}}(\text{lk}_M(S)) \text{sdv}^{\text{drv}}(\text{bd}(S)) \rightarrow \text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ between the link in the derived and the join of the derived boundary of S and the derived link of S in M obtained from Lemma 3.23. Using the combinatorial equivalence, which is the identity on the common subcomplex $\text{sdv}^{\text{drv}}(\text{bd}(S))$, we can describe the link $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ as the join of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ and $\text{sdv}^{\text{drv}}(\text{bd}(S))$.

The proof idea is the following, with the modifications for boundary cells S added in parentheses: The boundary of S is a $(\lambda - 1)$ -sphere for $\lambda = \dim(S)$, and for cells $T \in \text{bd } S$, the vertices v_T in $\text{sdv}^{\text{drv}}(\text{bd}(S))$ have a value $f(v_T) < f(v_S)$; since M is a combinatorial manifold, the link of S in M is a $(d - \lambda - 1)$ -sphere (-ball), and for $U \in \text{lk}_M(S)$, the vertices v_{SU} of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ have a value $f(v_{SU}) > f(v_S)$. On the other hand, the reference map f_λ^d is negative, even constantly -1 , on the $(\lambda - 1)$ -dimensional 1-norm unit sphere $S_1^{\lambda-1}$ around the origin in the λ -dimensional subspace with coordinates $X_i = 0$ for $\lambda + 1 \leq i \leq d$; and it is positive, even constantly 1 , on the $(d - \lambda - 1)$ -dimensional 1-norm unit (half-)sphere $S_1^{d-\lambda-1}$ around the origin in the $(d - \lambda)$ -dimensional subspace (sub-half-space) with coordinates $X_i = 0$ for $1 \leq i \leq \lambda$ (and $X_d \geq 0$). (Note that the

case $\lambda = \dim(S) = d$ cannot occur if S is a boundary simplex of a combinatorial d -manifold.) Since the $(d-1)$ -dimensional 1-norm unit (half-)sphere S_1^{d-1} around the origin in \mathbb{R}^d is the join of $S_1^{\lambda-1}$ and $S_1^{d-\lambda-1}$, it suffices to choose combinatorially equivalent subdivisions with matching signs at corresponding vertices for $\text{sdv}^{\text{drv}}(\text{bd}(S))$ and $S_1^{\lambda-1}$ on the one hand and for $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ and $S_1^{d-\lambda-1}$ on the other hand. The respective joins of these subdivisions yield combinatorially equivalent subdivisions with matching signs at corresponding vertices of the joins $\text{sdv}^{\text{drv}}(\text{bd}(S))\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S))) = \text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ and $S_1^{\lambda-1}S_1^{d-\lambda-1} = S_1^{d-1}$.

Decomposing the Construction. As outlined above, our aim is to construct a link complex for the origin, namely a subdivision of the 1-norm unit (half-)sphere S_1^{d-1} that is combinatorially equivalent to some subdivision of $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ in such a way that f_λ^d and $f - f(v_S)$ have the same sign at corresponding vertices. Since each cell of S_1^{d-1} is contained in some cell of $\Omega_{(+)}^d$, in other words, in one of the orthants, the reference map f_λ^d is linear on cells of the cone $0L$ for any subdivision L of S_1^{d-1} . Once the link complex is constructed, Lemma 3.7 asserts that (M, v_S, f) and $(\Omega_{(+)}^d, 0, f_\lambda^d)$ are locally equivalent and hence v_S is a non-degenerate critical point of f of index λ by definition.

The construction of the combinatorially equivalent subdivisions of S_1^{d-1} and $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ can be decomposed into two independent constructions of pairs of combinatorially equivalent subdivisions by representing both complexes as joins: The 1-norm unit (half-)sphere S_1^{d-1} is the join of the unit sphere $S_1^{\lambda-1}$ and the unit (half-)sphere $S_1^{d-\lambda-1}$ in orthogonal subspaces as described above. The link $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ is the join of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ and $\text{sdv}^{\text{drv}}(\text{bd}(S))$. Now it suffices to construct combinatorially equivalent subdivisions of $S_1^{\lambda-1}$ and $\text{sdv}^{\text{drv}}(\text{bd}(S))$ on the one hand and of $S_1^{d-\lambda-1}$ and $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ on the other hand such that both equivalences respect the sign at corresponding vertices because the joins of the subdivisions yield subdivisions of the joins with the desired properties: Clearly, for any complexes K and L with subdivisions K' and L' , the join $K'L'$ of the subdivisions is a subdivision of the join KL , because if $A' \in K'$ is contained in $A \in K$ and $B' \in L'$ is contained in $B \in L$ then $A'B'$ is contained in AB . We noted already in Fact 1.28 that a combinatorial equivalence of K and K' and of L and L' implies a combinatorial equivalence of the joins KL and $K'L'$. The vertex set of the join is the union of the original vertex sets, and the combinatorial equivalence in terms of vertices for the joins arises from the union of the two combinatorial equivalences. Therefore the joined combinatorial equivalence respects the signs at corresponding vertices if and only if the original combinatorial equivalences do so.

The Boundary Part. For proving the existence of combinatorially equivalent subdivisions of $S_1^{\lambda-1}$ and $\text{sdv}^{\text{drv}}(\text{bd}(S))$ with matching signs of f_λ^d and $f - f(v_S)$ at corresponding vertices, we show first that both maps have negative values throughout the complexes, which ensures matching signs at corresponding vertices for any combinatorially equivalent subdivisions, and afterwards the existence of the combinatorial equivalence. Since $S_1^{\lambda-1}$ is considered as embedded in the λ -dimensional subspace with coordinates $X_i = 0$ for $\lambda+1 \leq i \leq d$ and since the map f_λ^d restricted to this subspace is the negative of the 1-norm by definition, the map f_λ^d is constantly -1 on the 1-norm unit sphere $|S_1^{\lambda-1}|$. A cell T in the boundary of S fulfils $T \sqsubset S$, hence $f(v_T) < f(v_S)$ holds for vertices v_T in $\text{sdv}^{\text{drv}}(\text{bd}(S))$. Therefore $f - f(v_S)$ is negative at all vertices of $\text{sdv}^{\text{drv}}(\text{bd}(S))$.

and consequently negative on the whole complex since the values at the other points are obtained by linear interpolation between the negative values at the vertices. The boundary of the λ -simplex S is a combinatorial $(\lambda - 1)$ -sphere and so is the derived of the boundary. The unit sphere $S_1^{\lambda-1}$ is a combinatorial $(\lambda - 1)$ -sphere as well, hence $S_1^{\lambda-1}$ and $\text{sdv}^{\text{drv}}(\text{bd}(S))$ have combinatorially equivalent subdivisions because they are PL-homeomorphic combinatorial manifolds.

The Link Part. The construction of combinatorially equivalent subdivisions for the complex $S_1^{d-\lambda-1}$ and $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ works similarly. The (half-)sphere $S_1^{d-\lambda-1}$ is embedded in the $(d - \lambda)$ -dimensional subspace (sub-half-space) with coordinates $X_i = 0$ for $1 \leq i \leq \lambda$ (and $X_d \geq 0$), where the map f_λ^d agrees with the 1-norm and consequently attains the constant value 1 on the unit (half-)sphere $|S_1^{d-\lambda-1}|$. For a cell U in $\text{lk}_M(S)$, the join SU is a proper co-face of S , hence $f(v_{SU}) > f(v_S)$ holds for vertices v_{SU} in $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$. Therefore f_λ^d is positive on $|S_1^{d-\lambda-1}|$ and $f - f(v_S)$ is positive on $|\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))|$. The link $\text{lk}_M(S)$ is a $(d - \lambda - 1)$ -dimensional combinatorial sphere (ball) because M is a combinatorial manifold, and this remains true for the derived subdivision $\text{sdv}^{\text{drv}}(\text{lk}_M(S))$ and its counterpart $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ under the combinatorial equivalence ϕ^{comb} . The complex $S_1^{d-\lambda-1}$ is a combinatorial $(d - \lambda - 1)$ -sphere (-ball) as well, so there are combinatorially equivalent subdivisions of $S_1^{d-\lambda-1}$ and $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ with positive signs for f_λ^d and $f - f(v_S)$ at all vertices. This completes the proof. \square

Constructing PL Morse Functions from Discrete Morse Functions.

For any of the PL Morse functions f considered in the proof of the theorem, every vertex is critical with respect to f . This property is reminiscent of the existence proof for discrete Morse functions by observing that the map $g(S) = \dim(S)$ assigning to each cell its dimension is a discrete Morse function for which each cell is critical [19, p. 108]. Moreover, if every cell is critical with respect to a discrete Morse function g , then g fulfils $g(T) < g(S)$ whenever T is a proper face of S , which corresponds exactly to our assumption used in the proof on the values of the map f at the vertices of the derived, which implied that f is a PL Morse function where every vertex is critical, namely that $f(v_T) < f(v_S)$ holds whenever T is a proper face of S . This observation suggests to conjecture a more general result that allows to transform a discrete Morse function g on a combinatorial manifold M into a PL Morse function f on the derived of M by interpolating simplex-wise linearly between the values $f(v_S) = g(S)$ at each vertex v_S of the derived corresponding to a cell S of M .

A simple counterexample that refutes this conjecture in its most general form is the combinatorial 1-manifold formed by an edge e with vertices u and w together with the discrete Morse function g given by the values $g(u) = 0$ and $g(w) = g(e) = 1$. When f is defined as proposed, we obtain an edge between v_e and w in the derived of e where f is constantly 1, so that all points on that edge are degenerate critical with respect to f . But if the discrete Morse function g fulfils a general position assumption such as the one suggested in Section 1.3, namely that all values are distinct and that matched pairs in the Morse matching are the only exception to monotonicity, then the induced map f defined as above turns out to be a PL Morse function. Furthermore the critical

points of f are exactly the vertices v_S for which the cell S is critical with respect to g or a boundary cell of M which is critical with respect to the restriction of g to $\text{bd}(M)$. The index of such a critical vertex v_S is the dimension of S . Since there is for any discrete Morse function g an equivalent discrete Morse function g' in general position, we can find for any discrete Morse function g on a combinatorial manifold M a PL Morse function f linear on cells of $\text{sdv}^{\text{drv}}(M)$ such that the critical points of f are exactly the vertices v_S with S being a critical cell of g or of the restriction of g to the boundary of M .

Before we prove the conjecture under the additional assumption on g , let us shortly review two earlier related results. For the first result due to Bloch [8], consider a finite regular CW-complex M with discrete Morse function g and a Euclidean space \mathbb{R}^n of sufficiently large dimension n with a non-constant linear map $f: \mathbb{R}^n \rightarrow \mathbb{R}$, which can be viewed as map linear on cells for any complex embedded in \mathbb{R}^n . (In [8], the map f is described as the projection onto a line spanned by a unit vector in the space.) The result states that the barycentric subdivision of M can be embedded in \mathbb{R}^n in such a way that a cell S in M is critical with respect to g if and only if the barycentre v_S of S is critical with respect to f , however Bloch uses for v_S and f a notion of criticality that is due to Banchoff [4, 5]. This notion differs from the one introduced here in Definition 3.12 and will be reviewed later in this section. The other result deals with the map f linear on cells of a derived obtained from a discrete Morse function g on a simplicial complex K by simplex-wise linear interpolation of the values $f(v_S) = g(S)$. Bauer [6, Theorem 2.27, p. 32] shows that for every level $h \in \mathbb{R}$, the level subcomplex $K(h)$ generated by the cells S of K with $g(S) \leq h$ and the lower level set $|\text{sdv}^{\text{drv}}(K)|_{f \leq h}$ are homotopy equivalent.

For some of the cases in the proof of the conjecture, we need the following preliminary observations:

Lemma 3.25. *If K is a simplicial complex with a map f linear on cells such that all vertices have positive f -values except for one vertex v_- with negative f -value, then for any stellar subdivision of K there is a combinatorially equivalent subdivision of K with the original vertices of K being fixed under the combinatorial equivalence such that v_- is still the only vertex with non-positive f -value. Note that the roles of positive and negative values can be switched by considering $-f$, and that the level $h = 0$ separating the value of the exceptional vertex from the values of the other vertices can be replaced by any arbitrary level $h \in \mathbb{R}$ by shifting the map f .*

When we apply Theorem 1.31 for two PL-homeomorphic simplicial complexes K and L to obtain combinatorially equivalent subdivisions K' of K and L' of L , we can assume without loss of generality that one of the subdivisions is stellar.

Proof. For the first claim, it suffices to consider a single elementary starring, the rest follows from induction on the number of elementary starrings performed in the stellar subdivision. Since K is simplicial, an elementary starring at v_- itself is a vacuous operation and the statement trivial. Otherwise, the point x at which the elementary starring is performed is interior to a cell X containing at least one vertex v with $f(v) > 0$. By continuity of f any interior point y of X sufficiently close to v fulfils $f(y) > 0$. We noted already in Fact 1.28 that starring at different points x and y in the interior of the same cell X lead to combinatorially equivalent results, and that the combinatorial equivalence in

terms of vertices is given by extending the identity on $\text{Vrt}(K)$ by the assignment of x to y .

For the second claim, we use the following fact cited from Hudson [26, Lemma 1.10, p. 17]: For simplicial complexes K and L with PL map $\phi: |K| \rightarrow |L|$, there is a subdivision K' of K and a stellar subdivision L' of L such that ϕ is simplicial with respect to K' and L' . Applying this result to a PL homeomorphism ϕ , we can observe that K' and L' are combinatorially equivalent because ϕ is a simplicial bijection between these subdivisions. For showing that we can alternatively assume that K' is stellar, we consider the inverse homeomorphism ϕ^{-1} instead. \square

Lemma 3.26. *Assume that g is a discrete Morse function on the combinatorial d -manifold M such that all values $g(S)$ of g for $S \in M$ are distinct and for any pair $T \sqsubset S$ such that (T, S) is not an immediate face/co-face pair, the inequality $g(T) < g(S)$ holds. Define a map $f: |\text{sdv}^{\text{drv}}(M)| \rightarrow \mathbb{R}$ linear on cells of a derived of M by simplex-wise linear interpolation between the values at the vertices of the derived given by $f(v_S) = g(S)$ for each vertex v_S chosen in the interior of a cell $S \in M$. Then f is a PL Morse function such that:*

1. *All non-vertices of $\text{sdv}^{\text{drv}}(M)$ are regular with respect to f .*
2. *If S is critical with respect to g and an interior cell of M , then v_S is non-degenerate critical with respect to f with index $\lambda = \dim(S)$.*
3. *If S is critical with respect to g and a boundary cell of M , then v_S is non-degenerate boundary-critical of lower type with index $\lambda = \dim(S)$.*
4. *If S is regular with respect to g and an interior cell of M , then v_S is regular with respect to f .*
5. *If S is regular with respect to g and a boundary cell of M matched with a face or co-face in the boundary, in other words also regular with respect to the restriction of g to $\text{bd}(M)$, then v_S is regular with respect to f .*
6. *If S is regular with respect to g and a boundary cell of M matched with a co-face not in the boundary, in other words critical with respect to the restriction of g to $\text{bd}(M)$, then v_S is non-degenerate boundary-critical of upper type with respect to f with index $\lambda = \dim S$.*

Proof. If all cells are critical with respect to g , the induced map f is of the type considered in Theorem 3.24. The proof given there for non-vertices relied only on the fact that f is non-constant on any simplex of dimension at least 1. And the proof for a vertex v_S relied only on the inequalities $f(v_T) < f(v_S)$ for proper faces T of S and $f(v_W) > f(v_S)$ for proper co-faces W of S . Therefore the same proofs cover the first three cases when applied to a non-vertex of the derived or a vertex v_S such that S is critical with respect to g , even if not all cells are critical with respect to g :

The First Three Cases. The non-vertices are regular by Lemma 3.22 because all values at vertices are distinct and consequently any non-vertex is contained in the interior of a cell where f is non-constant. For critical cells S , the link $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ of v_S in the derived is the join of a $(\lambda - 1)$ -sphere formed by vertices v_T with $T \sqsubset S$ and hence $f(v_T) < f(v_S)$ and a $(d - \lambda - 1)$ -sphere (-ball) formed by vertices v_W with $S \sqsubset W$ and hence $f(v_W) > f(v_S)$,

which proves local equivalence with $(\Omega_{(+)}^d, 0, f_\lambda^d)$ by considering the 1-norm unit (half-)sphere S_1^{d-1} as join of the $(\lambda - 1)$ -sphere formed by vertices of S_1^{d-1} with negative f_λ^d -value and the $(d - \lambda - 1)$ -sphere (-ball) formed by vertices with positive f_λ^d -value. Therefore v_S is critical with index $\lambda = \dim(S)$ for interior cells S and boundary-critical of lower type with index $\lambda = \dim(S)$ for boundary cells S .

We will use the representation of $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ used in the proof of Theorem 3.24 as the join of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ and $\text{sdv}^{\text{drv}}(\text{bd}(S))$ for the other cases as well, where ϕ^{comb} is the combinatorial equivalence in terms of cells between the join $\text{sdv}^{\text{drv}}(\text{lk}_M(S)) \text{sdv}^{\text{drv}}(\text{bd}(S))$ and $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ obtained from Lemma 3.23.

Cases 4 and 5. For the next two cases, assume that S is regular and either interior or paired with another boundary cell. Let $\lambda = \dim(S)$. As usual, modifications for the boundary case are indicated by parentheses. Our aim is to construct a link complex for the origin in $\mathbb{R}_{(+)}^d$ combinatorially equivalent to a subdivision of $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$ such that π_1^d and $f - f(v_S)$ have the same sign at corresponding vertices. There are two alternatives to consider: Either S is paired with an immediate face, or it is paired with an immediate co-face in the Morse matching.

Paired with a Face. We first deal with the alternative that S is paired with an immediate face $T \sqsubset S$ implying $\lambda \geq 1$. Then the inequalities $f(v_T) > f(v_S)$, $f(v_W) > f(v_S)$ for $S \sqsubset W$, and $f(v_F) < f(v_S)$ for $T \neq F \sqsubset S$ hold. Again the construction of the link complex of the origin is decomposed into two constructions in orthogonal subspaces: On the one hand, we construct a link complex combinatorially equivalent to $\text{sdv}^{\text{drv}}(\text{bd}(S))$ in the subspace defined by $X_i = 0$ for $\lambda + 1 \leq i \leq d$ that respects the signs after some perturbation within that subspace. On the other hand we construct a link complex in the subspace (sub-half-space) defined by $X_i = 0$ for $1 \leq i \leq \lambda$ (and $X_d \geq 0$) combinatorially equivalent to a subdivision of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$. Then we join the two link complexes and shift the vertices of the latter one slightly in positive X_1 -direction to obtain matching signs for its vertices as well.

For the construction of the link complex combinatorially equivalent to the derived of $\text{bd}(S)$ we proceed as follows: Choose as link complex L'_T for the origin in the $(\lambda - 1)$ -dimensional subspace defined by $X_i = 0$ for $i = 1 \vee \lambda + 1 \leq i \leq d$ the boundary of a $(\lambda - 1)$ -simplex T' containing the origin in its interior. Later on, some perturbed copies of the vertices of T' will correspond to the vertices of T under the combinatorial equivalence. As link complex L'_1 for the origin in the 1-dimensional subspace defined by $X_i = 0$ for $i \neq 1$, we choose a vertex v'_T with positive X_1 -coordinate and a vertex t' with negative X_1 -coordinate. The vertex v'_T will correspond to v_T and the vertex t' will correspond to the vertex t of S lying opposite to the face T , in other words, the vertex fulfilling $S = tT$. By Lemma 3.19, the join $L'_1 L'_T$ is a link complex for the origin in the subspace defined by $X_i = 0$ for $\lambda + 1 \leq i \leq d$. The suggested correspondences for the vertices yield a combinatorial equivalence between $L'_1 L'_T$ and the complex obtained from $\text{bd}(S)$ by performing an elementary starring at v_T . This elementary starring can be considered as the first step in constructing the derived subdivision of $\text{bd}(S)$ by starring at each chosen point v_F for $F \in \text{bd}(S)$ in order of decreasing dimension of F . By Fact 1.29, there is an analogous stellar subdivision L'_- of $L'_1 L'_T$ combinatorially equivalent to $\text{sdv}^{\text{drv}}(\text{bd}(S))$ such that the previously sug-

gested correspondences between vertices are maintained. Since in the remaining steps of the construction of $\text{sdv}^{\text{drv}}(\text{bd}(S))$ only points that are not interior to cells containing the first new vertex v_T are chosen as new vertices, all corresponding vertices in L'_- are not interior to cells of $L'_1 L'_T$ containing v'_T . This implies that v'_T is the only vertex of L'_- with positive X_1 -coordinate because all other vertices are interior to cells of $L'_1 L'_T$ not containing v'_T and v'_T is the only vertex of $L'_1 L'_T$ with positive X_1 -coordinate by construction. Hence, for each vertex v_F of $\text{sdv}^{\text{drv}}(\text{bd}(S))$ with $F \neq T$, the corresponding vertex v'_F in L'_- has non-positive X_1 -coordinate $\pi_1^d(v'_F) \leq 0$. Since we are aiming for negative values for the vertices corresponding to a vertex v_F of $\text{sdv}^{\text{drv}}(\text{bd}(S))$ with $F \neq T$, we perturb L'_- by shifting all vertices v'_F with $\pi_1^d(v'_F) = 0$ by an ϵ amount in negative X_1 -direction. By Lemma 3.21 we obtain a link complex L''_- of the origin in the subspace defined by $X_i = 0$ for $\lambda + 1 \leq i \leq d$ combinatorially equivalent to $\text{sdv}^{\text{drv}}(\text{bd}(S))$ with matching signs of π_1^d and $f - f(v_S)$ at corresponding vertices, provided $\epsilon > 0$ is small enough.

The existence of a link complex of the origin in the subspace (sub-half-space) defined by $X_i = 0$ for $1 \leq i \leq \lambda$ (and $X_d \geq 0$) combinatorially equivalent to a subdivision of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ follows immediately from the fact that any link complex for the origin in the considered subspace and the complex $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ are both combinatorial $(d - \lambda - 1)$ -spheres (-balls) and hence PL-homeomorphic. Since PL-homeomorphic spaces have combinatorially equivalent subdivisions, we can choose an arbitrary link complex L'_+ for the origin and there exists a subdivision L''_+ of L'_+ combinatorially equivalent to some subdivision L_+ of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$. Recall that the vertices of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ are of the form v_W for some proper co-face W of S and that $f(v_W) > f(v_S)$ holds for those vertices. Therefore the shifted map $f - f(v_S)$ is positive throughout $|\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))|$, in particular it is positive at the vertices of L_+ . According to Lemma 3.19, joining L''_- and L''_+ produces a link complex $L'' = L''_- L''_+$ for the origin in $\mathbb{R}_{(+)}^d$ combinatorially equivalent to the join of $\text{sdv}^{\text{drv}}(\text{bd}(S))$ and L_+ . This join in turn is a subdivision of $\text{sdv}^{\text{drv}}(\text{bd}(S)) \phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S))) = \text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$. By our construction so far, the sign of π_1^d at the vertices in L''_- already matches with the sign of $f - f(v_S)$ at corresponding vertices of $\text{sdv}^{\text{drv}}(\text{bd}(S))$, and the vertices in L''_+ all have π_1^d -value 0, whereas the $(f - f(v_S))$ -value of the corresponding vertices in L_+ is positive. Shifting the vertices of L''_+ by an ϵ amount in positive X_1 -direction yields an ϵ -perturbation L''' of L'' fulfilling all desired properties for some sufficiently small $\epsilon > 0$.

Paired with a Co-Face. For the alternative that S is paired with a co-face W (in the boundary), implying $\lambda < d$ ($\lambda < d - 1$), the inequalities $f(v_W) < f(v_S)$, $f(v_Y) > f(v_S)$ for $S \sqsubset Y \neq W$, and $f(v_T) < f(v_S)$ for $T \sqsubset S$ hold. The link complex for the origin is again formed by joining link complexes in complementary subspaces and shifting their vertices. On the one hand, we construct a link complex for the origin in the subspace (sub-half-space) defined by $X_i = 0$ for $2 \leq i \leq \lambda + 1$ (and $X_d \geq 0$) combinatorially equivalent to a subdivision L_+ of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ such that in the link complex, the vertex corresponding to v_W is the only vertex with negative X_1 -coordinate whereas all other vertices have positive X_1 -coordinate, and in L_+ all vertices except v_W have f -value strictly larger than $f(v_S)$. On the other hand, in the subspace defined by $X_i = 0$ for $i = 1 \vee \lambda + 2 \leq i \leq d$, we can choose a simplex S' contain-

ing the origin in its interior, then the derived of its boundary is combinatorially equivalent to $\text{sdv}^{\text{drv}}(\text{bd}(S))$ and after joining $\text{sdv}^{\text{drv}}(\text{bd}(S'))$ with the first link complex, we shift the vertices belonging to $\text{sdv}^{\text{drv}}(\text{bd}(S'))$ slightly in negative X_1 -direction. In the following, both constructions are explained in more detail.

For the first construction, choose an arbitrary link complex for the origin in the subspace (sub-half-space) defined by $X_i = 0$ for $2 \leq i \leq \lambda + 1$ (and $X_d \geq 0$). Since the chosen link complex and $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ are both combinatorial $(d - \lambda - 1)$ -spheres (-balls), there are combinatorially equivalent subdivisions L_+ of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ and L'_+ of the chosen link complex. Using both observations from Lemma 3.25, we can assume that L_+ is stellar and v_W remains the only vertex with non-positive $(f - f(v_S))$ -value. After performing a suitable rotation (fixing the boundary of the sub-half-space set-wise), we can assume without generality that the vertex v'_W in L'_+ corresponding to v_W lies on the negative X_1 -axis.

We modify L'_+ according to Fact 1.32 by moving its vertices along their rays emanating from the origin until all vertices have the same Euclidean distance to the origin. This yields a link complex L''_+ of the origin combinatorially equivalent to L'_+ and L_+ with vertices in convex position such that the vertex v''_W corresponding to v_W has negative X_1 -coordinate $\pi_1^d(v''_W)$ and all other vertices of L''_+ have X_1 -coordinate strictly larger than $\pi_1^d(v''_W)$.

Now we can shift L''_+ in positive X_1 -direction by an amount of $\pi_1^d(v''_W) - \delta$ for sufficiently small $\delta > 0$ such that the shifted copy of v''_W still has negative X_1 -coordinate but all other shifted vertices have positive X_1 -coordinate. Let L'''_+ denote the shifted copy. The convex position of the vertices and the fact that the vertex v'''_W corresponding to v_W obtained from shifting v''_W along the X_1 -axis still lies on the negative X_1 -axis ensures that $0L'''_+$ is a valid cone in the subspace (sub-half-space) defined by $X_i = 0$ for $2 \leq i \leq \lambda + 1$ (and $X_d \geq 0$). By Lemma 3.18, the complex L'''_+ is a link complex for the origin in that subspace (sub-half-space). Moreover it has the desired property that it is combinatorially equivalent to a subdivision of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ with v_W being the only vertex fulfilling $f(v_W) \leq f(v_S)$ such that the vertex corresponding to v_W has negative X_1 -coordinate but all other vertices have positive X_1 -coordinate.

The second construction in the subspace defined by $X_i = 0$ for $i = 1 \vee \lambda + 2 \leq i \leq d$ is easy. As sketched above, we choose a λ -simplex S' in that subspace containing the origin in its interior. Denote by L'_- the derived of the boundary of S' , then L'_- is a link complex for the origin in the subspace and combinatorially equivalent to $\text{sdv}^{\text{drv}}(\text{bd}(S))$. The join $L' = L'_- L'''_+$ is now a link complex for the origin by Lemma 3.19 and combinatorially equivalent to the subdivision $\text{sdv}^{\text{drv}}(\text{bd}(S))_{L_+}$ of $\text{lk}_{\text{sdv}^{\text{drv}}(M)}(v_S)$. After shifting the vertices from L'_- by some sufficiently small ϵ amount in negative X_1 -direction, the obtained ϵ -perturbation of L' has all desired properties.

The Last Case. It remains to consider the last case that S is a regular boundary cell matched with a co-face W not in the boundary. As in the second alternative of the previous two cases, the inequalities $f(v_W) < f(v_S)$, $f(v_Y) > f(v_S)$ for $S \sqsubset Y \neq W$, and $f(v_T) < f(v_S)$ for $T \sqsubset S$ hold. Again the proof splits into two constructions of link complexes of the origin in a subspace and a sub-half-space. Since the proposed reference map is g_λ^d , we use subdivisions of the 1-norm unit (half-)sphere in both parts, so that their join L' is a subdivision of the 1-norm unit half-sphere in \mathbb{R}_+^d and g_λ^d is linear on cells of $0L'$. The

easier part is the subspace $X_i = 0$ for $\lambda + 1 \leq i \leq d$, in which we construct a link complex for the origin combinatorially equivalent to some subdivision of $\text{sdv}^{\text{drv}}(\text{bd}(S))$. Note that $\lambda < d$ because S is a boundary cell. The harder part of the proof is the construction of a link complex in the sub-half-space defined by $X_i = 0$ for $1 \leq i \leq \lambda$ and $X_d \geq 0$ combinatorially equivalent to a subdivision of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ such that g_λ^d and $f - f(v_S)$ have the same sign at corresponding vertices.

So let us deal first with the construction of a link complex in the subspace $X_i = 0$ for $\lambda + 1 \leq i \leq d$. The reference map g_λ^d has the constant value -1 when restricted to the 1-norm unit sphere $S_1^{\lambda-1}$ in that subspace. On the other hand, the values attained by f on $|\text{sdv}^{\text{drv}}(\text{bd}(S))|$ are strictly smaller than $f(v_S)$ because of the inequality $f(v_T) < f(v_S)$ for $T \sqsubset S$. Therefore finding combinatorially equivalent subdivisions of $S_1^{\lambda-1}$ and $\text{sdv}^{\text{drv}}(\text{bd}(S))$ suffices. Such an equivalence exists because both complexes are combinatorial $(\lambda - 1)$ -spheres.

Now we turn to the construction of the link complex in the sub-half-space defined by $X_i = 0$ for $1 \leq i \leq \lambda$ and $X_d \geq 0$. For our proof, we need the combinatorial annulus theorem [44, Corollary 3.19, p. 36], which states that if A is a PL n -ball contained in the interior of another PL n -ball B , then the closure of $B \setminus A$ is PL-homeomorphic to an annulus, that is, the Cartesian product of a PL $(n - 1)$ -sphere and a compact interval. If necessary, perform as a preliminary step a stellar subdivision on $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ such that the star of v_W is contained completely in the interior of the $|\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))|$. Since W is an interior cell of M , the vertex v_W is an interior point of the $(d - \lambda - 1)$ -ball $|\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))|$ and such a stellar subdivision exists, for example the derived subdivision will do. By Lemma 3.25, we can assume that all vertices except v_W of the obtained stellar subdivision L_+ of $\phi^{\text{comb}}(\text{sdv}^{\text{drv}}(\text{lk}_M(S)))$ have f -value strictly larger than $f(v_S)$. The star of v_W in L_+ is a $(d - \lambda - 1)$ -ball, and we made sure that it is contained in the interior of the $(d - \lambda - 1)$ -ball $|L_+|$. Applying the combinatorial annulus theorem yields that $|\text{del}_{L_+}(v_W)|$, which is the closure of $|L_+| \setminus |\text{st}_{L_+}(v_W)|$, is a $(d - \lambda - 1)$ -dimensional annulus. We can perform an analogous construction on the 1-norm unit half-sphere H in the sub-half-space defined by $X_i = 0$ for $1 \leq i \leq \lambda$ and $X_d \geq 0$. Note that the vertex e_d with coordinates $X_d = 1$ and $X_i = 0$ for $i \neq d$ is the only vertex of H whose g_λ^d -value is -1 , all other vertices have g_λ^d -value 1. Choose a stellar subdivision H' of H such that all vertices except e_d have positive g_λ^d -value and $|\text{st}_{H'}(e_d)|$ is contained in the interior of $|H'|$. Then $|\text{del}_{H'}(e_d)|$ is a $(d - \lambda - 1)$ -dimensional annulus as well.

Find PL homeomorphisms from $|\text{del}_{H'}(e_d)|$ and from $|\text{del}_{L_+}(v_W)|$ to the standard annulus $S^{d-\lambda-2} \times [0, 1]$ such that the links $|\text{lk}_{H'}(e_d)|$ and $|\text{lk}_{L_+}(v_W)|$ are mapped to $S^{d-\lambda-2} \times \{1\}$. Then a composition of one PL homeomorphism with the inverse of the other yields a PL homeomorphism between $|\text{del}_{H'}(e_d)|$ and $|\text{del}_{L_+}(v_W)|$ mapping $|\text{lk}_{H'}(e_d)|$ and $|\text{lk}_{L_+}(v_W)|$ to each other. Consequently, there are combinatorially equivalent subdivisions A of $|\text{del}_{L_+}(v_W)|$ and A' of $|\text{del}_{H'}(e_d)|$ that restrict to combinatorially equivalent subdivisions of A_L of $|\text{lk}_{L_+}(v_W)|$ and A'_L of $|\text{lk}_{H'}(e_d)|$. By construction, $f - f(v_S)$ is positive throughout $|\text{del}_{L_+}(v_W)|$ and g_λ^d is positive throughout $|\text{del}_{H'}(e_d)|$ implying that the signs at corresponding vertices of A and A' agree. As final step for the second part, we add the cones $v_W A_L$, a subdivision of $|\text{st}_{L_+}(v_W)|$, to A and $e_d A'_L$, a subdivision of $|\text{st}_{H'}(e_d)|$, to A' . Then $A \cup v_W A_L$ is a subdivision of L_+ combinatorially equivalent to the subdivision $A' \cup e_d A'_L$ of H' and the signs of $f - f(v_S)$ and g_λ^d

agree at corresponding vertices. Now we only have to join $A' \cup e_d A'_L$ with the link complex in the other subspace considered in the first part to obtain a link complex for the origin in \mathbb{R}_+^d with all desired properties. \square

3.2.5 Comparison with Previous Notions

In the last part of this chapter, we compare the definition of regular and critical points for piecewise linear functions suggested in Definition 3.12 with similar notions of regularity and criticality introduced previously by other authors. The first two notions by Kosinski [29] and by Eells and Kuiper [17] are also based on the existence of a local coordinate system, in terms of which the function has a certain standard representation. Kosinski's approach even uses a definition of two PL maps being equivalent at a pair of points for describing the local coordinate system, which resembles our notion of local equivalence. Since the notions of Kosinski, the notions of Eells and Kuiper, and the notions suggested here are all based on the idea of using a piecewise linear analogue of the Morse lemma as definition for defining regular and non-degenerate critical points, it comes as no surprise that the three notions are in principle equivalent. Some of the differences lie in more or less general assumptions such as whether the complexes are finite or locally finite, which does not really matter because all notions classify a point only depending on local properties near the point, or whether manifolds with boundaries are considered or not. Our aim is to check that despite of different choices for the standard reference functions and other minor differences in the definitions, the obtained classification is the same whenever all three notions apply.

In contrast to the three notions where local PL homeomorphisms are required for identifying points as regular or critical, the other two notions reviewed here for classifying the points with respect to a piecewise linear function do not take all possible local topological features into account. We observe for the notion of Brehm and Kühnel [9] and for the notion of Banchoff [4] that the definitions can be rephrased in such a way that they can be described as classifying the points with respect to different notions of local equivalence that are coarser than the notion of local equivalence suggested in Definition 3.1. The classification by Brehm and Kühnel is based on homology, namely on the Betti numbers of the lower level link $(\text{lk}(v))_{\leq f(v)}$. Banchoff's classification is even coarser because it considers only the Euler characteristic of the lower level link.

All four notions described below consider piecewise linear maps on combinatorial manifolds. It is worth mentioning here two other related types of Morse theory that have been discussed previously and have some overlaps with piecewise linear Morse theory. Agrachev, Pallaschke, and Scholtes [3] have studied Morse theory for piecewise smooth functions on smooth manifolds. Amongst other results, they obtain a level preserving isotopy for intervals without critical values and the Morse inequalities. Rote [42] applies their characterisation of regular points to the piecewise linear case and illustrates with a 2-dimensional example that a point can be critical in that sense, although a piecewise linear level-preserving isotopy across the level of that point exists. Another type of Morse theory that is applicable in principle to the piecewise linear setting is stratified Morse theory by Goresky and MacPherson [22]. In this theory, Rote [42] observes, all vertices of the complex would be treated as critical points.

Whether there are deeper connections between these types of Morse theory and our approach has not been investigated further.

Definition of Kosinski

The approach of Kosinski in [29] splits into two parts. The first part introduces a notion of equivalence for maps at points similar to the notion of local equivalence from Definition 3.1. Based on that notion, the second part calls a map regular at a point or says that it has a non-degenerate singularity of a certain index at the point, if the map at that point is equivalent to some standard map at the origin of \mathbb{R}^d . Since our approach is also divided into parts along the same idea, we first compare Kosinski's equivalence relation with local equivalence, and show that they yield the same relation for real-valued PL maps on manifolds. For the second part, we check that Kosinski's standard maps are locally equivalent to our standard maps at the origin.

Kosinski's Equivalence Relation. For the definition of the equivalence of maps at points Kosinski uses open neighbourhoods and a notion of a continuous map being PL at a point. Since we prefer to consider cone neighbourhood complexes and maps linear on cells of such complexes, we start by explaining how to switch between the two models. Kosinski uses manifolds as domains for the maps, but the notion of a manifold differs slightly from Definition 1.36 for combinatorial manifolds: Locally finite simplicial complexes whose domain is a closed subset of \mathbb{R}^n are allowed as complexes as well, and the complex is required to be connected. A continuous map $f: V \rightarrow |M_1|$ from an open subset V of a manifold $|M|$ to another manifold $|M_1|$ is called PL at a point $p \in V$ if p has an open neighbourhood U and $|M|$ and $|M_1|$ have a triangulations K and K_1 such that for any simplex S of K intersecting U the intersection $S \cap V$ is mapped linearly into a simplex of K_1 . We want to sketch how this definition of f being PL at p can be seen as equivalent to the condition that p has a cone neighbourhood complex pL in $|M|$ such that f is linear on cells of pL .

Let us assume first that f is linear on cells of a cone neighbourhood pL . Recall from Subsection 1.2.4 that f being linear on cells of pL implies that there are subdivisions N of pL and K_1 of M_1 such that f is simplicial with respect to the subdivisions. Choose a subdivision K of M that contains some subdivision of N , and let $U = V$ be some open neighbourhood of p contained in $|pL|$. Then each simplex of K intersecting U is contained in a simplex of N and is mapped linearly into some simplex of K_1 . Now assume that f is PL at p in the sense defined above. Choose some cone neighbourhood complex pL' contained in $U \cap V$, form the intersection complex of pL' with K , and triangulate the result to obtain a simplicial subdivision N' of pL' . Let $L = \text{lk}_{N'}(p)$. Then pL is a cone neighbourhood complex for p in $|M|$, and each of its simplices is completely contained in U , in V , and in a simplex of K . Hence each simplex of pL is mapped linearly into some simplex of the triangulation K_1 of $|M_1|$, and therefore f is linear on cells of pL .

Kosinski's definition of maps being equivalent at a pair of points is not restricted to real-valued functions, also higher dimensional Euclidean spaces are allowed as co-domains. Consider two manifolds $|M|$ and $|M'|$ with PL maps $f: |M| \rightarrow \mathbb{R}^k$ and $f': |M'| \rightarrow \mathbb{R}^k$ and two points $p \in |M|$ and $p' \in |M'|$. Then Kosinski calls f at p *equivalent* to f' at p' if there are neighbourhoods U_p of p

$$\begin{array}{ccc}
U_p & \xrightarrow{\phi} & U_{p'} \\
\downarrow f & & \downarrow f' \\
\mathbb{R}^k & \xrightarrow{h} & \mathbb{R}^k
\end{array}$$

Figure 3.4: Commuting diagram for Kosinski's notion of f at p being equivalent to f' at p' .

in $|M|$ and $U_{p'}$ of p' in $|M'|$ and topological homeomorphisms $h: \mathbb{R}^k \rightarrow \mathbb{R}^k$ and $\phi: U_p \rightarrow U_{p'}$ such that ϕ maps p to p' and is PL at p , the map h is orientation preserving and PL at $f(p)$ (in [29], it says PL at $f'(p')$, which seems to be a mistake), and the compositions $h \circ f$ and $f' \circ \phi$ agree on U_p as illustrated by the commuting diagram from Figure 3.4.

Comparison with Local Equivalence. For comparing Kosinski's equivalence relation with the local equivalence relation from Definition 3.1, let us first check that local equivalence implies equivalence in the sense of Kosinski. If f at p is locally equivalent to f' at p' , there are neighbourhoods N_p of p and $N_{p'}$ of p' , a PL homeomorphism $\phi: N_p \rightarrow N_{p'}$ mapping p to p' , and a constant $c \in \mathbb{R}$ such that $f' \circ \phi = f + c$ holds on N_p . As a PL homeomorphism, ϕ is clearly PL at p . Choosing for $h: \mathbb{R} \rightarrow \mathbb{R}$ the map defined by $h(t) = t + c$, we see immediately that h is PL at every point and orientation preserving, and that $f' \circ \phi = h \circ f$ holds on N_x . Therefore we obtain Kosinski's definition by restricting ϕ to some open neighbourhood U_p of p contained in N_p and its image $U_{p'} = \phi(U_p)$ as open neighbourhood of p' contained in $N_{p'}$.

Now we show for real-valued maps f and f' that f at p being equivalent to f' at p' in the sense of Kosinski implies local equivalence of f at p and f' at p' . So assume that the equivalence is witnessed by neighbourhoods U_p and $U_{p'}$ with homeomorphisms $\phi: U_p \rightarrow U_{p'}$ and $h: \mathbb{R} \rightarrow \mathbb{R}$ such that the requirements of the definition are fulfilled. Since h is PL at $f(p)$ and orientation preserving, we can find positive real numbers λ_- , λ_+ , and ϵ such that $h(t) = \lambda_-(t - f(p)) + f'(p')$ for $t \in [f(p) - \epsilon, f(p)]$ and $h(t) = \lambda_+(t - f(p)) + f'(p')$ for $t \in [f(p), f(p) + \epsilon]$. Let K and K' be triangulations of $|M|$ and $|M'|$ such that f and f' are linear on cells of the respective triangulations. Choose a suitable link complex L inducing a sufficiently small cone neighbourhood pL of p such that the following assumptions hold: The homeomorphism g , which is PL at p , restricts to a simplicial bijection between pL and a cone neighbourhood $p'L'$ for some link complex $L' = \phi(L)$ of p' , in particular L is contained in U_p and L' is contained in $U_{p'}$; the link complex L is contained in $|\text{st}_K(p)|$, and L' is contained in $|\text{st}_{K'}(p')|$; for all $x \in |pL|$ the inequality $|f(x) - f(p)| < \epsilon$ holds.

When we prove that $f(x) - f(p)$ and $f'(\phi(x)) - f'(p')$ have the same sign for any $x \in |L|$, then Lemma 3.9 shows that f at p is locally equivalent to f' at p' . Since $f' \circ \phi$ agrees with $h \circ f$ on $|L| \subseteq U_p$, we have $f'(\phi(x)) - f'(p') = h(f(x) - f(p)) - f'(p') = h(f(x) - f(p)) - f'(p')$. Since we assumed that $f(x)$ lies in an ϵ -neighbourhood of $f(p)$ for $x \in |L|$, where h is represented by the terms mentioned above, we obtain $h(f(x) - f(p)) - f'(p') = \lambda(f(x) - f(p))$ for $\lambda = \lambda_- > 0$ if $f(x) \leq f(p)$ and $\lambda = \lambda_+ > 0$ if $f(x) \geq f(p)$. In any case, we have $f'(\phi(x)) - f'(p') = \lambda(f(x) - f(p))$ for

Table 3.2: Rules for calculating a term $\Theta_i(X_i)$ from $f(e_i)$ and $f(-e_i)$ to represent Kosinski's standard maps by a term $f(X_1, \dots, X_d) = \sum_{i=1}^d \Theta_i(X_i)$.

$f(e_i)$	$f(-e_i)$	$\Theta_i(X_i)$
1	1	$ X_i $
1	-1	X_i
-1	1	$-X_i$
-1	-1	$- X_i $

some positive number λ , showing that $f(x) - f(p)$ and $f'(\phi(x)) - f'(p')$ have the same sign.

Kosinski's Standard Maps. For the definition of the standard reference maps, Kosinski considers the 1-norm unit ball in \mathbb{R}^d triangulated as a cone $0S_1^{d-1}$ with apex at the origin 0 and the natural triangulation of the 1-norm unit sphere S_1^{d-1} as base. The maps are obtained by simplex-wise linear interpolation between given values at the vertices. As before, we denote the vertices of S_1^{d-1} by e_i and $-e_i$ where e_i is the point with coordinates $X_i = 1$ and $X_j = 0$ for $j \neq i$.

The standard reference maps of Kosinski are those maps f that fulfil $f(0) = 0$ and $f(v) \in \{-1, 1\}$ for vertices v of S_1^{d-1} . It is easy to check that these maps can be characterised as those maps that are representable by a term $f(X_1, \dots, X_d) = \sum_{i=1}^d \Theta_i(X_i)$ where each $\Theta_i(X_i)$ is one of the terms X_i , $-X_i$, $|X_i|$, or $-|X_i|$. Which case applies for Θ_i depends on the values $f(e_i)$ and $f(-e_i)$ as listed in Table 3.2.

With these standard maps, Kosinski calls a piecewise linear map $f': |M'| \rightarrow \mathbb{R}$ on a d -manifold $|M'|$ *regular* at $p' \in |M'|$ if f' at p' is equivalent to a standard map f at 0, such that $f(e_i) = -f(-e_i)$ for at least one index i in the range $1 \leq i \leq d$, in other words $\Theta_i(X_i) \in \{X_i, -X_i\}$ for at least one index i . The map f' is said to have a *non-degenerate singularity* of index λ at p instead, if f' at p' is equivalent to a standard map f at 0, such that $f(e_i) = f(-e_i) = -1$ for exactly λ indices i in the range $1 \leq i \leq d$ and $f(e_i) = f(-e_i) = 1$ for the remaining indices, in other words $\Theta_i(X_i) = -|X_i|$ for exactly λ indices i and $\Theta_i(X_i) = |X_i|$ for the remaining indices.

Let us verify that Kosinski's notions agree with the corresponding notions from Definition 3.12: A PL map f' is regular at p' in Kosinski's sense if and only if p' is a regular (interior) point of f' in the sense of Definition 3.12; it has a non-degenerate singularity of index λ at p' if and only if p' is a non-degenerate critical (interior) point of f' of index λ . We only have to check that Kosinski's reference maps are locally equivalent at the origin to the corresponding reference maps from Definition 3.12.

This is easy for the reference maps characterising critical points. The map f_λ^d is one of the reference maps used by Kosinski for a singularity of index λ . All others are obtained by a permutation of the coordinates, that is, if f fulfils $f(e_i) = f(-e_i) = -1$ for exactly λ indices i in the range $1 \leq i \leq d$ and $f(e_i) = f(-e_i) = 1$ for the remaining indices, then there is a permutation τ of the indices

such that $f(X_1, \dots, X_d) = f_\lambda^d(X_{\tau(1)}, \dots, X_{\tau(d)})$. But a permutation of the coordinates is a linear homeomorphism from \mathbb{R}^d to itself fixing the origin, showing that f_λ^d is locally equivalent to any map $f(X_1, \dots, X_d) = f_\lambda^d(X_{\tau(1)}, \dots, X_{\tau(d)})$ obtained by a permutation τ of the coordinates.

For showing that Kosinski's reference maps f with $f(e_i) = -f(-e_i)$ for at least one index i for regular points are locally equivalent to π_1^d , we can first observe that the index i with $f(e_i) = -f(-e_i)$ can be assumed to be $i = 1$. A permutation of the coordinates as described in the previous case shows that any reference map f with $f(e_i) = -f(-e_i)$ for at least one index i is locally equivalent at the origin to a reference map f with $f(e_1) = -f(-e_1)$. Furthermore, the linear automorphism of \mathbb{R}^d mapping each point x to its negative $-x$ maps the origin to itself and witnesses that each reference map f with $f(e_1) = -f(-e_1)$ is locally equivalent at the origin to a reference map f with $f(e_1) = -f(-e_1) = 1$. Therefore it suffices to prove local equivalence with π_1^d only for such maps f with $f(e_1) = -f(-e_1) = 1$.

This is verified by an ϵ -perturbation of the 1-norm unit sphere S_1^{d-1} such that π_1^d has the same sign at the perturbed vertices as f has at the original vertices. The vertices e_1 and $-e_1$ have already matching signs, and all other vertices lie in the hyperplane $X_1 = 0$ and can be shifted slightly in negative/positive X_1 -direction such that the sign of π_1^d at the shifted vertex agrees with the sign of f at the original vertex.

Definition of Eells and Kuiper

The notion that we want to examine next appeared in [17], where Eells and Kuiper consider manifolds with non-degenerate maps that have exactly three critical points in the topological, the combinatorial, and the differentiable category. For the combinatorial category, they define a PL function $f: |M| \rightarrow \mathbb{R}$ as C^{omb} -non-degenerate if each point of $|M|$ is either C^{omb} -ordinary or C^{omb} -critical of index λ in the following sense: A p -centred local coordinate system is a PL homeomorphism $X: U \rightarrow \mathbb{R}^d$ defined on some neighbourhood U of p such that $X(p)$ is the origin; the coordinate system is represented by a family $X_i: U \rightarrow \mathbb{R}$ of coordinate maps with $X(y) = (X_1(y), \dots, X_d(y))$ for $y \in U$. A point p is C^{omb} -ordinary with respect to f if there are a p -centred local coordinate system X_i in a neighbourhood U and a positive number μ such that $X_d(y) = \mu(f(y) - f(p))$ holds for $y \in U$. Likewise, a point p is C^{omb} -critical of index λ if there are a p -centred local coordinate system X_i in a neighbourhood U and a positive number μ such that $-\max_{i=1}^\lambda |X_i(y)| + \max_{i=\lambda+1}^d |X_i(y)| = \mu(f(y) - f(p))$ holds for $y \in U$. In the following, we show that the point classification is equivalent to the classification from Definition 3.12 regarding interior points: A C^{omb} -ordinary point is regular and vice versa, a C^{omb} -critical point of index λ is a critical point of index λ and vice versa.

The two conditions can be rearranged to express the shifted map $f - f(p)$ in terms of the coordinates X_i . A C^{omb} -ordinary point fulfils Equation (3.1), and a C^{omb} -critical point of index λ fulfils Equation (3.2):

$$f(y) - f(p) = \mu^{-1} X_d(y) \tag{3.1}$$

$$f(y) - f(p) = \mu^{-1} \left(-\max_{i=1}^\lambda |X_i(y)| + \max_{i=\lambda+1}^d |X_i(y)| \right) \tag{3.2}$$

In this form, the definition of Eells and Kuiper can be easily seen as equivalent to a rephrasing using local equivalence: The point p is C^{omb} -ordinary if f at p is locally equivalent to $\mu^{-1}X_d$ at the origin, and it is C^{omb} -critical of index λ if f at p is locally equivalent to $\mu^{-1}(-\max_{i=1}^{\lambda}|X_i| + \max_{i=\lambda+1}^d|X_i|)$ at the origin. The scaling factor μ^{-1} can be omitted because a scaling is a linear homeomorphism of \mathbb{R}^d to itself fixing the origin. The map X_d at the origin is locally equivalent to X_1 at the origin as witnessed by a permutation of the coordinates mapping the last to the first coordinate. Therefore a point is C^{omb} -ordinary if and only if it is a regular (interior) point in the sense of Definition 3.12.

It remains to prove the following statement:

Lemma 3.27. *For each λ in the range $0 \leq \lambda \leq d$, the reference map $h_{\lambda}^d = -\max_{i=1}^{\lambda}|X_i| + \max_{i=\lambda+1}^d|X_i|$ for C^{omb} -critical points of index λ and the reference map $f_{\lambda}^d = -\sum_{i=1}^{\lambda}|X_i| + \sum_{i=\lambda+1}^d|X_i|$ for PL critical points of index λ are locally equivalent at the origin.*

Proof. Our aim is to find a subdivision of the ∞ -norm unit sphere S_{∞}^{d-1} , which is a hypercube, combinatorially equivalent to the derived of the standard triangulation of the 1-norm unit sphere S_1^{d-1} such that for a suitable choice of the vertices of the derived, the sign of h_{λ}^d at a vertex in the hypercube triangulation agrees with the sign of f_{λ}^d at the corresponding vertex of the derived of S_1^{d-1} . Of course we obey the requirement from Lemma 3.7 allowing to conclude local equivalence that h_{λ}^d is linear on cells of the cone with apex at the origin and the chosen triangulation of the hypercube as base. The analogous requirement for f_{λ}^d is already fulfilled because we use a subdivision of the standard triangulation of S_1^{d-1} where each cell is contained in an orthant and f_{λ}^d is linear on orthants.

Subdividing the Hypercube. A subdivision of \mathbb{R}^d such that h_{λ}^d is linear on cells of the subdivision is induced by cutting the space along the following hyperplanes: Each coordinate hyperplane $X_i = 0$, which separates the half-space $|X_i| = X_i$ from the half-space $|X_i| = -X_i$, each diagonal hyperplane $X_i = X_j$ for $i \neq j$, which separates the half-space $\max(X_i, X_j) = X_i$ identical with $\max(-X_i, -X_j) = -X_j$ from the half-space $\max(X_i, X_j) = X_j$ identical with $\max(-X_i, -X_j) = -X_i$, and each diagonal hyperplane $X_i = -X_j$ for $i \neq j$, which separates the half-space $\max(X_i, -X_j) = X_i$ identical with $\max(-X_i, X_j) = X_j$ from the half-space $\max(X_i, -X_j) = -X_j$ identical with $\max(-X_i, X_j) = -X_i$. Since all hyperplanes contain the origin we obtain a polyhedral fan, which is a subdivision of the polyhedral fan formed by the orthants. We will see that this polyhedral fan induces the barycentric subdivision of the standard triangulation on S_1^{d-1} and a combinatorially equivalent triangulation of S_{∞}^{d-1} .

For any hyperplane $X_i = 0$, the whole situation is symmetric with respect to a reflection at the hyperplane $X_i = 0$. Therefore we can restrict our attention to the first orthant $X_i \geq 0$ for all i . The construction for the first orthant can be extended to the whole space by taking symmetric copies of the construction in the other orthants.

Consider first the ∞ -norm unit sphere S_{∞}^{d-1} and the ∞ -norm unit ball B_{∞}^d bounded by it. The ball B_{∞}^d is the cone $0S_{\infty}^{d-1}$ on the ∞ -norm unit sphere with apex at the origin and it has the same domain as the hypercube $[-1, 1]^d$. Its intersection with the first orthant is the hypercube $[0, 1]^d$. So the part of S_{∞}^{d-1} lying in the first orthant can be viewed as a subcomplex of $[0, 1]^d$ consisting of

the faces not containing the origin. Consequently, restricting the subdivision of $[0, 1]^d$ induced by the diagonal hyperplanes $X_i = X_j$ (the other diagonal hyperplanes $X_i = -X_j$ intersect the first orthant only in one of its proper faces and do not affect the subdivision) to this subcomplex is another way to describe the subdivision induced by the diagonal hyperplanes on the part of S_∞^{d-1} lying in the first orthant.

The subdivision of the hypercube $[0, 1]^d$ obtained from cutting it along the diagonal hyperplanes is a triangulation of the hypercube reviewed in [16, Section 6.3.2, pp. 314–315], where several alternative characterisations and additional properties of the triangulation are observed. Important for us are the following properties of the triangulation stated there:

The triangulation does not introduce additional vertices to the hypercube; in fact, it can be obtained by performing an elementary starring at each vertex of the hypercube provided that the order of the vertices respects the following restriction: Whenever two vertices $v = (v_1, \dots, v_d)$ and $u = (u_1, \dots, u_d)$ of the hypercube $[0, 1]^d$ fulfil $v_i \leq u_i$ for all indices i , then the starring at v is performed before the starring at u .

The set of maximal dimensional simplices of the triangulation is in bijection with the set of permutations of the index set $\{1, \dots, d\}$. The bijection is given by representing a maximal dimensional simplex as a set $T_\tau = \{(x_1, \dots, x_d) \mid 0 \leq x_{\tau(1)} \leq \dots \leq x_{\tau(d)} \leq 1\}$ for some permutation τ . When we choose for vertices $v^0(\tau), \dots, v^d(\tau)$ the coordinates $v^i(\tau)_{\tau(j)} = 0$ if $j \leq i$ and $v^i(\tau)_{\tau(j)} = 1$ for $j > i$, the convex hull of these vertices is the simplex T_τ .

We can draw the following additional conclusions from these observations: The half-spaces defining T_τ are $0 \leq X_{\tau(1)}, X_{\tau(1)} \leq X_{\tau(2)}, \dots, X_{\tau(d-1)} \leq X_{\tau(d)}, X_{\tau(d)} \leq 1$. The position of T_τ with regard to the other subdividing hyperplanes is implied, T_τ lies in the following half-spaces: $X_i \geq 0, X_i \leq 1$, and $X_i \geq -1$ for all i ; $X_i \geq -X_j$ for all $i \neq j$; and $X_{\tau(i)} \leq X_{\tau(j)}$ for all $i < j$. All maximal simplices contain the origin as vertex, hence the triangulation is a cone with apex at the origin whose base is the triangulation we are looking for, a triangulation of the part of S_∞^{d-1} lying in the first orthant.

Let us denote this triangulation by S_∞^+ and a maximal simplex of S_∞^+ whose cone with the origin is T_τ by S_τ , so that $T_\tau = 0S_\tau$ holds for each permutation τ . Clearly we can describe S_τ as intersection of T_τ with one of its bounding hyperplanes, namely with the hyperplane $X_{\tau(d)} = 1$. In contrast, omitting the half-space $X_{\tau(d)} \leq 1$ from the family of half-spaces defining T_τ yields the extended cone S_τ^+ . From the characterisation of the maximal simplices we can infer that a set of vertices forms a simplex of S_∞^+ if and only if there is a permutation τ and a subset $J \subseteq \{0, \dots, d-1\}$ such that the set agrees with $\{v^i(\tau) \mid i \in J\}$.

The corresponding description of such a simplex of S_∞^+ as intersection of T_τ with some of its bounding hyperplanes is obtained as follows: The bounding hyperplane $X_{\tau(d)} = 1$ always belongs to the intersection; the hyperplane $0 = X_{\tau(1)}$ belongs to the intersection if and only if all vertices $v^i(\tau)$ with $i \in J$ have coordinate $v^i(\tau)_{\tau(1)} = 0$; the hyperplane $X_{\tau(k)} = X_{\tau(k+1)}$ belongs to the intersection if and only if for each vertex $v^i(\tau)$ with $i \in J$ the $\tau(k)$ -th coordinate $v^i(\tau)_{\tau(k)}$ and the $\tau(k+1)$ -th coordinate $v^i(\tau)_{\tau(k+1)}$ agree.

Combinatorial Equivalence by Pseudo-Radial Projection. Now let S_1^+ be the barycentric subdivision of the simplex Δ of S_1^{d-1} forming its intersection with the first orthant. We want to convince ourselves that S_1^+ and

S_∞^+ are combinatorially equivalent by pseudo-radial projection. As before, we denote the vertices of Δ by e_i with e_i having i -th coordinate 1 and all other coordinates zero. For each non-empty face F of Δ represented as convex hull $F = \text{conv}\{e_k \mid k \in K\}$ for some non-empty subset $K \subseteq \{1, \dots, d\}$, the barycentre b_F is given by $b_F = \|\sum_{k \in K} e_k\|_1^{-1} \sum_{k \in K} e_k$. On the other hand, the terms $\sum_{k \in K} e_k$ for non-empty $K \subseteq \{1, \dots, d\}$ enumerate the vertices of S_∞^+ . Since for each non-empty subset K , the two points $\|\sum_{k \in K} e_k\|_1^{-1} \sum_{k \in K} e_k$ and $\sum_{k \in K} e_k$ lie on the same ray emanating from the origin, pseudo-radial projection yields as combinatorial equivalence $\phi^{\text{Vrt}}: \text{Vrt}(S_\infty^+) \rightarrow \text{Vrt}(S_1^+)$ in terms of vertices a map defined by the assignments $\phi^{\text{Vrt}}(\sum_{k \in K} e_k) = \|\sum_{k \in K} e_k\|_1^{-1} \sum_{k \in K} e_k$ for each non-empty $K \subseteq \{1, \dots, d\}$.

Consider a non-empty face F_∞^+ of S_∞^+ defined by relations between the coordinates $0 R_0 X_{\tau(1)}, X_{\tau(1)} R_1 X_{\tau(2)}, \dots, X_{\tau(d-1)} R_{d-1} X_{\tau(d)}, X_{\tau(d)} = 1$ where each relation R_i is either the order relation \leq or the equality $=$ and at least one relation is not the equality. As noted above, the extended cone F_∞^+ is defined by omitting the last relation $X_{\tau(d)} = 1$. The vertices of F_∞^+ are of the form $v^i(\tau)$ with $i \in J$ where J is the subset of $\{0, \dots, d-1\}$ fulfilling $J = \{i \mid R_i \text{ is the order relation } \leq\}$. In fact, $v^i(\tau)$ is the point defined by the relations $0 R_0^i X_{\tau(1)}, X_{\tau(1)} R_1^i X_{\tau(2)}, \dots, X_{\tau(d-1)} R_{d-1}^i X_{\tau(d)}, X_{\tau(d)} = 1$ where R_j^i is the order relation \leq and all other R_j^i with $j \neq i$ are the equality $=$.

By definition of the $v^i(\tau)$, they can be represented as $v^i(\tau) = \sum_{j>i} e_{\tau(j)}$. Define for each permutation τ and each index i in the range $0 \leq i \leq d-1$ a face $F^i(\tau)$ of Δ by taking the convex hull of the vertices $e_{\tau(j)}$ with $j > i$. It follows immediately that $F^{i+1}(\tau)$ is a face of $F^i(\tau)$ for any applicable index i , meaning that the $F^i(\tau)$ with $i \in J$ form a chain in the face poset of Δ . Moreover the barycentre of $F^i(\tau)$ is the point $b_{F^i(\tau)} = \|v^i(\tau)\|_1^{-1} v^i(\tau)$. Since the $F^i(\tau)$ form a chain, the corresponding barycentres $b_{F^i(\tau)}$ with $i \in J$ span a cell F_1 of the barycentric subdivision.

A barycentre $b_{F^i(\tau)}$ can be represented by the relations $0 R_0^i X_{\tau(1)}, X_{\tau(1)} R_1^i X_{\tau(2)}, \dots, X_{\tau(d-1)} R_{d-1}^i X_{\tau(d)}, \sum_{i=1}^d X_{\tau(i)} = 1$ with the same relations R_j^i as above. Since the intersection of F_∞^+ with Δ is given by the relations $0 R_1 X_{\tau(1)}, X_{\tau(1)} R_2 X_{\tau(2)}, \dots, X_{\tau(d-1)} R_d X_{\tau(d)}, \sum_{i=1}^d X_{\tau(i)} = 1$, we can conclude that pseudo-radial projection of the cell F_∞^+ to Δ is exactly the cell F_1 spanned by the barycentres $b_{F^i(\tau)}$ with $i \in J$. Note that we can find for any cell of S_1^+ spanned by barycentres b_{G_k} with the $(G_k)_{k \in K}$ forming a chain in the face poset a permutation τ and a subset J of $\{0, \dots, d-1\}$ such that the chains $(G_k)_{k \in K}$ and $(F^i(\tau))_{i \in J}$ consist of the same set of faces.

Let us denote the subdivision of the ∞ -norm unit sphere and hypercube boundary $S_\infty^{d-1} = \text{bd}([-1, 1]^d)$ induced by cutting along the hyperplanes $X_i = 0$, $X_i = X_j$, and $X_i = -X_j$ by S'_∞ and the barycentric subdivision of the standard triangulation of the 1-norm unit sphere S_1^{d-1} by S'_1 . With the above considerations, we have verified that S'_∞ and S'_1 are combinatorially equivalent via pseudo-radial projection with the combinatorial equivalence in terms of vertices $\phi'^{\text{Vrt}}: \text{Vrt}(S'_\infty) \rightarrow \text{Vrt}(S'_1)$ scaling the distance from the origin for each vertex from unit length with respect to ∞ -norm to unit length with respect to 1-norm.

A Derived with Matching Signs. However, with the barycentric subdivision the signs of h_λ^d at the vertices of S'_∞ and the signs of f_λ^d at the vertices of S'_1 do not always agree at corresponding vertices. To overcome this obstacle, we can use the fact that a barycentric subdivision is combinatorially equivalent to any

other derived subdivision. Therefore we can alter S'_1 and ϕ'_{Vrt} to a derived subdivision S_1^* of S_1^{d-1} and a combinatorial equivalence $\phi^{*\text{Vrt}}: \text{Vrt}(S'_\infty) \rightarrow \text{Vrt}(S_1^*)$ by replacing the barycentres $b_F \in \text{Vrt}(S'_1)$ of the cells F in S_1^{d-1} by other points v_F in the interior of the respective cell F . We show that it is possible to choose for each cell F a vertex v_F such that f_λ^d has the same sign at v_F as h_λ^d at the corresponding vertex in S'_∞ .

For defining the vertices, we divide the coordinates of the vertices in S'_∞ into the head block of the first λ coordinates (X_1, \dots, X_λ) and the tail block of the last $d - \lambda$ coordinates $(X_{\lambda+1}, \dots, X_d)$. If the tail block of a vertex $u = (u_1, \dots, u_d)$ of S'_∞ is zero, then its h_λ^d -value is negative and $(u_1, \dots, u_\lambda, 0, \dots, 0)$ is mapped by ϕ'^{Vrt} to the barycentre of a simplex F in S'_1 whose vertices are among the e_k with $k \leq \lambda$ and hence have the f_λ^d -value -1 . We choose an arbitrary point v_F in the interior of F for $\phi^{*\text{Vrt}}(u_1, \dots, u_\lambda, 0, \dots, 0)$, which inherits the negative f_λ^d -value -1 from the vertices by interpolation.

Similarly, if the head block is zero, then $(0, \dots, 0, u_{\lambda+1}, \dots, u_d)$ has positive h_λ^d -value and is mapped by ϕ'^{Vrt} to the barycentre of a simplex F in S'_1 whose vertices are among the e_k with $k \geq \lambda + 1$ and hence have the f_λ^d -value 1 . We choose an arbitrary point v_F in the interior of F for $\phi^{*\text{Vrt}}(0, \dots, 0, u_{\lambda+1}, \dots, u_d)$, which necessarily has the positive f_λ^d -value 1 .

The remaining vertices have non-zero head and tail blocks and consequently their h_λ^d -value is zero. They are mapped to the barycentre of a simplex F in S_1 that has vertices e_k with $k \geq \lambda + 1$ and f_λ^d -value 1 and vertices e_j with $j \leq \lambda$ and f_λ^d -value -1 . This implies that the term $F_{f_\lambda^d=0}$ is a canonical representation of its level set and hence we can choose an interior point v_F of the simplex with f_λ^d -value $f_\lambda^d(v_F) = 0$ for $\phi^{*\text{Vrt}}(u_1, \dots, u_d)$. This defines a derived subdivision S_1^* of S_1^{d-1} and a combinatorial equivalence $\phi^{*\text{Vrt}}: \text{Vrt}(S'_\infty) \rightarrow \text{Vrt}(S_1^*)$ with matching signs of h_λ^d and f_λ^d at corresponding vertices as desired.

Alternative Proof. Note that there is an alternative proof for the local equivalence of f_λ^d and h_λ^d at the origin by joining link complexes for the head and tail block subspaces. For a fixed index λ , a subdivision of \mathbb{R}^d obtained by cutting along the coordinate hyperplanes $X_i = 0$ and only those diagonal hyperplanes $X_i = X_j$ and $X_i = -X_j$ where the indices i and j either both belong to the head block or both belong to the tail block suffices for h_λ^d being linear on cells. When we construct a subdivision of $S_1^{\lambda-1}$ in the head block subspace and a subdivision of $S_1^{d-\lambda-1}$ in the tail block subspace such that h_λ^d is linear on cells of the cone on these subdivisions with the origin as apex, then h_λ^d is linear on cells of the cone on the join of the subdivisions as well. Since h_λ^d and f_λ^d are both negative on the whole $S_1^{\lambda-1}$ in the head block subspace and are both positive on the whole of $S_1^{d-\lambda-1}$ in the tail block subspace, matching signs at the vertices of the join are obtained immediately for this construction. \square

Definition of Brehm and Kühnel

Another notion of criticality suggested by Brehm and Kühnel in [9] is based on homology. An alternative reference for the notion is Kühnel's treatment in [30]. Brehm and Kühnel consider maps $f: |M| \rightarrow \mathbb{R}$ linear on cells of a combinatorial d -manifold M and in general position. In their terminology such maps are called *regular simplex-wise linear* and due to the general position assumption criticality occurs only at vertices. They call a vertex v in $|M|$ *critical* if not all

of the following singular relative homology groups vanish.

$$H_i(|M|_{\leq f(v)}, |M|_{\leq f(v)} \setminus \{v\})$$

Furthermore they remark that these singular homology groups are isomorphic to simplicial relative homology groups as follows.

$$H_i(|M|_{\leq f(v)}, |M|_{\leq f(v)} \setminus \{v\}) \cong H_i(\text{scp}_{\leq f(v)}(\text{st}(v)), \text{scp}_{\leq f(v)}(\text{lk}(v)))$$

For the latter relative homology group an isomorphic non-relative reduced homology group can be found using an observation in Hatcher [25, p. 125] stating for arbitrary spaces $Y \subseteq X$ that $H_i(X, Y)$ is isomorphic to $\tilde{H}_i(X \cup \text{cone}(Y))$.

$$\begin{aligned} H_i(\text{scp}_{\leq f(v)}(\text{st}(v)), \text{scp}_{\leq f(v)}(\text{lk}(v))) \\ \cong \tilde{H}_i(\text{scp}_{\leq f(v)}(\text{st}(v)) \cup \text{cone}(\text{scp}_{\leq f(v)}(\text{lk}(v)))) \end{aligned}$$

Since $\text{scp}_{\leq f(v)}(\text{st}(v))$ is a cone with apex v and base $\text{scp}_{\leq f(v)}(\text{lk}(v))$, the complex on the right hand side is combinatorially equivalent to the suspension of $\text{scp}_{\leq f(v)}(\text{lk}(v))$.

$$\tilde{H}_i(\text{scp}_{\leq f(v)}(\text{st}(v)) \cup \text{cone}(\text{scp}_{\leq f(v)}(\text{lk}(v)))) \cong \tilde{H}_i(\text{susp}(\text{scp}_{\leq f(v)}(\text{lk}(v))))$$

Another result stated in Hatcher [25, Section 2.1, Exercise 20, p. 132] relates the homology of a suspension $\text{susp}(X)$ to the homology of the original space X by $\tilde{H}_i(X) = \tilde{H}_{i+1}(\text{susp}(X))$. Furthermore we can use the fact that $\text{lk}(v)_{\leq f(v)}$ collapses to $\text{scp}_{\leq f(v)}(\text{lk}(v))$ and hence the two spaces have the same homology.

$$\tilde{H}_i(\text{susp}(\text{scp}_{\leq f(v)}(\text{lk}(v)))) \cong \tilde{H}_{i-1}(\text{scp}_{\leq f(v)}(\text{lk}(v))) \cong \tilde{H}_{i-1}(\text{lk}(v)_{\leq f(v)})$$

Summing up these isomorphisms, we obtain that v is critical in the sense of Brehm and Kühnel if and only if $\tilde{H}_{i-1}(\text{lk}(v)_{\leq f(v)}) \neq 0$ for some i , in other words, if and only if the space $\text{lk}(v)_{\leq f(v)}$ is not acyclic. Brehm and Kühnel also define the notion of an index of a critical point. A critical point is said to have *index* λ with *multiplicity* k_λ if k_λ is the rank of the homology group $H_i(|M|_{\leq f(v)}, |M|_{\leq f(v)} \setminus \{v\})$, which agrees with the rank of the reduced homology group $\tilde{H}_{i-1}(\text{lk}(v)_{\leq f(v)})$ by the above isomorphisms. In other words, the multiplicity k_λ of the index λ is the $(\lambda - 1)$ -th reduced Betti number of the lower link, which can be alternatively understood as the lower level link $\text{lk}(v)_{\leq f(v)}$ or the lower subcomplex $\text{scp}_{\leq f(v)}(\text{lk}(v))$ of the link, yielding the equalities $k_\lambda = \tilde{\beta}_{\lambda-1}(\text{lk}(v)_{\leq f(v)}) = \tilde{\beta}_{\lambda-1}(\text{scp}_{\leq f(v)}(\text{lk}(v)))$.

Comparison. For comparing the notion of criticality of Brehm and Kühnel with the notion from Definition 3.12, observe that local equivalence of f at a point v with one of the standard maps π_1^d or f_λ^d at the origin implies a PL homeomorphism between $\text{lk}(v)_{\leq f(v)}$ and $(S_1^{d-1})_{\pi_1^d \leq 0}$ or $(S_1^{d-1})_{f_\lambda^d \leq 0}$ respectively. We showed in the proof of Lemma 3.17 that $(S_1^{d-1})_{\pi_1^d \leq 0}$ is a ball and that $(S_1^{d-1})_{f_\lambda^d \leq 0}$ is homotopy equivalent to a $(\lambda - 1)$ -sphere, with the empty set being considered as (-1) -sphere. Reading off the Betti numbers from Table 3.1 and converting them into reduced Betti numbers allows us to determine the classification of the points according to Brehm and Kühnel. We can conclude

that regular points in our sense are not critical in the sense of Brehm and Kühnel, because the lower level link $\text{lk}(v)_{\leq f(v)}$ of a regular point v is a ball and balls are acyclic. The lower level link $\text{lk}(v)_{\leq f(v)}$ of non-degenerate critical point v of index λ in our sense is homotopy equivalent to a $(\lambda - 1)$ -sphere, which has the reduced Betti numbers $\tilde{\beta}_{\lambda-1}(S^{\lambda-1}) = 1$ and $\tilde{\beta}_n(S^{\lambda-1}) = 0$ for $n \neq \lambda - 1$. Therefore such points are critical points in the sense of Brehm and Kühnel that have index λ with multiplicity $k_\lambda = 1$ and all other indices with multiplicity 0.

The definition of Brehm and Kühnel can be interpreted as a classification of points in a combinatorial manifold with a map linear on cells with respect to an equivalence relation that considers two triples (M, x, f) and (M', x', f') as equivalent if and only if $\tilde{\beta}_n(\text{lk}_M(x)_{f \leq f(x)}) = \tilde{\beta}_n(\text{lk}_{M'}(x')_{f' \leq f'(x')})$ holds for all n . Note that replacing the reduced Betti numbers by the unreduced ones yields the same equivalence relation. By Characterisation 9. from Theorem 3.11 for local equivalence, the triples (M, x, f) and (M', x', f') are locally equivalent if and only if there is a PL homeomorphism from $\text{lk}_M(x)$ to $\text{lk}_{M'}(x')$ mapping upper level link to upper level link and lower level link to lower level link, which induces in particular a PL homeomorphism between $\text{lk}_M(x)_{f \leq f(x)}$ and $\text{lk}_{M'}(x')_{f' \leq f'(x')}$. Since homeomorphic spaces have the same Betti numbers, the equivalence relation defined by matching Betti numbers of the lower link is at least as coarse as the local equivalence relation. When we do not insist on general position it is easy to construct an example showing that matching Betti numbers of the lower link is even strictly coarser than local equivalence: Take the suspension $\text{susp}(S^{d-1})$ on some combinatorial $(d-1)$ -sphere S^{d-1} with $d \geq 2$ and define $f(a_1) = f(a_2) = f(u) = 0$ for the two apices a_1 and a_2 and for a single vertex u in S^{d-1} and define $f(w) > 0$ for all other vertices w in S^{d-1} . Then a_1 is not critical in the sense of Brehm and Kühnel because the lower level link consists of a single point and is acyclic, but it is degenerate critical in our sense because the lower level link is not homeomorphic to a $(d-1)$ -ball.

When we restrict our attention to maps in general position and combinatorial manifolds without boundary, it seems to be harder to design an example witnessing that the two equivalence relations differ. For example, if we wanted to construct a vertex in a combinatorial d -manifold as above that is not critical in the sense of Brehm and Kühnel but critical in our sense, any candidate (M, x, f) would have a homology $(d-1)$ -ball as lower level link $\text{lk}_M(x)_{f \leq f(x)}$ embedded in the $(d-1)$ -sphere $\text{lk}_M(x)$. This raises the question whether there are, in sufficiently large dimension, PL homology $(d-1)$ -balls different from the standard ball that embed piecewise linearly into some PL $(d-1)$ -sphere, because this might lead to an example as desired, provided that we can represent such an embedded homology $(d-1)$ -ball as the lower level set $S_{\leq 0}^{d-1}$ of some map linear on cells of some triangulation of the $(d-1)$ -sphere S^{d-1} such that the value 0 itself is not attained at vertices. Unfortunately it seems that this question has not been addressed yet in the literature, not even in the smooth or general topology categories.

So instead of an example relying on embedded homology balls, we could look for an example relying on a standard ball that is embedded in a non-standard way, that is, $\text{lk}_M(x)_{f \leq f(x)}$ is PL-homeomorphic to the lower half-sphere $(S_1^{d-1})_{\pi_1^d \leq 0}$ but no homeomorphism between them extends to a homeomorphism between $\text{lk}_M(x)$ and S_1^{d-1} mapping $\text{lk}_M(x)_{f \geq f(x)}$ to the upper half-sphere $(S_1^{d-1})_{\pi_1^d \geq 0}$. But if f is in general position, the boundary sphere of

$\text{lk}_M(x)_{f \leq f(x)}$ corresponding to the $(d-2)$ -sphere $(S_1^{d-1})_{\pi_1^d=0}$ under the homeomorphism is the level set $\text{lk}_M(x)_{f=f(x)}$. Recall from the facts cited in Subsection 1.2.12 that a $(d-2)$ -sphere embedded in a $(d-1)$ -sphere decomposes the $(d-1)$ -sphere into two uniquely determined parts, and if one of the parts is a PL-ball, both parts are PL-balls and the embedding is unknotted. In our example, the $(d-2)$ -sphere $\text{lk}_M(x)_{f=f(x)}$ separates the $(d-1)$ -sphere $\text{lk}_M(x)$ into $\text{lk}_M(x)_{f \leq f(x)}$ and $\text{lk}_M(x)_{f \geq f(x)}$ with the former part being a PL-ball implying that the latter part is a PL-ball as well. Using Fact 1.39 for extending homeomorphisms between boundary spheres of balls to homeomorphisms of the balls themselves, we can extend the homeomorphism between $\text{lk}_M(x)_{f=f(x)}$ and $(S_1^{d-1})_{\pi_1^d \geq 0}$ in two ways: On the one hand to a homeomorphism between $\text{lk}_M(x)_{f \leq f(x)}$ and $(S_1^{d-1})_{\pi_1^d \leq 0}$ and on the other hand to a homeomorphism between $\text{lk}_M(x)_{f \geq f(x)}$ and $(S_1^{d-1})_{\pi_1^d \geq 0}$. The union of these two homeomorphisms witnesses local equivalence of (M, x, f) with $(\mathbb{R}^d, 0, \pi_1^d)$, thus x is a regular point of f .

The above considerations show that the desired example for a point that is not critical in the sense of Brehm and Kühnel, but critical in our sense cannot be obtained by trying to embed a standard ball into the sphere in a non-standard way. But the idea of embedding a standard lower level link in a non-standard way should work in the cases where non-standard embeddings are known to exist. One such case is the existence of knotted spheres in co-dimension 2 [44, p. 52], which we use to sketch a construction of a critical point in dimension $d \geq 4$ that has index $d-2$ with multiplicity 1 and all other indices with multiplicity 0 as defined by Brehm and Kühnel, just as a non-degenerate critical point of index $d-2$ in our sense would, but is degenerate critical in that sense. Consider a non-trivially embedded thickened knot $S^{d-3} \times B^2$ in S^{d-1} , thought of as the lower level link embedded into the link of the point to be constructed. Since the thickened knot $S^{d-3} \times B^2$ deformation retracts onto its core knot S^{d-3} , it has the same homology as a $(d-3)$ -sphere. When we consider the map h_{d-2}^d as defined in Lemma 3.27 on the boundary sphere S_∞^{d-1} of the hypercube, we can verify that $(S_\infty^{d-1})_{h_{d-2}^d \leq 0}$ is exactly the thickened unknot $S_\infty^{d-3} \times [-1, 1]^2$. We assumed $S^{d-3} \times B^2$ to be a non-trivial knot in S^{d-1} , so there is no homeomorphism from S^{d-1} to S_∞^{d-1} taking $S^{d-3} \times B^2$ to the unknot $S_\infty^{d-3} \times [-1, 1]^2$. If we manage to represent $S^{d-3} \times B^2$ as lower level set $S_{f \leq 0}^{d-1}$ for a map f linear on cells of a triangulation of S^{d-1} not attaining the value 0 at vertices, we can extend this to a combinatorial d -manifold with a map f linear on cells, such that the triangulation of S^{d-1} is the link of a vertex and $S^{d-3} \times B^2$ the lower level link. Then we have obtained a degenerate critical vertex where the Betti numbers of the lower link agree with the Betti numbers of a $(d-3)$ -sphere as desired.

Definition of Banchoff

The last notion of criticality we review here is due to Banchoff's expositions in [4] and [5]. Banchoff defines in [4] an index for all vertices of a polytopal complex with respect to a height function in general position. Rephrased in our notation, the index $a(v, f)$ of v with respect to a function f is defined by $a(v, f) = \sum_{S \in M} (-1)^{\dim S} A(S, v, f)$, where $A(S, v, f)$ is an indicator function that is 1 if and only if v is a vertex of S and f attains its maximum on S at v and 0 oth-

erwise. In other words, $A(S, v, f) = 1$ holds if and only if $S \in \text{scp}_{\leq f(v)}(\text{st}(v)) \setminus \text{scp}_{\leq f(v)}(\text{lk}(v))$. Thus the index of v agrees with the following difference of Euler characteristics: $a(v, f) = \chi(\text{scp}_{\leq f(v)}(\text{st}(v))) - \chi(\text{scp}_{\leq f(v)}(\text{lk}(v)))$. Since $\text{scp}_{\leq f(v)}(\text{st}(v))$ is a cone with apex v and base $\text{scp}_{\leq f(v)}(\text{lk}(v))$, its Euler characteristic is 1 and we obtain $a(v, f) = 1 - \chi(\text{scp}_{\leq f(v)}(\text{lk}(v))) = 1 - \chi(\text{lk}(v)_{\leq f(v)})$.

We can derive an equivalence relation on triples (M, v, f) from Banchoff's definition by considering two triples (M, v, f) and (M', v', f') as equivalent if the index $a(v, f)$ of v in M with respect to f agrees with the index $a(v', f')$ of v' in M' with respect to f' . Compared to the classification of Brehm and Kühnel, Banchoff's classification is a coarsening: Brehm and Kühnel classify vertices by the Betti numbers of the lower link, whereas Banchoff considers only its Euler characteristic, that is the alternating sum of all Betti numbers. In particular, vertices v that are not critical with respect to f in the sense of Brehm and Kühnel have index $a(v, f) = 0$. It is easy to design examples of lower links whose Euler characteristics agree but whose individual Betti numbers differ: Define a piecewise linear map f on a 2-sphere that is negative in the interior two disjoint regions, namely in a strip around the equator and a ball around the north pole, zero on their boundaries, and positive elsewhere, for example by taking the derived of S_1^2 in \mathbb{R}^3 and choosing negative values for vertices on the equator and at the north pole and positive values for all other vertices of the derived. The disjoint union of a ball and an annulus has the non-zero Betti numbers $\beta_0 = 2$ and $\beta_1 = 1$, but its Euler characteristic 1 agrees with the Euler characteristic of a single ball. After completing the example to a combinatorial 3-manifold with PL map f where the example occurs as link $\text{lk}(v)$ with embedded lower level link $\text{lk}(v)_{\leq f(v)}$, we obtain a vertex v with index $a(v, f) = 0$ that is critical in the sense of Brehm and Kühnel having indices 0 and 1 with multiplicity 1 and all other indices with multiplicity 0. Using the fact that the Euler characteristic $\chi(X)$ can be expressed in terms of the reduced Betti numbers by $\chi(X) = 1 + \sum (-1)^n \tilde{\beta}_n(X)$, we see that the index $a(v, f)$ in the sense of Banchoff can also be computed as alternating sum of the multiplicities k_λ of the indices λ the vertex v has according to Brehm and Kühnel:

$$\begin{aligned} a(v, f) &= 1 - \chi(\text{lk}(v)_{\leq f(v)}) = - \sum (-1)^n \tilde{\beta}_n(\text{lk}(v)_{\leq f(v)}) \\ &= \sum (-1)^\lambda \tilde{\beta}_{\lambda-1}(\text{lk}(v)_{\leq f(v)}) = \sum (-1)^\lambda k_\lambda \end{aligned}$$

Recall that the lower level link of a regular point in our sense is a ball and has Euler characteristic 1, and that the lower link of critical points of index λ in our sense has the homology of a $(\lambda - 1)$ -sphere and hence Euler characteristic $1 - (-1)^{\lambda-1}$. This yields as index for regular points the value 0, and for critical points of index λ the value $(-1)^\lambda$. Banchoff does not explicitly define the notions of critical and regular vertices in [4]. But the paper contains a comparison of the smooth and PL situation taking the n -torus as example. In this example, the same correspondence as above is obtained: Vertices with index $a(v, f) = 0$ correspond to regular points in the smooth example, critical vertices of index λ in the smooth sense correspond to vertices with index $a(v, f) = (-1)^\lambda$. In another paper [5], Banchoff restricts the study to 2-dimensional closed manifolds embedded in \mathbb{R}^3 and defines the index in a different but equivalent fashion. The account illustrates the results by comparing them to smooth analogues, and it is

also observed that ordinary points where the level plane of the height function through the point cuts the star into two pieces have index 0. Bloch [8] suggests to adopt the condition that the index is non-zero as definition for critical points in the sense of Banchoff.

Chapter 4

More Characterisations for Regular Points

In the previous chapter, we defined regular points as those points where the map is locally equivalent to the linear coordinate map π_1^d at the origin. Since local equivalence can be characterised in many different fashions, we already obtained several characterisations for regular points. This chapter adds some more characterisations of regular points, at least if general position is assumed.

The original motivation for the studies in this thesis are questions regarding the properties of regular points, in particular properties that are characterising such points and properties expressing the absence of topological changes in the level sets. Chiang, Lenz, Lu, and Rote raise such kind of questions in a paper on the construction of contour trees [13]. These questions are collected and supplemented by references to related results and problems by Rote in [42]. In both papers general position for the map $f: |M| \rightarrow \mathbb{R}$ linear on cells of a combinatorial manifold M is assumed, so that critical behaviour occurs only at vertices, and a vertex v is called regular if all level sets $|M|_{=h}$ for h in a sufficiently small interval $[f(v) - \epsilon, f(v) + \epsilon]$ are homeomorphic [13, Definition 3]. For the three-dimensional case, it is shown that there even is a level preserving isotopy between the level sets of an interval range $[a, b]$ whose preimage contains only regular points, namely a continuous map $\Phi: |M|_{=b} \times [a, b] \rightarrow M$ such that for each $h \in [a, b]$ the restriction of Φ to $|M|_{=b} \times \{h\}$ yields a PL homeomorphism between $|M|_{=b}$ and $|M|_{=h}$ [13, Theorem 2]. However, the map Φ constructed in the proof is not piecewise linear when considered on its whole domain $|M|_{=b} \times [a, b]$, and one part of the proof does not carry over to higher dimensions, as Rote points out in [42].

In this chapter, we prove the existence of a piecewise linear f -level-preserving isotopy $\Phi: |M|_{=b} \times [a, b] \rightarrow M$ as above for arbitrary dimension, whenever f is in general position and the preimage $f^{-1}[a, b]$ contains only regular points in the sense of Definition 3.12. The result establishes a piecewise linear analogue to Fact 1.3 from smooth Morse theory and is presented in Theorem 4.20. Moreover, this property can be used to characterise regular vertices as stated in Theorem 4.19: For a map in general position, a vertex v is regular if and only if a piecewise linear f -level-preserving isotopy $\Phi: |M|_{=b} \times [a, b] \rightarrow M$ in some interval $[a, b]$ containing $f(v)$ in its interior exists. Even without the general position

assumption, regular points can be characterised according to Theorem 4.10 by the existence of an f -level-preserving isotopy in some neighbourhood of the point.

Most parts of the chapter are dedicated to the study of f -level-preserving isotopies leading to the results outlined above. But the first observations establish a link between the characterisation of regular points and the Schoenflies conjecture using the observation in Lemma 4.1 that the definition of regular points can be rephrased in terms of unknotted sphere or ball pairs. This yields additional characterisations of regular points for situations where the Schoenflies conjecture is already known to be true and conjectured candidates for a characterisation in situations where the Schoenflies conjecture is still open.

4.1 Unknotted Sphere Pairs

Characterisations of Regular Points Implied by the Different Characterisations of Local Equivalence. Let us first review the characterisations of regular points that follow from Definition 3.12 in connection with the alternative characterisations of local equivalence from Theorem 3.11. Each criterion for local equivalence from the list in Theorem 3.11 yields a criterion for $x \in |M|$ being a regular point of a map $f: |M| \rightarrow \mathbb{R}$ linear on cells of a polytopal complex M . The following list states these criteria, supplemented by a few remarks. Again all but the last two criteria also work for polyhedral complexes. In each criterion, the considered neighbourhood or link complex for the origin is a neighbourhood or link complex in \mathbb{R}^d with $d > 0$ for regular interior points and a neighbourhood or link complex in \mathbb{R}_+^d with $d > 1$ for regular boundary points; for the last two criteria, S^{d-1} can stand for any link complex for the origin, but is intended to represent some standard link complex such as the 1-norm or ∞ -norm unit sphere around the origin for interior points or the respective half-sphere for boundary points.

1. There is a PL homeomorphism between some cone neighbourhood of x and some cone neighbourhood of the origin mapping x to the origin such that $f - f(x)$ and the linear coordinate map π_1^d commute with the PL homeomorphism. Rephrasing this criterion gives the classical characterisation via local coordinate systems: The point x is regular if there is an x -centred PL local coordinate system (X_1, \dots, X_d) such that f is represented by $f(x) + X_1$ in these coordinates.
2. There are combinatorially equivalent cone neighbourhoods of x and of the origin such that x corresponds to 0, the map f is linear on cells of the cone neighbourhood, and the values of $f - f(x)$ and π_1^d agree at corresponding vertices.
3. There is a link complex for x contained in $\text{st}(x)$ and PL-homeomorphic to a link complex for the origin such that $f - f(x)$ and π_1^d commute with the PL homeomorphism.
4. There is a link complex for x contained in $\text{st}(x)$ and combinatorially equivalent to a link complex for the origin such that f is linear on cells of the link complex and $f - f(x)$ and π_1^d agree at corresponding vertices.

5. There is a link complex for x contained in $\text{st}(x)$ and PL-homeomorphic to a link complex for the origin such that $f - f(x)$ and π_1^d have the same sign at corresponding points.
6. There is a link complex for x contained in $\text{st}(x)$ and combinatorially equivalent to a link complex for the origin such that f is linear on cells of the link complex and $f - f(x)$ and π_1^d have the same sign at corresponding vertices.
7. Any pair of a link complex for x contained in $\text{st}(x)$ and a link complex for the origin is PL-homeomorphic in such a way that $f - f(x)$ and π_1^d have the same sign at corresponding points.
8. Any pair of a link complex for x contained in $\text{st}(x)$ and a link complex for the origin has combinatorially equivalent subdivisions such that $f - f(x)$ and π_1^d have the same sign at corresponding vertices.
9. There is a PL homeomorphism between $\text{lk}(x)$ and S^{d-1} mapping the upper level link $\text{lk}(x)_{\geq f(x)}$ to the upper half-sphere (for interior points, quarter-sphere for boundary points) $S_{\pi_1^d \geq 0}^{d-1}$ and the lower level link $\text{lk}(x)_{\leq f(x)}$ to the lower half-sphere (quarter-sphere) $S_{\pi_1^d \leq 0}^{d-1}$, that is, $f - f(x)$ and π_1^d have the same sign at corresponding points. This characterisation is the one that suits the purposes of this chapter best.
10. The link $\text{lk}(x)$ and S^d have combinatorially equivalent subdivisions such that $f - f(x)$ and π_1^d have the same sign at corresponding vertices.

Since the second to last characterisation is crucial for this chapter, let us elaborate on it a bit further. In a nutshell, it requires that the separation of the link into upper and lower level link has to be homeomorphic to the separation of a standard sphere into two half spheres by an equatorial hyperplane for an interior regular point, and that it is homeomorphic to the separation of a standard half-sphere into two quarter-spheres by an equatorial hyperplane for regular boundary points. Note that this requirement matches the suggested characterisation for regular interior vertices in 4-dimensional manifolds in [13] and [42], namely that the level set through the vertex should decompose the link into two 3-dimensional balls intersecting in their common boundary, because cone constructions as in Fact 1.39 show that the union of two balls intersecting in their boundaries is homeomorphic to the union of two opposite half-spheres of a standard sphere. Our results on local equivalence from the previous chapter show that it does not matter which link complex for the origin in \mathbb{R}^d or \mathbb{R}_+^d exactly we treat as the standard sphere or half-sphere S^{d-1} : For two link complexes L and L' of the origin, pseudo-radial projection between the subdivisions $L_{\pi_1^d \leq 0} \cup L_{\pi_1^d \geq 0}$ and $L'_{\pi_1^d \leq 0} \cup L'_{\pi_1^d \geq 0}$ yields a homeomorphism respecting the separation into upper and lower part.

Characterisation by Unknotted Sphere or Ball Pairs. A homeomorphism between $\text{lk}(x)$ and S^{d-1} witnessing regularity of x by mapping upper part to upper part and lower part to lower part maps in particular the separating level set $\text{lk}(x)_{=f(x)}$ to the equator $S_{\pi_1^d = 0}^{d-1}$ of the sphere or half-sphere. Hence for a regular point x , the pair $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is homeomorphic to the

pair $(S^{d-1}, S_{\pi_1^d=0}^{d-1})$ which in turn is homeomorphic to the standard $(d-1, d-2)$ -sphere pair for interior x or to the standard $(d-1, d-2)$ -ball pair for boundary points x , as one easily checks. This yields as necessary condition for x being regular that $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is an unknotted sphere or ball pair. For obtaining a necessary and sufficient condition, we only have to add a trivial supplementary assumption.

Lemma 4.1. *A point x in a polytopal complex M is a regular point of a map $f: |M| \rightarrow \mathbb{R}$ if and only if its link $\text{lk}(x)$ contains vertices u_- and u_+ with $f(u_-) < f(x)$ and $f(u_+) > f(x)$ and the pair $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is an unknotted $(d-1, d-2)$ -sphere or -ball pair with $d > 0$ for sphere pairs, which characterise regular interior points, and $d > 1$ for ball pairs, which characterise regular boundary points.*

Proof. Assume first that x is regular for showing necessity of the condition.

Necessity. The existence of the vertices u_- and u_+ is already guaranteed if we find arbitrary points y_- and y_+ in $|\text{lk}(x)|$ with $f(y_-) < f(x)$ and $f(y_+) > f(x)$ because the values of y_- and y_+ are obtained by interpolation of values of vertices of $\text{lk}(x)$. By regularity, there is a homeomorphism between $\text{lk}(x)$ and the 1-norm unit (half-)sphere S_1^{d-1} such that $f - f(x)$ and π_1^d have the same sign at corresponding points. The point $e_1 = (1, 0, \dots, 0)$ in S_1^{d-1} has positive π_1^d -value, thus its counterpart y_+ in $\text{lk}(x)$ fulfils $f(y_+) > f(x)$; analogously, the point $-e_1 = (-1, 0, \dots, 0)$ in S_1^{d-1} has negative π_1^d -value, thus its counterpart y_- in $\text{lk}(x)$ fulfils $f(y_-) < f(x)$.

For necessity, it remains to show that $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is an unknotted sphere or ball pair. Since the pair $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is homeomorphic to any pair $(S^{d-1}, S_{\pi_1^d=0}^{d-1})$ where S^{d-1} is a link complex of the origin, we are free to choose a link complex S^{d-1} in \mathbb{R}^d such that $(S^{d-1}, S_{\pi_1^d=0}^{d-1})$ is easily recognised as unknotted sphere pair and a link complex S^{d-1} in \mathbb{R}_+^d such that $(S^{d-1}, S_{\pi_1^d=0}^{d-1})$ is easily recognised as unknotted ball pair. In Subsection 1.2.12, we adopted the standard pairs suggested in [44], namely the standard $(d-1, d-2)$ -sphere pair $(\text{bd}([-1, 1]^d), \text{bd}([-1, 1]^{d-1}) \times \{0\})$ and the standard $(d-1, d-2)$ -ball pair $([-1, 1]^{d-1}, [-1, 1]^{d-2} \times \{0\})$. We sketch how to construct for a suitable choice of the link complex S^{d-1} for the origin a homeomorphism between $(S^{d-1}, S_{\pi_1^d=0}^{d-1})$ and the respective standard pair.

Interior Points. For regular interior points, we choose the ∞ -norm unit sphere $S_\infty^{d-1} = \text{bd}([-1, 1]^d)$ as link complex S^{d-1} for the origin in \mathbb{R}^d . With this choice, the pair $(S^{d-1}, S_{\pi_1^d=0}^{d-1})$ agrees with $(\text{bd}([-1, 1]^d), \{0\} \times \text{bd}([-1, 1]^{d-1}))$, which is almost the standard sphere pair. We only have to apply the linear transformation of \mathbb{R}^d that swaps the first and last coordinate to obtain a homeomorphism between $(\text{bd}([-1, 1]^d), \{0\} \times \text{bd}([-1, 1]^{d-1}))$ and the standard pair.

Boundary Points. For regular boundary points we choose for S^{d-1} the 1-norm unit half-sphere $S_1^{d-1} \cap \mathbb{R}_+^d$ as link complex for the origin in \mathbb{R}_+^d . Observe that $S_1^{d-1} \cap \mathbb{R}_+^d$ is the cone on $S_1^{d-2} \times \{0\}$ with apex at $e_d = (0, \dots, 0, 1)$ and the level set $(S_1^{d-1} \cap \mathbb{R}_+^d)_{\pi_1^d=0}$ is the cone on $\{0\} \times S_1^{d-3} \times \{0\}$ with apex e_d . Taking as apex the origin instead of e_d , we obtain the combinatorially equivalent pair $(B_1^{d-1} \times \{0\}, \{0\} \times B_1^{d-2} \times \{0\})$, where $B_1^k = 0S_1^{k-1}$ is the 1-norm unit ball around the origin in \mathbb{R}^k . We can omit the last coordinate, which is constantly

zero. Swapping the first and the $(d-1)$ -th coordinate gives the homeomorphic ball pair $(B_1^{d-1}, B_1^{d-2} \times \{0\})$.

The boundaries of the balls form the 1-norm $(d-2, d-3)$ -sphere pair $(S_1^{d-2}, S_1^{d-3} \times \{0\})$, which is homeomorphic to the ∞ -norm $(d-2, d-3)$ -sphere pair $(S_\infty^{d-2}, S_\infty^{d-3} \times \{0\})$ by pseudo-radial projection. Extending this boundary homeomorphism to the balls by the cone construction as in Fact 1.39 yields a homeomorphism of the 1-norm $(d-1, d-2)$ -ball pair $(B_1^{d-1}, B_1^{d-2} \times \{0\})$ to the ∞ -norm $(d-1, d-2)$ -ball pair $(B_\infty^{d-1}, B_\infty^{d-2} \times \{0\})$ which agrees with the standard $(d-1, d-2)$ -ball pair $([-1, 1]^{d-1}, [-1, 1]^{d-2} \times \{0\})$. This completes the proof for the necessity of the criterion.

Sufficiency. Now we prove that the criterion is also sufficient.

Interior Points. Assume first that the pair $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is an unknotted $(d-1, d-2)$ -sphere pair. We saw in the first part of the proof that the standard sphere pair is homeomorphic to the pair $(S_\infty^{d-1}, (S_\infty^{d-1})_{\pi_1^d=0})$. Thus there is a PL homeomorphism $\phi: |\text{lk}(x)| \rightarrow |S_\infty^{d-1}|$ with $\phi(|\text{lk}(x)|_{=f(x)}) = |S_\infty^{d-1}|_{\pi_1^d=0}$. It suffices to show that either ϕ itself or its composition with the reflection of \mathbb{R}^d at the hyperplane $\pi_1^d = 0$, which swaps the upper half-sphere $(S_\infty^{d-1})_{\pi_1^d \geq 0}$ and the lower half-sphere $(S_\infty^{d-1})_{\pi_1^d \leq 0}$, maps the upper level link $\text{lk}(x)_{\geq f(x)}$ to the upper half-sphere and the lower level link $\text{lk}(x)_{\leq f(x)}$ to the lower half-sphere.

Applying the results from [31] cited in Subsection 1.2.12 concerning the separation of a sphere by a sphere of co-dimension 1 yields: The $(d-2)$ -sphere $|\text{lk}(x)|_{=f(x)}$ separates the $(d-1)$ -sphere $|\text{lk}(x)|$ in a unique non-trivial way into two parts B_1 and B_2 with $B_1 \cap B_2 = |\text{lk}(x)|_{=f(x)}$ and $B_1 \cup B_2 = |\text{lk}(x)|$. The existence of u_- and u_+ guarantees that the decomposition given by the choices $B_1 = |\text{lk}(x)|_{\geq f(x)}$ and $B_2 = |\text{lk}(x)|_{\leq f(x)}$ not only fulfils $B_1 \cap B_2 = |\text{lk}(x)|_{=f(x)}$ and $B_1 \cup B_2 = |\text{lk}(x)|$ but also differs from the trivial decomposition into $|\text{lk}(x)|$ and $|\text{lk}(x)|_{=f(x)}$. On the other hand, since ϕ is a homeomorphism, the choices $B_1 = \phi^{-1}(|S_\infty^{d-1}|_{\pi_1^d \geq 0})$ and $B_2 = (|S_\infty^{d-1}|_{\pi_1^d \leq 0})$ yield a non-trivial decomposition with these properties as well. Therefore uniqueness of the decomposition implies that either the equalities $\phi(|\text{lk}(x)|_{\geq f(x)}) = |S_\infty^{d-1}|_{\pi_1^d \geq 0}$ and $\phi(|\text{lk}(x)|_{\leq f(x)}) = |S_\infty^{d-1}|_{\pi_1^d \leq 0}$ or the equalities $\phi(|\text{lk}(x)|_{\geq f(x)}) = |S_\infty^{d-1}|_{\pi_1^d \leq 0}$ and $\phi(|\text{lk}(x)|_{\leq f(x)}) = |S_\infty^{d-1}|_{\pi_1^d \geq 0}$ hold. In the former case, ϕ itself is the desired homeomorphism that maps the upper level link to the upper half-sphere and the lower level link to the lower half-sphere, in the latter case the composition of ϕ and the reflection at the hyperplane $\pi_1^d = 0$ has this property.

Boundary Points. For completing the proof, it remains to consider the case of an unknotted $(d-1, d-2)$ -ball pair $(\text{lk}(x), \text{lk}(x)_{=f(x)})$. From the first part of the proof we use the fact that the standard ball pair is homeomorphic to the pair $(S_1^{d-1} \cap \mathbb{R}_+^d, (S_1^{d-1} \cap \mathbb{R}_+^d)_{\pi_1^d=0})$. Therefore a PL homeomorphism $\phi: |\text{lk}(x)| \rightarrow |S_1^{d-1} \cap \mathbb{R}_+^d|$ with $\phi(|\text{lk}(x)|_{=f(x)}) = |S_1^{d-1} \cap \mathbb{R}_+^d|_{\pi_1^d=0}$ exists. Our aim is again to prove that either ϕ itself or its composition with the reflection at the hyperplane $\pi_1^d = 0$ maps the upper level link to the upper quarter-sphere and the lower level link to the lower quarter-sphere.

We transfer the problem from the realm of ball pairs to the realm of sphere pairs by taking the double of all involved objects, so that we can apply the separation result for spheres again. The double of a manifold N is the quotient space obtained from N and a copy N' of N by identifying each boundary point p in N with its copy p' in N' . Hence for a combinatorial manifold N , the double

of N can be written as the union of N and a combinatorially equivalent complex N' such that the intersection of N and N' is exactly the boundary of N and at the same time the boundary of N' and the combinatorial equivalence restricts to the identity on the common intersection.

When we take the double $D = \text{lk}(x) \cup (\text{lk}(x))'$ of $\text{lk}(x)$, we can extend f to D by letting $f(p') = f(p)$ for each point $p \in \text{lk}(x)$ and its copy $p' \in (\text{lk}(x))'$. The double of $S_1^{d-1} \cap \mathbb{R}_+^d$ whose boundary coincides with the intersection with the hyperplane $X_d = 0$ can be identified with the union of the complex and its mirror image under reflection at the hyperplane. This union of the half-sphere $S_1^{d-1} \cap \mathbb{R}_+^d$ and its mirror image is exactly the whole sphere S_1^{d-1} and the π_1^d -values of a point and its mirror image agree. We can also extend the homeomorphism ϕ to a homeomorphism between D and S_1^{d-1} in a natural way by letting $\phi(p') = (\phi(p))'$ for a copied point $p' \in (\text{lk}(x))'$ of a point $p \in \text{lk}(x)$ where $(\phi(p))'$ is the mirror image of the point $\phi(p) \in S_1^{d-1} \cap \mathbb{R}_+^d$. The following properties of this construction can be easily checked: The level sets $D_{\leq f(x)}$, $D_{\geq f(x)}$, and $D_{=f(x)}$ are doubles of $\text{lk}(x)_{\leq f(x)}$, $\text{lk}(x)_{\geq f(x)}$, and $\text{lk}(x)_{=f(x)}$ respectively. Similarly, the upper and lower half spheres of S_1^{d-1} are the doubles of the upper and lower quarter spheres of $S_1^{d-1} \cap \mathbb{R}_+^d$, and their common intersection $(S_1^{d-1})_{\pi_1^d=0}$ is the double of $(S_1^{d-1} \cap \mathbb{R}_+^d)_{\pi_1^d=0}$. Furthermore the extended homeomorphism ϕ maps $D_{=f(x)}$ to $(S_1^{d-1})_{\pi_1^d=0}$.

These observations yield that the pair $(D, D_{=f(x)})$ is homeomorphic by the extended homeomorphism ϕ to the pair $(S_1^{d-1}, (S_1^{d-1})_{\pi_1^d=0})$, which is obviously homeomorphic to the standard sphere pair by pseudo-radial projection and a coordinate swap. Hence $(D, D_{=f(x)})$ is an unknotted sphere pair and $|D|_{>f(x)}$ and $|D|_{<f(x)}$ are both non-empty, because their subsets $|\text{lk}(x)|_{>f(x)}$ and $|\text{lk}(x)|_{<f(x)}$ already are by assumption. Using again the results from Subsection 1.2.12, the same arguments as above show that either the extended homeomorphism ϕ or its composition with the reflection at the hyperplane $\pi_1^d = 0$ maps $D_{\geq f(x)}$ to $(S_1^{d-1})_{\pi_1^d \geq 0}$ and $D_{\leq f(x)}$ to $(S_1^{d-1})_{\pi_1^d \leq 0}$. But then the original homeomorphism ϕ , which is the restriction of the extended homeomorphism ϕ to $\text{lk}(x)$ on the one hand and $S_1^{d-1} \cap \mathbb{R}_+^d$ on the other hand, or its composition with the reflection necessarily maps $\text{lk}(x)_{\geq f(x)}$ to $(S_1^{d-1} \cap \mathbb{R}_+^d)_{\pi_1^d \geq 0}$ and $\text{lk}(x)_{\leq f(x)}$ to $(S_1^{d-1} \cap \mathbb{R}_+^d)_{\pi_1^d \leq 0}$ as desired. \square

Corollary 4.2. *If a map $f: |M| \rightarrow \mathbb{R}$ linear on cells of a polytopal complex M is in general position, then a point $x \in |M|$ is regular, if and only if $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is an unknotted $(d-1, d-2)$ -sphere or -ball pair with $d > 1$ or $\text{lk}(x)$ is a 0-sphere such that one of its vertices has an f -value less than $f(x)$ and the other vertex an f -value greater than $f(x)$.*

Proof. We only have to check that the condition that $\text{lk}(x)$ contains vertices u_- and u_+ with $f(u_-) < f(x)$ and $f(u_+) > f(x)$ from the previous lemma is automatically fulfilled if $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is an unknotted $(d-1, d-2)$ -sphere or -ball pair with $d > 1$ and f is in general position.

When x is a vertex, no vertex of $\text{lk}(x)$ has the value $f(x)$ by general position. Since $\text{lk}(x)_{=f(x)}$ is a $(d-2)$ -sphere or -ball with $d > 1$, it is non-empty. But for $\text{lk}(x)_{=f(x)}$ being non-empty, in the absence of vertices with value $f(x)$, it is necessary that $\text{lk}(x)$ contains vertices u_- and u_+ with $f(u_-) < f(x)$ and $f(u_+) > f(x)$.

When x is not a vertex, consider the cell X containing x in its interior. Then the value $f(x)$ is attained in the interior of X , namely at x . By Lemma 2.21, the map f is either constant on X or the cell X contains vertices u_- and u_+ with $f(u_-) < f(x)$ and $f(u_+) > f(x)$. But the former alternative is excluded because X contains at least two vertices and their f -values differ by the general position assumption. Since the vertices of X also belong to the vertices of $\text{lk}(x)$, the vertices u_- and u_+ of the latter alternative satisfy the requirements. \square

The Influence of the Schoenflies Conjecture. When M is a combinatorial manifold and x an interior point of $|M|$, then $\text{lk}(x)$ is a sphere. For checking whether x is regular based on the characterisation from Lemma 4.1, disregarding the simple test on vertices u_+ and u_- with values greater than and less than $f(x)$ respectively, it remains to check whether $\text{lk}(x)_{=f(x)}$ is a sphere and is unknotted in $\text{lk}(x)$. If $\text{lk}(x)_{=f(x)}$ is known to be a sphere, the remaining question leads to the Schoenflies problem. Recall from Subsection 1.2.12 that the d -dimensional Schoenflies conjecture postulates that any embedding of S^{d-1} in S^d is unknotted. For the piecewise linear case, the conjecture is true for $d \leq 3$. For higher dimensions the problem is still open, but the conjecture is true for all dimensions if and only if it holds in dimension $d = 4$. If we add the assumption that the embedding of S^{d-1} in S^d is locally flat, then the conjecture is known to be true in all dimensions except $d = 4$, where even this weaker version is still unsolved. The current status of the conjecture allows to draw the following conclusions from Lemma 4.1 and Corollary 4.2:

Lemma 4.3. *Let M be a combinatorial d -manifold with interior point $x \in |M|$ and a map $f: |M| \rightarrow \mathbb{R}$ linear on cells of M . Assume either that $\text{lk}(x)$ contains vertices u_- and u_+ with $f(u_-) < f(x)$ and $f(u_+) > f(x)$ or that $d > 1$ and f is in general position.*

1. *If the dimension d is smaller than or equal to 4, then x is regular if and only if $\text{lk}(x)_{=f(x)}$ is PL-homeomorphic to the $(d-2)$ -sphere. If the Schoenflies conjecture is true in all dimensions, the restriction $d \leq 4$ can be omitted.*
2. *If the dimension d is greater than or equal to 6 (or $d \leq 4$) and for all vertices u of $\text{lk}(x)$ the value $f(u)$ differs from $f(x)$, then x is regular if and only if $\text{lk}(x)_{=f(x)}$ is PL-homeomorphic to the $(d-2)$ -sphere. If the Schoenflies conjecture is true for locally flat embeddings of S^3 into S^4 , the claim is also true for $d = 5$.*
3. *If we assume general position for f , $d > 1$, and $d \neq 5$, then x is regular if and only if x is not a vertex or $\text{lk}(x)_{=f(x)}$ is PL-homeomorphic to the $(d-2)$ -sphere. If the Schoenflies conjecture is true for locally flat embeddings of S^3 into S^4 , the claim is also true for $d = 5$.*

Proof. All three claims listed are just reformulations of the characterisation in Lemma 4.1 where the requirement that the pair $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is unknotted is replaced by using or assuming the truth of the Schoenflies conjecture. The existence of the vertices u_- and u_+ is explicitly assumed or follows from general position as in Corollary 4.2. The assumption that M is a d -manifold and that x is an interior point implies that $\text{lk}(x)$ is a $(d-1)$ -sphere. It is also assumed

in all claims that $\text{lk}(x)_{=f(x)}$ is a $(d-2)$ -sphere, so that $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is a $(d-1, d-2)$ -sphere pair. The truth of the Schoenflies conjecture for $(d-1, d-2)$ -sphere pairs with $d \leq 4$ or assuming its truth for higher dimension asserts that the pair is unknotted and yields the first claim. The second claim relies on the Schoenflies conjecture for locally flat sphere pairs. Although the claim also holds for $d \leq 4$, this case is not interesting because the stronger general version of the Schoenflies conjecture provides the simpler characterisation stated in the first claim. Recall from Corollary 2.40 that for any combinatorial manifold N , the level set $N_{=h}$ is locally flat in N if h is not attained at vertices. Hence, if $f(u) \neq f(x)$ for all vertices of $\text{lk}(x)$, then $\text{lk}(x)_{=f(x)}$ is locally flat in $\text{lk}(x)$. Since the Schoenflies conjecture for locally flat $(d-1, d-2)$ -sphere pairs is true for all $d \neq 5$ and open for $d = 5$, we obtain the second claim. For the last claim, the case of non-vertices is covered by the consequence of Lemma 3.22 that non-vertices are regular for maps in general position. For vertices, general position ensures the assumption made in the second claim that $f(u) \neq f(x)$ for all vertices of $\text{lk}(x)$ is fulfilled. Thus the truth of the Schoenflies conjecture for locally flat $(d-1, d-2)$ -sphere pairs with $d \neq 5$ or assuming it for $d = 5$ gives the last claim. \square

4.2 Level Preserving Isotopies

Now we turn to the relation between regular points and level preserving isotopies. For the study we first define the notion of an f -level-preserving isotopy and introduce some additional technical terms involving such isotopies. Then we derive a characterisation of critical points by local f -level-preserving isotopies. Finally we construct for maps in general position f -level-preserving isotopies across vertices, so that we obtain an isotopy on the whole manifold for any interval without critical values.

4.2.1 Definition and Basic Properties

Definition 4.4 (f -level-preserving PL isotopy). Let M be a polyhedral complex and $f: |M| \rightarrow \mathbb{R}$ linear on cells. An f -level-preserving PL isotopy of level sets between $M_{=a}$ and $M_{=b}$ is a PL homeomorphism $\phi: |M|_{=h} \times [a, b] \rightarrow |M|_{\in [a, b]}$ with $h \in [a, b]$ such that for any $t \in [a, b]$ the restriction of ϕ to $|M|_{=h} \times \{t\}$ is a PL homeomorphism between $|M|_{=h} \times \{t\}$ and $|M|_{=t}$. If we denote by $\phi_t: |M|_{=h} \rightarrow |M|_{=t}$ the map defined by $\phi_t(x) = \phi(x, t)$ for $x \in |M|_{=h}$, then the last condition can be rephrased by saying that for all $t \in [a, b]$ the map ϕ_t is a PL homeomorphism between $|M|_{=h}$ and $|M|_{=t}$. We sometimes call an f -level-preserving PL isotopy $\phi: |M|_{=h} \times [a, b]$ a PL isotopy on M , if the exact boundaries of the interval $[a, b]$ are irrelevant or implied from the context. Similarly, we call ϕ a PL isotopy between a and b , if M is clear from the context.

A *local f -level-preserving PL isotopy* in a neighbourhood of $x \in |M|$ is an f -level-preserving PL isotopy $\phi: |N|_{=h} \times [a, b] \rightarrow |N|_{\in [a, b]}$ where the domain $|N|$ of N is a neighbourhood of x in $|M|$ and $[a, b]$ is a neighbourhood of $f(x)$ in \mathbb{R} , that is $f(x) \in (a, b)$. Note that for polytopal M , we can always find a subdivision of M such that N is a subcomplex of that subdivision [44, Addendum 2.12, p. 16].

If K is a subcomplex of M and the image of the restriction of an f -level-preserving PL isotopy $\phi: |M|_{=h} \times [a, b] \rightarrow |M|_{\in [a, b]}$ to the subcomplex $|K|_{=h} \times$

$[a, b]$ is the level set $|K|_{\in[a, b]}$, then the restricted map $\phi|_{|K|_{=h} \times [a, b]}: |K|_{=h} \times [a, b] \rightarrow |K|_{\in[a, b]}$ is an f -level-preserving PL isotopy on K between $K_{=a}$ and $K_{=b}$. We abbreviate $\phi|_{|K|_{=h} \times [a, b]}$ by $\phi|_K$ and call it the restriction of ϕ to K .

A PL isotopy between complexes X and Y is often defined as a level-preserving embedding $\Phi: X \times I \rightarrow Y \times I$, where level-preserving means that Φ leaves the I -part unchanged, in other words, $\Phi(x, t) = (\phi(x, t), t)$ for a suitable function $\phi: X \times I \rightarrow Y$ [44, Definitions 3.21, p. 37]. Hudson [26, p. 128] remarks that the condition that Φ is level-preserving is equivalent to the condition on ϕ , that for each t the map $\phi_t: X \rightarrow Y$ is an embedding, where ϕ_t is defined by $\phi_t(x) = \phi(x, t)$. Hence, in our definition of an f -level-preserving PL isotopy, the usual notion of being level-preserving is implied by the condition that the restriction of ϕ to $|M|_{=h} \times \{t\}$ is a PL homeomorphism. The property of being f -level preserving is an additional requirement independent of the property of being level-preserving in that sense.

We start the study of f -level preserving isotopies with some simple observations that are inherent to the concept of an isotopy.

Lemma 4.5. *If $\phi: |M|_{=h} \times [a, b] \rightarrow |M|_{\in[a, b]}$ is an f -level-preserving PL isotopy and h' a value in the interval $[a, b]$, then there is an f -level-preserving PL isotopy $\psi: |M|_{=h'} \times [a, b] \rightarrow |M|_{\in[a, b]}$ that is the identity at level h' , meaning that $\psi_{h'} = \text{id}_{|M|_{=h'}}$ holds.*

More generally, if a self-homeomorphism χ of $|M|_{=h'}$ is given, there is an f -level-preserving PL isotopy $\psi: |M|_{=h'} \times [a, b] \rightarrow |M|_{\in[a, b]}$ that restricts to χ at level h' , namely $\psi_{h'} = \chi$.

Proof. The first claim follows from the second by choosing $\chi = \text{id}_{|M|_{=h'}}$. So let us prove the second claim. Define the map ψ by the assignment $\psi(x, t) = \phi(\phi_{h'}^{-1}(\chi(x)), t)$. Since χ and $\phi_{h'}^{-1}$ are PL homeomorphisms, the assignment $(x, t) \mapsto (\phi_{h'}^{-1}(\chi(x)), t)$ defines a PL homeomorphism from $|M|_{=h'} \times [a, b]$ to $|M|_{=h} \times [a, b]$. Then ψ is a PL homeomorphism because it is the composition of this PL homeomorphism and the PL homeomorphisms ϕ . Furthermore $\psi_t = \phi_t \circ \phi_{h'}^{-1} \circ \chi$ is a PL-homeomorphism between $|M|_{=h'}$ and $|M|_{=t}$. Thus ψ is an f -level-preserving PL isotopy and for $t = h'$ we obtain $\psi_{h'} = \phi_{h'} \circ \phi_{h'}^{-1} \circ \chi = \chi$ as desired. \square

Lemma 4.6. *If there is an f -level-preserving PL isotopy between $M_{=a}$ and $M_{=b}$ and there is an f -level-preserving PL isotopy between $M_{=b}$ and $M_{=c}$ for $a \leq b \leq c$, then there is an f -level-preserving PL isotopy between $M_{=a}$ and $M_{=c}$.*

Proof. By the previous lemma, we can assume that both PL isotopies are the identity at level b : $\phi: |M|_{=b} \times [a, b] \rightarrow |M|_{\in[a, b]}$ with $\phi_b = \text{id}_{|M|_{=b}}$ and $\psi: |M|_{=b} \times [b, c] \rightarrow |M|_{\in[b, c]}$ with $\psi_b = \text{id}_{|M|_{=b}}$. Then ϕ and ψ agree on their common domain $|M|_{=b} \times \{b\}$ and we can combine the maps ϕ and ψ to a single map $\chi: |M|_{=b} \times [a, c] \rightarrow |M|_{\in[a, c]}$ by letting $\chi = \phi$ for arguments in $|M|_{=b} \times [a, b]$ and $\chi = \psi$ for arguments in $|M|_{=b} \times [b, c]$. Since the images of ϕ and ψ intersect only in the image $|M|_{=b}$ of their common domain $|M|_{=b} \times \{b\}$, the map χ is a PL-homeomorphism. Furthermore, χ_t agrees with ϕ_t for $t \in [a, b]$ and with ψ_t for $t \in [b, c]$. Therefore, χ is an f -level-preserving PL isotopy. \square

Lemma 4.7. *The restriction of an f -level-preserving PL isotopy to some subinterval is again an f -level-preserving PL isotopy. More precisely, if $\phi: |M|_{=h} \times [a, b] \rightarrow |M|_{\in[a, b]}$ is an f -level-preserving PL isotopy for the interval $[a, b]$ and $[a', b'] \subseteq [a, b]$ a subinterval containing h , then the restriction of ϕ to $|M|_{=h} \times [a', b']$ is an f -level-preserving PL isotopy between $M_{=a'}$ and $M_{=b'}$.*

Proof. Since ϕ is f -level-preserving, the image of $|M|_{=h} \times [a', b']$ under ϕ is the level set complex $|M|_{\in[a', b']}$. Thus the restriction $\psi: |M|_{=h} \times [a', b'] \rightarrow |M|_{\in[a', b]}$ of the PL homeomorphism ϕ to $|M|_{=h} \times [a', b']$ is a PL homeomorphism as well. This homeomorphism is f -level-preserving because $\psi_t = \phi_t$ holds for each $t \in [a', b']$ and ϕ is f -level-preserving. \square

4.2.2 Characterising Regular Points by Local Level Preserving Isotopies

Now assume that M is a combinatorial manifold, that an f -level-preserving PL isotopy $\phi: |M|_{=h} \times [a, b] \rightarrow |M|_{\in[a, b]}$ exists, and that $x \in |M|$ has a value $f(x)$ in the open interval (a, b) . We want to show that x is regular. Actually, it is enough that an f -level-preserving PL isotopy in some neighbourhood of x exists.

Lemma 4.8. *Let M be a d -dimensional combinatorial manifold and $f: |M| \rightarrow \mathbb{R}$ linear on cells. Assume that in some neighbourhood of $x \in |M|$, a local f -level-preserving PL isotopy exists. Then x is a regular point.*

Proof. Let N be a neighbourhood of x in $|M|$ with a local f -level-preserving PL isotopy $\phi: |N|_{=h} \times [a, b] \rightarrow |N|_{\in[a, b]}$ for some interval $[a, b]$ containing $f(x)$ in its interior. For simplifying the proof, we first argue that we can assume $f(x) = 0$ without loss of generality. The general case can be reduced to this situation by shifting f , the interval $[a, b]$, and the level h . The local equivalence class of (M, x, f) is not affected by shifting the map f , thus x is regular with respect to f if and only if x is regular with respect to $f - f(x)$. The f -level-preserving PL isotopy ϕ for the interval $[a, b]$ induces an $(f - f(x))$ -level-preserving isotopy ϕ' for the shifted interval $[a - f(x), b - f(x)]$ in the obvious way: The assignment $\phi'(p, t) = \phi(p, t + f(x))$ defines an $(f - f(x))$ -level-preserving isotopy $\phi': |N|_{f-f(x)=h-f(x)} \times [a - f(x), b - f(x)] \rightarrow |N|_{f-f(x) \in [a - f(x), b - f(x)]}$. Note that the domains of the involved level sets in $|N|$, namely $|N|_{f-f(x)=h-f(x)} = |N|_{f=h}$ and $|N|_{f-f(x) \in [a - f(x), b - f(x)]} = |N|_{f \in [a, b]}$, remain unchanged. One easily checks that ϕ' is a homeomorphism if and only if ϕ is and that $\phi'_{t-f(x)}$ agrees with ϕ_t . Hence ϕ' is a local $(f - f(x))$ -level-preserving isotopy if and only if ϕ is a local f -level-preserving isotopy.

So let us assume $f(x) = 0$. For treating interior and boundary points simultaneously, we use again the term $\mathbb{R}_{(+)}^k$ to denote alternatively \mathbb{R}^k for interior points or \mathbb{R}_+^k for boundary points. By isotopy $|N|_{=f(x)}$ is homeomorphic to any other level set $|N|_{=t}$ for $t \in [a, b]$, and we can choose t in such a way that $|N|_{=t}$ does not contain vertices of N . Recall from Corollary 2.40 that such a level set is a $(d - 1)$ -dimensional manifold, hence $|N|_{=t}$ is a $(d - 1)$ -manifold for each $t \in [a, b]$ by isotopy. Let $y = (x', f(x))$ be the point in $|N|_{=h} \times [a, b]$ corresponding to x under the isotopy, that is, the point fulfilling $\phi(y) = \phi(x', f(x)) = x$. Since $|N|_{=h}$ is a manifold, there is a neighbourhood U' of x' in $|N|_{=h}$ with a coordinate chart $\chi': |U'| \rightarrow |V| \subseteq \mathbb{R}_{(+)}^{d-1}$ mapping x' to the origin and $|U'|$ to

some neighbourhood $|V|$ of the origin that witnesses that fact. This induces a coordinate chart $\chi: |U'| \times [a, b] \rightarrow [a, b] \times |V| \subseteq \mathbb{R}_{(+)}^d$ defined by the assignment $\chi(p, t) = (t, \chi'(p))$ for $p \in |U'|$ and $t \in [a, b]$ which maps y to the origin, because we assumed $f(x) = 0$, and the neighbourhood $U = U' \times [a, b]$ of y to the neighbourhood $\chi(|U|) = [a, b] \times |V|$ of the origin in $\mathbb{R}_{(+)}^d$. Since ϕ is f -level-preserving, $f \circ \phi \circ \chi^{-1}$ agrees on $\chi(|U|)$ with the coordinate map π_1^d . Therefore some neighbourhood of x in $|M|$, namely $\phi(|U|)$, is PL-homeomorphic to the neighbourhood $\chi(|U|)$ of the origin in $\mathbb{R}_{(+)}^d$ via $\chi \circ \phi^{-1}$ such that f corresponds to the coordinate function π_1^d . In other words, (M, x, f) is locally equivalent to $(\mathbb{R}_{(+)}^d, 0, \pi_1^d)$, showing that x is regular. \square

The criterion for regular points in the previous lemma is not only sufficient but also necessary. For establishing this, we assume that x is regular, and we want to construct an f -level-preserving PL isotopy. The definition of regularity gives such an isotopy in some neighbourhood of x almost immediately. But a construction that yields an isotopy for the special neighbourhood $\text{st}(x)$ or for the whole manifold M is more involved, and we prove these stronger conclusions only for vertices under suitable general position assumptions on f .

Lemma 4.9. *If $x \in |M|$ is regular with respect to a map $f: |M| \rightarrow \mathbb{R}$ linear on cells of a combinatorial manifold M , then there is a local f -level-preserving isotopy in some neighbourhood of x .*

Proof. From the previous proof, we adopt the observation that we can assume $f(x) = 0$ without loss of generality and the notation $\mathbb{R}_{(+)}^d$ for \mathbb{R}^d if x is a regular interior point or for \mathbb{R}_+^d if x is a regular boundary point. Since x is regular, (M, x, f) is locally equivalent to $(\mathbb{R}_{(+)}^d, 0, \pi_1^d)$. This means that there are neighbourhoods N_x of x in $|M|$ and N_0 of the origin in $\mathbb{R}_{(+)}^d$ and a PL homeomorphism $\psi: |N_0| \rightarrow |N_x|$ such that f corresponds to the coordinate function π_1^d via ψ . Choose ϵ small enough so that the ∞ -norm ball B of radius ϵ in $\mathbb{R}_{(+)}^d$ around the origin is contained in $|N_0|$. This ball is of the form $B = [-\epsilon, \epsilon] \times B'$ with B' being the ∞ -norm ball of radius ϵ around the origin in $\mathbb{R}_{(+)}^{d-1}$. We show that ψ induces an f -level-preserving PL isotopy in the neighbourhood $|N| = \psi(|B|)$ of x between $|N|_{=-\epsilon}$ and $|N|_{=\epsilon}$. Note that since f and π_1^d correspond via ψ , the level set $|N|_{\in [-\epsilon, \epsilon]}$ agrees with $|N|$ and the restriction of ψ to $\{0\} \times |B'|$ yields a PL homeomorphism $\chi: |B'| \rightarrow |N|_{=0}$, so that for any $p \in |N|_{=0}$ and $t \in [-\epsilon, \epsilon]$, the pair $(t, \chi^{-1}(p))$ describes an element of $[-\epsilon, \epsilon] \times B' = B$. In fact, since χ^{-1} and $\text{id}_{[-\epsilon, \epsilon]}$ are PL homeomorphisms, the assignment $(p, t) \mapsto (t, \chi^{-1}(p))$ defines a PL homeomorphism from $|N|_{=0} \times [-\epsilon, \epsilon]$ to $|B|$. The isotopy $\phi: |N|_{=0} \times [-\epsilon, \epsilon] \rightarrow |N|_{\in [-\epsilon, \epsilon]}$ is defined by composing this homeomorphism with ψ , which maps $|B|$ to $|N| = |N|_{\in [-\epsilon, \epsilon]}$, in other words, ϕ is given by the assignment $\phi(p, t) = \psi(t, \chi^{-1}(p))$. For each $t \in [-\epsilon, \epsilon]$, the restricted map ϕ_t is a composition of a homeomorphism mapping $|N|_{=0}$ to $\{t\} \times |B'|$ and a restriction of ψ mapping $\{t\} \times |B'| = |B|_{\pi_1^d=t}$ homeomorphically to $|N|_{=t}$ because f and π_1^d commute with ψ . Thus, ϕ fulfils all requirements for being a local f -level-preserving PL isotopy in a neighbourhood of x . \square

We obtain the following alternative characterisation of regular points:

Theorem 4.10. *Let M be a combinatorial manifold with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. A point $x \in |M|$ is regular if and only if there is a local f -level preserving isotopy in some neighbourhood of x .*

Proof. One direction of the equivalence is stated in Lemma 4.8, the other in Lemma 4.9. \square

4.2.3 Constructing Level Preserving Isotopies across Vertices

Isotopies Defined by Combinatorially Equivalent Subdivisions. As a next step, we develop some sufficient conditions for a PL homeomorphism between $|M|_{=h} \times [a, b]$ and $|M|_{\in[a, b]}$ being f -level-preserving. We assume that the PL homeomorphism is given as a combinatorial equivalence between some subdivisions of $M_{=h} \times [a, b]$ and $M_{\in[a, b]}$. If the levels of corresponding vertices agree, that is, the combinatorial equivalence maps a vertex $v = (x, t)$ of $M_{=h} \times [a, b]$ to a vertex v' with $f(v') = t$, then an f -level-preserving isotopy exists.

Lemma 4.11. *Let T_1 and T_2 be combinatorially equivalent simplicial subdivisions of $M_{=h} \times [a, b]$ and $M_{\in[a, b]}$ for some $h \in [a, b]$ and a polytopal complex M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Denote by $\phi^{\text{Vrt}}: \text{Vrt}(T_1) \rightarrow \text{Vrt}(T_2)$ the combinatorial equivalence in terms of vertices between the vertex set of T_1 and the vertex set of T_2 . Assume that ϕ^{Vrt} is f -level-preserving in the sense that for all vertices $v = (y, t) \in \text{Vrt}(T_1)$ the function f at the corresponding vertex $v' = \phi^{\text{Vrt}}(y, t)$ in $\text{Vrt}(T_2)$ has the value $f(v') = t$. Then the PL homeomorphism $\phi: |M|_{=h} \times [a, b] \rightarrow |M|_{\in[a, b]}$ defined by simplex-wise linear interpolation of ϕ^{Vrt} is an f -level-preserving PL isotopy.*

Proof. Since simplex-wise linear interpolation of a combinatorial equivalence of simplicial complexes yields a PL homeomorphism, ϕ is indeed a PL homeomorphism. It remains to show that ϕ_t is a PL homeomorphism from $|M|_{=h}$ to $|M|_{=t}$ for each $t \in [a, b]$. Let x be an element of $|M|_{=h}$ and consider the pair $(x, t) \in |M|_{=h} \times [a, b]$ for some $t \in [a, b]$ which is used to define $\phi_t(x) = \phi(x, t)$. We can write $(x, t) = \sum_i \mu_i(x_i, t_i) = (\sum_i \mu_i x_i, \sum_i \mu_i t_i)$ as a convex combination of vertices (x_i, t_i) of a simplex in T_1 . With this representation, we obtain $f(\phi(\sum_i \mu_i(x_i, t_i))) = f(\sum_i \mu_i \phi(x_i, t_i)) = \sum \mu_i f(\phi^{\text{Vrt}}(x_i, t_i)) = \sum \mu_i t_i = t$ because ϕ is linear on simplices of T_1 and maps vertices of such a simplex to vertices of a simplex in T_2 , the map f is linear on simplices of T_2 , and $f(\phi^{\text{Vrt}}(y, t)) = t$ holds for all vertices $(y, t) \in \text{Vrt}(T_1)$ by our assumption that ϕ^{Vrt} is f -level-preserving. Hence we obtain for each $t \in [a, b]$ that the image $\phi_t(|M|_{=h}) = \phi(|M|_{=h} \times \{t\})$ is contained in $|M|_{=t}$. Since ϕ is bijective, any element of $|M|_{=t} \subseteq |M|_{\in[a, b]}$ has a preimage in $|M|_{=h} \times [a, b]$ under ϕ . But elements $(x', t') \in |M|_{=h} \times [a, b]$ with $t' \neq t$ are impossible as a preimage of a point in $|M|_{=t}$ because such elements are mapped to $|M|_{=t'}$ which is disjoint from $|M|_{=t}$. Thus, $\phi_t(|M|_{=h}) = \phi(|M|_{=h} \times \{t\})$ is not only contained in but also agrees with $|M|_{=t}$. This implies that the PL homeomorphism ϕ restricts to a PL homeomorphism between $|M|_{=h} \times \{t\}$ and $|M|_{=t}$, and since ϕ_t is obtained from this restricted homeomorphism by identifying $|M|_{=h} \times \{t\}$ with $|M|_{=h}$, the map ϕ_t is a PL homeomorphism as well. \square

We can weaken the assumptions in the previous lemma by requiring that

$M_{=h} \times [a, b]$ and $M_{\in[a, b]}$ have combinatorially equivalent subdivisions T_1 and T_2 that are polytopal complexes, but not necessarily simplicial complexes.

Lemma 4.12. *Let T_1 and T_2 polytopal complexes that are combinatorially equivalent subdivisions of $M_{=h} \times [a, b]$ and $M_{\in[a, b]}$ for some $h \in [a, b]$ and a polytopal complex M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Denote by $\phi^{\text{Vrt}}: \text{Vrt}(T_1) \rightarrow \text{Vrt}(T_2)$ the combinatorial equivalence in terms of vertices between the vertex sets. Assume that ϕ^{Vrt} is f -level-preserving in the same sense as in Lemma 4.11, that is $f(\phi^{\text{Vrt}}(y, t)) = t$ for all vertices (y, t) of T_1 . Then there is an f -level-preserving PL isotopy $\phi: M_{=h} \times [a, b] \rightarrow M_{\in[a, b]}$.*

Proof. Using Lemma 1.30, we can find combinatorially equivalent simplicial subdivisions T'_1 and T'_2 of T_1 and T_2 . When we arrange that the combinatorial equivalence $\psi^{\text{Vrt}}: \text{Vrt}(T'_1) \rightarrow \text{Vrt}(T'_2)$ is still f -level-preserving, applying Lemma 4.11 yields the desired result. This is trivial if we use starring at each vertex in according orders for constructing T'_1 and T'_2 , because in this case the vertex sets $\text{Vrt}(T'_1) = \text{Vrt}(T_1)$ and $\text{Vrt}(T'_2) = \text{Vrt}(T_2)$ do not change and $\psi^{\text{Vrt}} = \phi^{\text{Vrt}}$ is still f -level-preserving.

If we construct derived subdivisions T'_1 and T'_2 of T_1 and T_2 instead, a simple method to ensure that ψ^{Vrt} is f -level-preserving, is taking barycentric subdivisions. For those, the combinatorial equivalence ψ^{Vrt} is defined by the assignments $\psi^{\text{Vrt}}(\frac{1}{k}(p_1 + \dots + p_k)) = \frac{1}{k}(\phi^{\text{Vrt}}(p_1) + \dots + \phi^{\text{Vrt}}(p_k))$ for each set of vertices $\{p_1, \dots, p_k\}$ spanning a cell of T_1 . One easily checks that this combinatorial equivalence is f -level-preserving because ϕ^{Vrt} is f -level-preserving and for both maps f and the projection from $|M|_{=h} \times [a, b]$ to the interval $[a, b]$ in the Cartesian product, the value of the map at the barycentre of a cell is the barycentre of the values at the vertices. \square

Intervals without Vertices. As a first application, we show that for intervals $[a, b]$ whose preimage does not contain vertices an f -level-preserving PL isotopy exists.

Lemma 4.13. *Let M be a polytopal complex with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Suppose no vertex of M has a value $f(v) \in [a, b]$ under f . Then for any $h \in [a, b]$, the complexes $M_{=h} \times [a, b]$ and $M_{\in[a, b]}$ are combinatorially equivalent and there is an f -level-preserving PL isotopy $\phi: |M|_{=h} \times [a, b] \rightarrow |M|_{\in[a, b]}$. For any subcomplex K of M , ϕ restricts to an f -level-preserving PL isotopy $\phi|_K$ on $K: \phi|_K: |K|_{=h} \times [a, b] \rightarrow |K|_{\in[a, b]}$. The combinatorial equivalence is given by $S_{=h} \times Y \mapsto S_{\in Y}$ for S a cell of M and $Y \in \{\{a\}, \{b\}, [a, b]\}$. If $h = a$, it restricts on $M_{=a} \times \{a\}$ to the projection onto $M_{=a}$, and on $M_{=a} \times \{b\}$, it restricts to the natural combinatorial equivalence in the sense of Lemma 2.35 between $M_{=a}$, identified with $M_{=a} \times \{b\}$ by projection, and $M_{=b}$.*

Proof. The fact that the suggested assignment yields a combinatorial equivalence of $M_{=h} \times [a, b]$ and $M_{\in[a, b]}$ was proven in Lemma 2.36, including the observation that the equivalence restricts to the natural combinatorial equivalences at the endpoints a and b . Moreover, this combinatorial equivalence obviously restricts to a combinatorial equivalence between $K_{=h} \times [a, b]$ and $K_{\in[a, b]}$ for any subcomplex K . Using Lemma 4.12 we obtain the desired PL isotopy ϕ . Note that the map ϕ is not uniquely determined by the combinatorial equivalence, the interpolated values depend on the chosen simplicial subdivisions. But if we use

on $K_{=h} \times [a, b]$ and $K_{\in[a, b]}$ the subdivisions induced from the chosen subdivisions on $M_{=h} \times [a, b]$ and $M_{\in[a, b]}$, then the interpolated values of the PL isotopy on the subcomplex K agree with the interpolated values of the PL isotopy on the whole complex M . Hence the restriction to K of the f -level-preserving isotopy ϕ on M is an f -level-preserving isotopy on K . \square

Isotopies on the Star of Regular Vertices. Now we consider a regular vertex v and we want to construct an f -level-preserving PL isotopy on $\text{st}(v)$ for some suitable interval $[f(v) - \epsilon, f(v) + \epsilon]$. We can divide this problem into finding an isotopy for the interval $[f(v) - \epsilon, f(v)]$ in the lower level star and an isotopy for the interval $[f(v), f(v) + \epsilon]$ in the upper level star. By considering f and $-f$, we can observe that both problems are symmetric, and it suffices to find a PL isotopy for the interval $[f(v), f(v) + \epsilon]$.

So let $f(v) = a$ and assume general position in the sense that all vertices in $\text{lk}(v)$ have an f -value different from $f(v)$. Choose ϵ small enough such that for $b = f(v) + \epsilon$ no vertex in $\text{lk}(v)$ has an f -value in $[a, b]$. We want to show that with the general position assumption, an f -level preserving PL isotopy between $\text{st}(v)_{=a}$ and $\text{st}(v)_{=b}$ exists. As a first step, we list some complexes that are combinatorially equivalent in a natural way; this preliminary observation is used in the subsequent proofs.

Lemma 4.14. *Let v be a vertex in a combinatorial manifold M with $f: |M| \rightarrow \mathbb{R}$ linear on cells such that no vertex of $\text{lk}(v)$ has an f -value in the interval $[a, b] = [f(v), f(v) + \epsilon]$. Then there are natural combinatorial equivalences between the following polytopal complexes that are unions of complexes with a common subcomplex as intersection.*

$$\begin{array}{ll}
 \text{st}(v)_{=b} \cup \text{lk}(v)_{\in[a, b]} & \text{with intersection } \text{lk}(v)_{=b} \\
 \text{lk}(v)_{\geq b} \cup \text{lk}(v)_{\in[a, b]} & \text{with intersection } \text{lk}(v)_{=b} \\
 \text{st}(v)_{=b} \times \{b\} \cup \text{lk}(v)_{=b} \times [a, b] & \text{with intersection } \text{lk}(v)_{=b} \times \{b\} \\
 \text{lk}(v)_{\geq b} \times \{b\} \cup \text{lk}(v)_{=b} \times [a, b] & \text{with intersection } \text{lk}(v)_{=b} \times \{b\} \\
 \text{lk}(v)_{\geq a} \times \{b\} \cup \text{lk}(v)_{=a} \times [a, b] & \text{with intersection } \text{lk}(v)_{=a} \times \{b\}
 \end{array}$$

The equivalences are obtained by the natural combinatorial equivalences of the left hand terms of the unions and the natural combinatorial equivalences of the right hand terms which agree on the respective common intersections.

Proof. Complexes of the form $X \times \{b\}$ can be naturally identified with the complex X by projection. For the left hand terms of the unions, $\text{st}(v)_{=b}$ is naturally combinatorially equivalent to $\text{lk}(v)_{\geq b}$ by Lemma 2.38 and this equivalence restricts to the identity on $\text{lk}(v)_{=b}$. The complexes $\text{lk}(v)_{\geq b}$ and $\text{lk}(v)_{\geq a}$ are naturally combinatorially equivalent by Lemma 2.35, because no vertex of $\text{lk}(v)$ has a value in $[a, b]$. This equivalence restricts to the natural equivalence between $\text{lk}(v)_{=b}$ and $\text{lk}(v)_{=a}$.

For the right hand terms of the unions, by Lemma 2.36, the complexes $\text{lk}(v)_{\in[a, b]}$, $\text{lk}(v)_{=b} \times [a, b]$, and $\text{lk}(v)_{=a} \times [a, b]$ are combinatorially equivalent and the equivalences restrict to the natural equivalences between $\text{lk}(v)_{=b}$, $\text{lk}(v)_{=b} \times \{b\}$, and $\text{lk}(v)_{=a} \times \{b\}$. \square

For the existence of an f -level-preserving PL isotopy on $\text{st}(v)$ between the levels a and b , it is necessary that the level sets $\text{st}(v)_{=a}$ and $\text{st}(v)_{=b}$ are PL-

homeomorphic. The next lemma shows that the desired isotopy can be constructed from such a homeomorphism and an isotopy on the link, provided that the maps match with each other in the sense that they both induce the same homeomorphism between $\text{lk}(v)_{=a}$ and $\text{lk}(v)_{=b}$. The proof idea is illustrated by a simple 2-dimensional example in Figure 4.1. It is based on the proof given in [13] for the existence of a level-preserving isotopy across regular values in the 3-dimensional case.

Lemma 4.15. *Let v be a regular vertex of a combinatorial manifold M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells and let $f(v) = a < b$. Then there is an f -level-preserving PL isotopy between $\text{st}(v)_{=a}$ and $\text{st}(v)_{=b}$ if there is an f -level-preserving PL isotopy between $\text{lk}(v)_{=a}$ and $\text{lk}(v)_{=b}$ and there is a PL homeomorphism between $\text{st}(v)_{=a}$ and $\text{st}(v)_{=b}$ which agrees on $\text{lk}(v)$ with the isotopy.*

More precisely, if $\psi: |\text{lk}(v)|_{=a} \times [a, b] \rightarrow |\text{lk}(v)|_{\in[a, b]}$ is an f -level-preserving PL isotopy on the link and $\chi: |\text{st}(v)|_{=a} \rightarrow |\text{st}(v)|_{=b}$ a PL homeomorphism with $\chi|_{|\text{lk}(v)|_{=a}} = \psi_b$, then an f -level-preserving PL isotopy $\phi: |\text{st}(v)|_{=a} \times [a, b] \rightarrow \text{st}(v)_{\in[a, b]}$ with $\phi_b = \chi$ and $\phi|_{|\text{lk}(v)|_{=a} \times [a, b]} = \psi$ is obtained by a cone construction.

Proof. Let π be the projection $\pi: \text{st}(v)_{=a} \times \{b\} \rightarrow \text{st}(v)_{=a}$ that identifies $\text{st}(v)_{=a} \times \{b\}$ with $\text{st}(v)_{=a}$. Then the restrictions of the maps $\chi \circ \pi$, $\psi_b \circ \pi$, and ψ to the intersection of their domains, which is $|\text{lk}(v)|_{=a} \times \{b\}$, agree. Hence we can define a map on the union of the domains of ψ and $\chi \circ \pi$ by case distinction. This yields a PL map $\theta: |\text{lk}(v)|_{=a} \times [a, b] \cup |\text{st}(v)|_{=a} \times \{b\} \rightarrow |\text{lk}(v)|_{\in[a, b]} \cup |\text{st}(v)|_{=b}$ which is well defined by $\theta|_{|\text{lk}(v)|_{=a} \times [a, b]} = \psi$ and $\theta|_{|\text{st}(v)|_{=a} \times \{b\}} = \chi \circ \pi$ and can be considered as the union of the maps ψ and $\chi \circ \pi$. Since the images of ψ and $\chi \circ \pi$, namely $|\text{lk}(v)|_{\in[a, b]}$ and $|\text{st}(v)|_{=b}$, intersect in $|\text{lk}(v)|_{=b}$, which is exactly the image $\psi(|\text{lk}(v)|_{=a} \times \{b\}) = \chi \circ \pi(|\text{lk}(v)|_{=a} \times \{b\})$ of the intersection of the domains of ψ and $\chi \circ \pi$, the union θ of the PL homeomorphisms ψ and $\chi \circ \pi$ is a PL homeomorphism as well. Moreover, θ is f -level-preserving in the sense that $f(\theta(x, t)) = t$ for $x \in |\text{lk}(v)|_{=a}$ and $t \in [a, b]$ or $x \in |\text{st}(v)|_{=a}$ and $t = b$. The fact that θ is a PL homeomorphism implies in particular, that there are combinatorially equivalent simplicial subdivisions Θ_1 of $|\text{lk}(v)|_{=a} \times [a, b] \cup |\text{st}(v)|_{=a} \times \{b\}$ and Θ_2 of $|\text{lk}(v)|_{\in[a, b]} \cup |\text{st}(v)|_{=b}$ such that θ is a simplicial bijection. Therefore restricting the homeomorphism θ to the vertex set of Θ_1 yields a combinatorial equivalence in terms of vertices $\theta^{\text{Vrt}}: \text{Vrt}(\Theta_1) \rightarrow \text{Vrt}(\Theta_2)$ that is f -level-preserving.

Observe that $\text{st}(v)_{=a} \times [a, b]$ and $\text{st}(v)_{\in[a, b]}$ can be subdivided by performing elementary starrings at the vertex (v, a) of $\text{st}(v)_{=a} \times [a, b]$ and at the vertex v of $\text{st}(v)_{\in[a, b]}$, giving the subdivided complexes the structure of a cone: The cone with base $|\text{lk}(v)|_{=a} \times [a, b] \cup |\text{st}(v)|_{=a} \times \{b\}$ and apex (v, a) is a subdivision of $\text{st}(v)_{=a} \times [a, b]$ obtained by starring $\text{st}(v)_{=a} \times [a, b]$ at the vertex (v, a) , and the cone with base $|\text{lk}(v)|_{\in[a, b]} \cup |\text{st}(v)|_{=b}$ and apex v is a subdivision of $\text{st}(v)_{\in[a, b]}$ obtained by starring $\text{st}(v)_{\in[a, b]}$ at the vertex v . Define the cones $T_1 = (v, a)\Theta_1$ and $T_2 = v\Theta_2$ and extend the combinatorial equivalence θ^{Vrt} between the vertex sets of Θ_1 and Θ_2 to a combinatorial equivalence $\phi^{\text{Vrt}}: \text{Vrt}(T_1) \rightarrow \text{Vrt}(T_2)$ between the vertex sets of T_1 and T_2 by letting $\phi^{\text{Vrt}}(v, a) = v$ and $\phi^{\text{Vrt}}(u) = \theta^{\text{Vrt}}(u)$ for each vertex u of Θ_1 . Then T_1 and T_2 are combinatorially equivalent simplicial subdivisions of $\text{st}(v)_{=a} \times [a, b]$ and $\text{st}(v)_{\in[a, b]}$, and the combinatorial equivalence ϕ^{Vrt} is f -level-preserving. By Lemma 4.11, the linear interpolation ϕ of ϕ^{Vrt} is an f -level-preserving PL-isotopy as desired. Moreover, ϕ agrees with θ on the bases of the cones, because ϕ^{Vrt} is defined as θ^{Vrt} on the bases, and

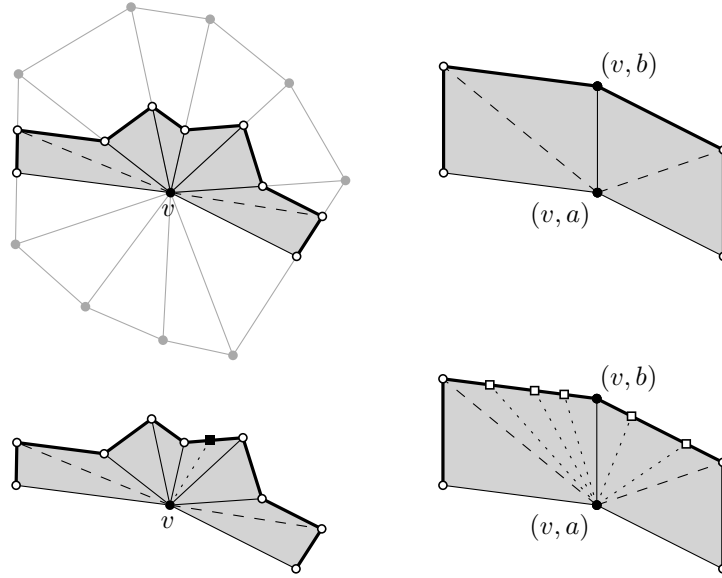


Figure 4.1: Illustration for Lemma 4.15. Depicted on top left is the star of a regular vertex v . The grey-shaded area is the domain of the level set $\text{st}(v)_{\in[a,b]}$ for the interval $[a, b] = [f(v), f(v) + \epsilon]$, the uninterrupted black edges are the edges of $\text{st}(v)_{\in[a,b]}$. The dashed edges are the additional edges introduced by an elementary starring of $\text{st}(v)_{\in[a,b]}$ at v . This starring subdivides $\text{st}(v)_{\in[a,b]}$ into a cone with apex v and the complex $\text{lk}(v)_{\in[a,b]} \cup \text{st}(v)_{=b}$ as base. This base is highlighted by thickened edges. The figure on top right depicts the complex $\text{st}(v)_{=a} \times [a, b]$. The dashed edges belong to the subdivision of this complex obtained by an elementary starring at (v, a) . This subdivision yields a cone with apex (v, a) and base $\text{lk}(v)_{=a} \times [a, b] \cup \text{st}(v)_{=a} \times \{b\}$. Again the base is highlighted by thickened edges. The bottom row displays combinatorially equivalent subdivisions of the two cones in the top row: The additional vertices indicated by squares yield combinatorially equivalent subdivisions of the bases, and together with the dotted lines combinatorially equivalent subdivisions of $\text{st}(v)_{\in[a,b]}$ and $\text{st}(v)_{=a} \times [a, b]$.

therefore ϕ restricts to ψ on $\text{lk}(v)_{=a} \times [a, b]$ and the maps ϕ_b and χ agree by the choice of θ as union of ψ and $\chi \circ \pi$. \square

The previous lemma showed how a PL isotopy on the star can be defined based on the assumption that a PL isotopy on the link and a matching PL homeomorphism between $\text{st}(v)_{=a}$ and $\text{st}(v)_{=b}$ exist. Now we prove for a regular vertex v and a map f in general position that there are pairs of maps fulfilling the assumption, which implies the existence of a PL isotopy on the star of such vertices.

Lemma 4.16. *Let v be a regular vertex of a combinatorial manifold M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Assume that no vertex of $\text{lk}(v)$ has the same f -value $f(v) = a$ as v and choose $b > a$ small enough such that $f(u) \notin [a, b]$ holds for all vertices $u \in \text{lk}(v)$. Then there is an f -level-preserving PL isotopy $\phi: |\text{st}(v)|_{=a} \times [a, b] \rightarrow |\text{st}(v)|_{\in [a, b]}$ on the star of v .*

If an f -level-preserving PL isotopy $\psi: |\text{lk}(v)|_{=a} \times [a, b] \rightarrow |\text{lk}(v)|_{\in [a, b]}$ on the link $\text{lk}(v)$ is given, then ϕ can be chosen in such a way that its restriction $\phi|_{\text{lk}(v)}$ to $\text{lk}(v)$ agrees with ψ .

Proof. If an f -level-preserving PL isotopy ψ for the interval $[a, b]$ on $\text{lk}(v)$ is not already given, we can choose one by applying Lemma 4.13 for the complex $\text{lk}(v)$ because no vertex of $\text{lk}(v)$ has an f -value in $[a, b]$. We show that we can construct a PL homeomorphism $\chi: |\text{st}(v)|_{=a} \rightarrow |\text{st}(v)|_{=b}$ as in Lemma 4.15.

Since v is regular, the complexes $\text{st}(v)_{=a}$ and $\text{lk}(v)_{\geq a}$ are combinatorial balls of dimension $\dim(M) - 1$. As noted in Lemma 4.14, there is a natural combinatorial equivalence between $\text{lk}(v)_{\geq a} \times \{b\} \cup \text{lk}(v)_{=a} \times [a, b]$ and $\text{st}(v)_{=b} \cup \text{lk}(v)_{\in [a, b]}$ that restricts to a combinatorial equivalence between $\text{lk}(v)_{\geq a}$ and $\text{st}(v)_{=b}$, thus $\text{st}(v)_{=b}$ is a combinatorial $(\dim(M) - 1)$ -ball as well.

If v is an interior vertex, then $\text{lk}(v)_{=a}$ is the boundary of the ball $\text{st}(v)_{=a}$ and $\text{lk}(v)_{=b}$ is the boundary of the ball $\text{st}(v)_{=b}$. The map ψ_b defined by the isotopy on the link is a PL homeomorphism between the boundaries $\text{lk}(v)_{=a}$ and $\text{lk}(v)_{=b}$ of these balls. Since any PL homeomorphism between boundaries of PL balls can be extended to the whole balls by Fact 1.39, we obtain a PL homeomorphism $\chi: |\text{st}(v)|_{=a} \rightarrow |\text{st}(v)|_{=b}$ such that χ and ψ fulfil the assumptions of Lemma 4.15. Applying the Lemma yields an f -level-preserving isotopy ϕ as desired.

It remains to consider the case of a boundary vertex $v \in \text{Vrt}(\text{bd}(M))$. First, we construct a PL isotopy on the star $\text{st}_{\text{bd}(M)}(v)$ of v in the boundary. The link $\text{lk}_{\text{bd}(M)}(v)$ in the boundary coincides with the boundary $\text{bd}(\text{lk}_M(v))$ of the link in the whole manifold. Recall from Corollary 2.37 that the assumption $f(u) \notin [a, b]$ for vertices $u \in \text{lk}(v)$ asserts that for each $h \in [a, b]$, the level set $(\text{lk}_M(v))_{=h}$ is a locally flat sub-manifold of $\text{lk}_M(v)$, which implies in particular that the boundary $\text{bd}((\text{lk}_M(v))_{=h})$ of such a level set coincides with the level set $(\text{lk}_{\text{bd}(M)}(v))_{=h}$ in the boundary. This means, because any PL homeomorphism ψ_h maps boundary points to boundary points, that ψ restricts to an f -level-preserving isotopy $\psi|_{\text{lk}_{\text{bd}(M)}(v)}$ on $\text{lk}_{\text{bd}(M)}(v)$. Considering v as an interior vertex of $\text{bd}(M)$, we can apply the construction for interior vertices described above. We obtain an f -level-preserving PL isotopy $\phi^{\text{bd}}: |\text{st}_{\text{bd}(M)}(v)|_{=a} \times [a, b] \rightarrow |\text{st}_{\text{bd}(M)}(v)|_{\in [a, b]}$ on the star $\text{st}_{\text{bd}(M)}(v)$ of v taken in the boundary of M such that the restriction $\phi^{\text{bd}}|_{\text{lk}_{\text{bd}(M)}(v)}$ agrees with the restriction $\psi|_{\text{lk}_{\text{bd}(M)}(v)}$.

Now consider the boundaries of the level sets $(\text{st}_M(v))_{=a}$ and $(\text{st}_M(v))_{=b}$ in the star, which are both combinatorial balls. They both decompose into the level set in the link and the level set in the boundary as follows: The boundary of the combinatorial ball $(\text{st}_M(v))_{=a}$ is the union of $(\text{lk}_M(v))_{=a}$ and $(\text{st}_{\text{bd}(M)}(v))_{=a}$ which intersect in $(\text{lk}_{\text{bd}(M)}(v))_{=a}$, and the boundary of the combinatorial ball $(\text{st}_M(v))_{=b}$ is the union of $(\text{lk}_M(v))_{=b}$ and $(\text{st}_{\text{bd}(M)}(v))_{=b}$ which intersect in $(\text{lk}_{\text{bd}(M)}(v))_{=b}$. Furthermore, we have the PL homeomorphism $\psi_b: |\text{lk}_M(v)|_{=a} \rightarrow |\text{lk}_M(v)|_{=b}$ between the level sets in the link forming the respective first parts of the two decompositions and the PL homeomorphism $\phi_b^{\text{bd}}: |\text{st}_{\text{bd}(M)}(v)|_{=a} \rightarrow |\text{st}_{\text{bd}(M)}(v)|_{=b}$ between the level sets in the boundary forming the respective second parts. Both homeomorphisms $\psi_b|_{(\text{lk}_{\text{bd}(M)}(v))_{=a}}$ and $\phi_b^{\text{bd}}|_{(\text{lk}_{\text{bd}(M)}(v))_{=a}}$ agree on their common domain $|\text{lk}_{\text{bd}(M)}(v)|_{=a}$ which is mapped homeomorphically to $|\text{lk}_{\text{bd}(M)}(v)|_{=b}$. Thus we can define a PL homeomorphism ψ_b^+ on the union of the domains by taking the union of the maps ψ_b and ϕ_b^{bd} : $\psi_b^+: |\text{lk}_M(v)|_{=a} \cup |\text{st}_{\text{bd}(M)}(v)|_{=a} \rightarrow |\text{lk}_M(v)|_{=b} \cup |\text{st}_{\text{bd}(M)}(v)|_{=b}$. This map ψ_b^+ is again a PL homeomorphism between the boundaries of the combinatorial balls $(\text{st}_M(v))_{=a}$ and $(\text{st}_M(v))_{=b}$ and can be extended to a PL homeomorphism $\chi: |\text{st}_M(v)|_{=a} \rightarrow |\text{st}_M(v)|_{=b}$ between the balls by Fact 1.39. Again we have constructed a map χ as desired and can proceed as above. \square

Extending Isotopies to the Whole Manifold and across Several Regular Vertices. The next result shows that for a map f in general position, a regular vertex v , and a sufficiently small interval $[f(v), f(v) + \epsilon]$, a PL isotopy between the global level sets $M_{=f(v)}$ and $M_{=f(v)+\epsilon}$ exists.

Lemma 4.17. *Let v be a regular vertex of a combinatorial manifold M with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Assume that v is the only vertex of M with f -value $f(v) = a$. Furthermore, choose $b > a$ such that $f(u) \notin [a, b]$ holds for all vertices $u \neq v$ of M . Then there is an f -level-preserving PL isotopy $\phi: |M|_{=a} \times [a, b] \rightarrow |M|_{\in [a, b]}$.*

Proof. We regard M as the union of the subcomplexes $\text{del}(v)$ and $\text{st}(v)$ intersecting in $\text{lk}(v)$ and combine PL isotopies on the two parts that agree on the intersection to a PL isotopy on the whole complex M .

An f -level-preserving isotopy ϕ^{del} on $\text{del}(v)$ that restricts to an f -level-preserving isotopy ϕ^{lk} on $\text{lk}(v)$ exists by Lemma 4.13. For the star, there is an f -level-preserving PL-isotopy ϕ^{st} on $\text{st}(v)$ that restricts to ϕ^{lk} on $\text{lk}(v)$ by Lemma 4.16. Hence, the union ϕ of ϕ^{del} and ϕ^{st} is well-defined and an f -level-preserving isotopy on M as desired. \square

Remark 4.18. The same proof works also for several regular vertices with the same regular value a as long as their stars intersect only in their links, in other words, since M is simplicial, if the vertices with value a are not adjacent by an edge of M .

From Lemma 4.8 and Lemma 4.16 or Lemma 4.17, we obtain alternative characterisation of regular vertices in manifolds where f -values of adjacent or even of all vertices are distinct.

Theorem 4.19. *Let M be combinatorial manifold and $f: |M| \rightarrow \mathbb{R}$ linear on cells such that no adjacent vertices have same f -value. Then a vertex v in M is regular if and only if there is an f -level-preserving PL isotopy on $\text{st}(v)$ for*

some interval $[a, b]$ containing $f(v)$ in its interior. If all f -values of vertices are distinct, then a vertex v is regular if and only if there is an f -level-preserving PL isotopy on M for some interval $[a, b]$ containing $f(v)$ in its interior.

Proof. We showed in Lemma 4.8 that an f -level-preserving PL isotopy in a neighbourhood of v implies its regularity. Since $\text{st}(v)$ and the whole manifold M are neighbourhoods of v , in both claims the criterion for regularity of v is sufficient. For necessity of the criterion in the first claim, we use Lemma 4.16. It yields under the given assumption on the f -values that there is an f -level-preserving PL isotopy on $\text{st}(v)$ for some interval $[f(v), f(v) + \epsilon_+]$. A symmetric version of the lemma yields an f -level-preserving PL isotopy on $\text{st}(v)$ for some interval $[f(v) - \epsilon_-, f(v)]$. By Lemma 4.6, both isotopies can be combined to an f -level-preserving PL isotopy on $\text{st}(v)$ for the interval $[f(v) - \epsilon_-, f(v) + \epsilon_+]$ as desired. For the second claim, Lemma 4.17 and Remark 4.18 and their symmetric variants give isotopies on M for intervals $[f(v), f(v) + \epsilon_+]$ and $[f(v) - \epsilon_-, f(v)]$ that can be combined to the desired f -level-preserving PL isotopy on M for the interval $[f(v) - \epsilon_-, f(v) + \epsilon_+]$. \square

We can also conclude that for any interval $[a, b]$ without critical values and a map f in general position an f -level-preserving PL isotopy exists. This result is a piecewise linear analogue of Fact 1.3 from smooth Morse theory.

Theorem 4.20. *Let M be a combinatorial manifold and $f: |M| \rightarrow \mathbb{R}$ linear on cells. Assume that the preimage $f^{-1}[a, b]$ contains no critical points and that any pair of vertices of M lying in that preimage has distinct f -values or is not adjacent. Then there is an f -level-preserving PL isotopy $\phi: |M|_{=a} \times [a, b] \rightarrow |M|_{\in[a, b]}$.*

Proof. Subdivide the interval such that each subinterval contains at most one value that is attained by f at a vertex and this value is one of the endpoints of the subinterval. Such a subdivision can be obtained for example by subdividing $[a, b]$ at every value that is attained by f at a vertex and afterwards halving the intervals between two consecutive such values. We obtain isotopies on each subinterval as follows: Apply Lemma 4.13 for intervals where both endpoints are not attained at vertices. For the remaining intervals, Lemma 4.17 together with Remark 4.18 applies, either in original form for intervals whose left endpoint is attained at a vertex or in a symmetric variant for intervals whose right endpoint is attained at a vertex. Combine the obtained isotopies of the subintervals to an isotopy of the whole interval according to Lemma 4.6. \square

Chapter 5

Piecewise Linear Analogues of Classical Results from Smooth Morse Theory

This chapter collects piecewise linear analogues of some fundamental results from classical smooth Morse theory. The selection of the results follows the treatment in Milnor's book on Morse theory [36], which we also used already as a guideline for the selected facts in the introductory review of smooth Morse theory in Section 1.1. We usually assume that M is a combinatorial manifold without boundary with a PL Morse function $f: |M| \rightarrow \mathbb{R}$ linear on cells and in general position. The first two theorems relate the change in homotopy type, when going from a level set $M_{\leq a}$ to a level set $M_{\leq b}$, to the critical points of M that have an f -value in the range $[a, b]$. Theorem 5.1 states that no change occurs when the range contains no critical value, Theorem 5.2 characterises the change occurring when passing a single critical point as the attachment of a cell. This leads to the description of the homotopy type of M in terms of a CW-complex with one λ -cell for each critical point of index λ , which is presented in Theorem 5.3. Finally we derive the Morse inequalities in Theorem 5.4 and the Reeb theorem in Theorem 5.5.

5.1 The Fundamental Results on the Behaviour of Level Sets

Absence of Topological Changes for Intervals without Critical Values.

The first result deals with the case that the interval $[a, b]$ contains no critical value. In this case the level sets $M_{\leq a}$ and $M_{\leq b}$ are homeomorphic and the latter deformation retracts to the former. This result is the PL analogue of Fact 1.2 from smooth Morse theory, which we cited from [36, Theorem 3.1, p. 12]. The proof given here relies on the existence of an f -level-preserving PL isotopy between $M_{=a-\delta}$ and $M_{=b}$ for some $\delta > 0$. The assumption that such an isotopy exists can be regarded as a replacement for the general position assumption and the condition that the interval contains no critical value, because Theorem 4.20 asserts the existence of the isotopy for a map in general position and an interval

Table 5.1: Transformation of $M_{\leq b}$ into $M_{\leq a}$ by PL homeomorphisms for an interval $[a, b]$ without critical values under the assumptions of Theorem 5.1.

$$\begin{aligned}
M_{\leq b} &\stackrel{\text{PL}}{\equiv} M_{\leq a-\epsilon/2} \cup M_{\in[a-\epsilon/2, b]} \\
&\stackrel{\text{PL}}{\equiv} M_{\leq a-\epsilon/2} \times \{a-\epsilon/2\} \cup M_{=a-\epsilon/2} \times [a-\epsilon/2, b] \\
&\stackrel{\text{PL}}{\equiv} M_{\leq a-\epsilon/2} \times \{a-\epsilon\} \cup M_{=a-\epsilon/2} \times [a-\epsilon, a] \\
&\stackrel{\text{PL}}{\equiv} M_{\leq a-\epsilon} \times \{a-\epsilon\} \cup M_{=a-\epsilon} \times [a-\epsilon, a] \\
&\stackrel{\text{PL}}{\equiv} M_{\leq a-\epsilon} \cup M_{\in[a-\epsilon, a]} \\
&\stackrel{\text{PL}}{\equiv} M_{\leq a}
\end{aligned}$$

without critical values.

Theorem 5.1. *Let M be a combinatorial manifold with a map $f: |M| \rightarrow \mathbb{R}$ linear on cells. Assume that for some $\delta > 0$ an f -level-preserving PL isotopy $\phi: M_{=a-\delta} \times [a-\delta, b] \rightarrow M_{\in[a-\delta, b]}$ exists. A sufficient condition for this assumption is that the preimage $f^{-1}[a, b]$ contains no critical points and that the vertices of M lying in that preimage have distinct f -values; the second part of the condition already implies that non-vertices of M lying in $f^{-1}[a, b]$ are not critical, so that only vertices have to be checked for the first part of the condition. Then $M_{\leq a}$ is PL-homeomorphic to $M_{\leq b}$, and $M_{\leq b}$ deformation retracts to $M_{\leq a}$.*

Proof. Let us first settle the claimed sufficient condition for the existence of the isotopy. If the preimage $f^{-1}[a, b]$ contains no critical points and only vertices with distinct f -values, this property also holds for the preimage $f^{-1}[a-\delta, b]$ of a slightly larger interval $[a-\delta, b]$ provided that δ is small enough. Namely, if we choose δ in such a way that no vertex has a value in the interval $[a-\delta, a)$, then all non-vertices in $f^{-1}[a-\delta, b]$, in particular the additional points in $f^{-1}[a-\delta, a)$, lie in cells where f is non-constant and hence are regular points by Lemma 3.22. The existence of the isotopy on the interval $[a-\delta, b]$ follows from the absence of critical values according to Theorem 4.20.

Homeomorphism. Now we show how to obtain a PL homeomorphism between $M_{\leq a}$ and $M_{\leq b}$. Choose a positive $\epsilon \leq \delta$ small enough such that $f^{-1}[a-\epsilon, a)$ contains no vertices. We start with $M_{\leq b}$ and transform it in several steps by applying PL homeomorphisms, which are sometimes given by subdivisions or combinatorial equivalences, into the complex $M_{\leq a}$. The homeomorphic complexes encountered in the intermediate steps are listed in Table 5.1. In the following we explain the PL homeomorphisms underlying these steps.

In the first step, we only perform a subdivision of $M_{\leq b}$. Hence the PL homeomorphism is simply the identity on $|M|_{\leq b}$. So consider the subdivision of $M_{\leq b}$ induced by the level set $M_{=a-\epsilon/2}$, namely the complex formed by the union of $M_{\leq a-\epsilon/2}$ and $M_{\in[a-\epsilon/2, b]}$.

Using Lemmata 4.5 and 4.7 we can construct from the PL isotopy ϕ between $M_{=a-\delta}$ and $M_{=b}$ a PL isotopy $\psi: |M|_{=a-\epsilon/2} \times [a-\epsilon/2, b] \rightarrow |M|_{\in[a-\epsilon/2, b]}$ for

the smaller interval $[a - \epsilon/2, b]$ such that $\psi_{a-\epsilon/2}$ is the identity on $|M|_{=a-\epsilon/2}$. The projection $\pi: |M|_{\leq a-\epsilon/2} \times \{a - \epsilon/2\} \rightarrow |M|_{\leq a-\epsilon/2}$ to the factor $|M|_{\leq a-\epsilon/2}$ of the Cartesian product is a PL homeomorphism because the other factor is a singleton and it agrees with ψ on the common intersection $|M|_{=a-\epsilon/2} \times \{a - \epsilon/2\}$ of the domains of ψ and π . Thus the union of ψ and π is a well-defined PL homeomorphism between $M_{\leq a-\epsilon/2} \times \{a - \epsilon/2\} \cup M_{=a-\epsilon/2} \times [a - \epsilon/2, b]$ and $M_{\leq a-\epsilon/2} \cup M_{\in[a-\epsilon/2, b]}$. The inverse of this homeomorphism yields the second step of the transformation.

There is a natural combinatorial equivalence of the intervals $[a - \epsilon/2, b]$ and $[a - \epsilon, a]$, which restricts to a combinatorial equivalence between the singletons $\{a - \epsilon/2\}$ and $\{a - \epsilon\}$. These combinatorial equivalences can be extended to the Cartesian products, so that $M_{\leq a-\epsilon/2} \times \{a - \epsilon/2\}$ is combinatorially equivalent to $M_{\leq a-\epsilon/2} \times \{a - \epsilon\}$ and $M_{=a-\epsilon/2} \times [a - \epsilon/2, b]$ is combinatorially equivalent to $M_{=a-\epsilon/2} \times [a - \epsilon, a]$. Since these combinatorial equivalences both restrict to the same combinatorial equivalence between $M_{=a-\epsilon/2} \times \{a - \epsilon/2\}$ and $M_{=a-\epsilon/2} \times \{a - \epsilon\}$ on their common intersection, they yield a combinatorial equivalence between $M_{\leq a-\epsilon/2} \times \{a - \epsilon/2\} \cup M_{=a-\epsilon/2} \times [a - \epsilon/2, b]$ and $M_{\leq a-\epsilon/2} \times \{a - \epsilon\} \cup M_{=a-\epsilon/2} \times [a - \epsilon, a]$. The combinatorial equivalence induces the PL homeomorphism underlying the third step.

Since no vertex of M has a value in $[a - \epsilon, a - \epsilon/2]$, Lemma 2.35 yields a combinatorial equivalence between $M_{\leq a-\epsilon/2}$ and $M_{\leq a-\epsilon}$ which restricts to a combinatorial equivalence between $M_{=a-\epsilon/2}$ and $M_{=a-\epsilon}$. This equivalence extends to the Cartesian products and the extensions match on the common intersection, so that we obtain a combinatorial equivalence and consequently a PL homeomorphism between the unions $M_{\leq a-\epsilon/2} \times \{a - \epsilon\} \cup M_{=a-\epsilon/2} \times [a - \epsilon, a]$ and $M_{\leq a-\epsilon} \times \{a - \epsilon\} \cup M_{=a-\epsilon} \times [a - \epsilon, a]$.

Use again Lemmata 4.5 and 4.7 to construct from the PL isotopy ϕ between $M_{=a-\delta}$ and $M_{=b}$ a PL isotopy $\chi: |M|_{=a-\epsilon} \times [a - \epsilon, a] \rightarrow |M|_{\in[a-\epsilon, a]}$ such that $\chi_{a-\epsilon}$ is the identity on $|M|_{=a-\epsilon}$. With this choice χ matches on $|M|_{=a-\epsilon} \times \{a - \epsilon\}$ with the projection of $|M|_{\leq a-\epsilon} \times \{a - \epsilon\}$ to $|M|_{\leq a-\epsilon}$, so that the union of χ and the projection yields a PL homeomorphism between $M_{\leq a-\epsilon} \times \{a - \epsilon\} \cup M_{=a-\epsilon} \times [a - \epsilon, a]$ and $M_{\leq a-\epsilon} \cup M_{\in[a-\epsilon, a]}$.

Similarly to the first step, the last step is given by the identity on the domain, because the obtained complex $M_{\leq a-\epsilon} \cup M_{\in[a-\epsilon, a]}$ is a subdivision of $M_{\leq a}$ induced by the level set $M_{=a-\epsilon}$.

Deformation Retraction. Now we show that $M_{\leq b}$ deformation retracts to $M_{\leq a}$. Consider the subdivision of $M_{\leq b}$ into $M_{\leq a} \cup M_{\in[a, b]}$. Modify the PL isotopy ϕ using Lemmata 4.5 and 4.7 to obtain a PL isotopy $\theta: |M|_{=a} \times [a, b] \rightarrow |M|_{\in[a, b]}$ such that θ_a is the identity of $|M|_{=a}$. The union of θ with the projection of $|M|_{\leq a} \times \{a\}$ to $|M|_{\leq a}$ yields a PL homeomorphism between $|M|_{\leq a} \times \{a\} \cup |M|_{=a} \times [a, b]$ and $|M|_{\leq a} \cup |M|_{\in[a, b]} = |M|_{\leq b}$ that maps $|M|_{\leq a} \times \{a\}$ to $|M|_{\leq a}$. Therefore a deformation retraction of $M_{\leq a} \times \{a\} \cup M_{=a} \times [a, b]$ to $M_{\leq a} \times \{a\}$ induces via the homeomorphism the desired deformation retraction of $M_{\leq b}$ to $M_{\leq a}$. We show how to define a deformation retraction of $M_{\leq a} \times \{a\} \cup M_{=a} \times [a, b]$ to $M_{\leq a} \times \{a\}$.

Triangulate $M_{\leq a} \times \{a\} \cup M_{=a} \times [a, b]$ without introducing new vertices according to Fact 1.25. Call the obtained complex K and triangulate $K \times [0, 1]$ again without introducing new vertices. Define the deformation retraction $\psi: K \times [0, 1] \rightarrow K$ by simplex-wise linear interpolation between the following values of ψ at vertices of $K \times [0, 1]$: For a vertex $v = (u, a)$ of $M_{\leq a} \times \{a\}$,

define $\psi(v, 0) = \psi(v, 1) = v$. This definition includes already the vertices of $M_{=a} \times \{a\} \subseteq M_{=a} \times [a, b]$. For the remaining vertices (u, b) of $M_{=a} \times \{b\} \subseteq M_{=a} \times [a, b]$, define $\psi(u, b, 0) = (u, b)$ and $\psi(u, b, 1) = (u, a)$. For verifying that ψ is well-defined, we have to check that for each cell of the triangulation of $K \times [0, 1]$ the ψ -values of its vertices are vertices of a single cell of $M_{\leq a} \times \{a\} \cup M_{=a} \times [a, b]$. So consider a cell T from the triangulation of $K \times [0, 1]$. It is contained in some cell $S \times [0, 1]$ of $K \times [0, 1]$ for S a cell of K and the vertices of T are among the vertices of $S \times [0, 1]$. Moreover the cell S is either contained in some cell U of $M_{\leq a} \times \{a\}$ or in some cell $U' \times [a, b]$ of $M_{=a} \times [a, b]$ and the vertices of S are among the vertices of U or $U' \times [a, b]$ respectively. The chosen values for the vertices assert that the vertices of T are mapped to vertices of U in the former case, and to vertices of $U' \times [a, b]$ in the latter case. Furthermore they assert that ψ is a homotopy relative to $|M|_{\leq a} \times \{a\}$, that ψ_0 is the identity on $|K| = |M|_{\leq a} \times \{a\} \cup |M|_{=a} \times [a, b]$, and that ψ_1 is a retraction of $|M|_{\leq a} \times \{a\} \cup |M|_{=a} \times [a, b]$ to $|M|_{\leq a} \times \{a\}$. Therefore ψ is a strong deformation retraction of $M_{\leq a} \times \{a\} \cup M_{=a} \times [a, b]$ to $M_{\leq a} \times \{a\}$ as desired. \square

Attachment of Cells as Effect of Passing Critical Values. Now we turn to the case that the interval $[a, b]$ contains a single critical value. Since we assume general position we can restrict our attention to a single non-degenerate critical vertex v such that v is the only vertex of M lying in $|M|_{\in[a, b]}$. As stated in Fact 1.4 the change in homotopy type when going from $M_{\leq a}$ to $M_{\leq b}$ in the corresponding smooth case is equivalent to the attachment of a cell. We will prove for the piecewise linear case that the change in homotopy type is equivalent to the attachment of a handle, which in turn is homotopy equivalent to the attachment of a cell. Let us first recall from [44] some terminology regarding handles and why the attachment of a handle is homotopy equivalent to the attachment of a cell.

The following terminology for handles is taken from [44, p. 74]. A *handle* of index λ on a d -manifold M with boundary is a d -ball H with $M \cap H \subseteq \text{bd}(M)$ together with a homeomorphism $h: [-1, 1]^\lambda \times [-1, 1]^{d-\lambda}$ such that $h(\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda})$ agrees with $M \cap H$. We say that the union of $M \cup H$ is obtained from M by attaching the handle H . Slightly generalising the concept, let (H, H') be a manifold pair homeomorphic to the pair $([-1, 1]^d, \text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda})$ with an *attaching map* $g: H' \rightarrow \text{bd}(M)$. Then the quotient space obtained from M and H by identifying p and $g(p)$ for each $p \in H'$ is described as M with the handle H attached by g . Some subsets of a λ -handle H that are images under h of certain subsets of the hypercube $[-1, 1]^d$ have special names: $h([-1, 1]^\lambda \times \{0\})$ is called the *core* of the handle, $h(\{0\} \times [-1, 1]^{d-\lambda})$ its *co-core*, $h(\text{bd}([-1, 1]^\lambda) \times \{0\})$ its *attaching sphere*, $h(\{0\} \times \text{bd}([-1, 1]^{d-\lambda}))$ its *belt sphere*, $h(\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda})$ its *attachment tube*, and $h([-1, 1]^\lambda \times \text{bd}([-1, 1]^{d-\lambda}))$ its *belt tube*.

The reason why the attachment of a handle and the attachment of a cell are homotopy equivalent is sketched in [44, p. 83] as follows: There is a deformation retraction from a λ -handle H to the union of its core and its attaching tube. This deformation retraction induces a deformation retraction from $M \cup H$ to the union of M and the core, because the attaching tube is contained in the boundary of M . The core of a λ -handle is a λ -ball, and it is attached to the boundary of M in the attachment sphere, which is the boundary sphere of that

ball. This means that M with a λ -handle attached deformation retracts and hence is homotopy equivalent to M with a λ -cell attached, namely to M with the core of the handle attached.

With these prerequisites we prove the following piecewise linear analogue of Fact 1.4 and its counterpart in Milnor's book [36, Theorem 3.2, p. 14].

Theorem 5.2. *Let M be a combinatorial manifold and $f: |M| \rightarrow \mathbb{R}$ linear on cells. Assume that v is an interior non-degenerate critical vertex of index λ with $f(v) = h$ and v is the only vertex with a value in $[h - \epsilon, h + \epsilon]$ for some $\epsilon > 0$. Then the complex $M_{\leq h+\epsilon}$ is homotopy equivalent to $M_{\leq h-\epsilon}$ with a λ -handle attached, which is again homotopy equivalent to $M_{\leq h-\epsilon}$ with a λ -cell attached.*

Proof. As cited above from [44, p. 83], the result of the attachment of a λ -handle is homotopy equivalent to the result of the attachment of its core, which is a λ -cell. This proves the last assertion, that $M_{\leq h-\epsilon}$ with a λ -handle attached is homotopy equivalent to $M_{\leq h-\epsilon}$ with a λ -cell attached.

The proof idea for the main part of the statement is to show that the subdivision $M_{\leq h-\epsilon} \cup M_{\in [h-\epsilon, h+\epsilon]}$ of $M_{\leq h+\epsilon}$ collapses to $M_{\leq h-\epsilon}$ with a λ -handle attached, namely $M_{\leq h-\epsilon} \cup \text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$, where $\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$ forms the λ -handle attached to $M_{\leq h-\epsilon}$ with the common intersection $\text{st}(v)_{=h-\epsilon}$ of $M_{\leq h-\epsilon}$ and $\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$ as attaching tube. For verifying that $\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$ is a λ -handle, we map it in several steps homeomorphically to the ball $[-1, 1]^d$ such that $\text{st}(v)_{=h-\epsilon}$ corresponds to $\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda}$.

As a first step, we show that the subdivision $M_{\leq h-\epsilon} \cup M_{\in [h-\epsilon, h+\epsilon]}$ of $M_{\leq h+\epsilon}$ collapses to $M_{\leq h-\epsilon} \cup \text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$ using Corollary 2.44. We consider the subcomplex $K = \text{st}(v)$ and the value $a = h - \epsilon < h + \epsilon$. By our choice of ϵ , for every cell $S \in M \setminus \text{st}(v)$ exactly one of the following three cases occurs:

1. All vertices of S have f -value strictly less than $h - \epsilon$. This implies $S_{\leq h+\epsilon} = S_{\leq h-\epsilon}$ is contained in $M_{\leq h-\epsilon}$.
2. All vertices of S have f -value strictly greater than $h + \epsilon$. This implies $S_{\leq h+\epsilon} = \emptyset$ is contained in $M_{\leq h-\epsilon}$.
3. The cell S has vertices with f -value strictly less than $h + \epsilon$ and vertices with f -value strictly greater than $h + \epsilon$. This implies that $S_{\leq h+\epsilon}$ and $S_{=h+\epsilon}$ are canonical.

Hence $K = \text{st}(v)$ fulfils the requirements in Corollary 2.44, whose application yields that $M_{\leq h-\epsilon} \cup M_{\in [h-\epsilon, h+\epsilon]}$ collapses to $M_{\leq h-\epsilon} \cup \text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$. Since $M_{\leq h-\epsilon} \cup M_{\in [h-\epsilon, h+\epsilon]}$ is a subdivision of $M_{\leq h+\epsilon}$, this means that $M_{\leq h+\epsilon}$ is homotopy equivalent to $M_{\leq h-\epsilon}$ with $\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$ attached. It remains to show that the attachment of $\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$ to $M_{\leq h-\epsilon}$ is indeed the attachment of a λ -handle.

The common intersection of $\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$ and $M_{\leq h-\epsilon}$ is $\text{st}(v)_{=h-\epsilon}$, which is contained in the boundary of $M_{\leq h-\epsilon}$. For verifying that $\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$ is a λ -handle attached to $M_{\leq h-\epsilon}$ via their common intersection $\text{st}(v)_{=h-\epsilon}$ as attaching tube of the handle, we map the pair $(\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}, \text{st}(v)_{=h-\epsilon})$ in several steps homeomorphically to the pair $([-1, 1]^d, \text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda})$. One easily checks that the cone with apex v and base $\text{st}(v)_{=h-\epsilon} \cup \text{st}(v)_{=h+\epsilon} \cup \text{lk}(v)_{\in [h-\epsilon, h+\epsilon]}$ is a subdivision of $\text{st}(v)_{\in [h-\epsilon, h+\epsilon]}$. The base of the cone contains the subcomplexes $\text{st}(v)_{=h-\epsilon}$ and $\text{st}(v)_{=h+\epsilon}$. By Lemma 2.38, the former complex $\text{st}(v)_{=h-\epsilon}$

is combinatorially equivalent to $\text{lk}(v)_{\leq h-\epsilon}$ leaving $\text{lk}(v)_{=h-\epsilon}$ fixed, and the latter complex $\text{st}(v)_{=h+\epsilon}$ is combinatorially equivalent to $\text{lk}(v)_{\geq h+\epsilon}$ leaving $\text{lk}(v)_{=h+\epsilon}$ fixed. Hence we can replace the two subcomplexes by their combinatorially equivalent counterparts. We obtain that $\text{st}(v)_{\in[h-\epsilon, h+\epsilon]}$ is homeomorphic to the cone with apex v and base $\text{lk}(v)_{\leq h-\epsilon} \cup \text{lk}(v)_{\geq h+\epsilon} \cup \text{lk}(v)_{\in[h-\epsilon, h+\epsilon]}$ with the attaching tube $\text{st}(v)_{=h-\epsilon}$ corresponding to $\text{lk}(v)_{\leq h-\epsilon}$. Note that the cone with apex v and base $\text{lk}(v)_{\leq h-\epsilon} \cup \text{lk}(v)_{\geq h+\epsilon} \cup \text{lk}(v)_{\in[h-\epsilon, h+\epsilon]}$ is a subdivision of $\text{st}(v)$, so that we obtain the pair $(\text{st}(v), \text{lk}(v)_{\leq h-\epsilon})$ as homeomorphic pair for $(\text{st}(v)_{\in[h-\epsilon, h+\epsilon]}, \text{st}(v)_{=h-\epsilon})$.

For the next step, we observe the base $\text{lk}(v)_{\leq h-\epsilon} \cup \text{lk}(v)_{\geq h+\epsilon} \cup \text{lk}(v)_{\in[h-\epsilon, h+\epsilon]}$ is a subdivision of $\text{lk}(v)_{\leq h-\epsilon} \cup \text{lk}(v)_{\geq h-\epsilon}$. Since no vertex of $\text{lk}(v)$ has a value in the interval $[h-\epsilon, h]$, the complex $\text{lk}(v)_{\leq h-\epsilon} \cup \text{lk}(v)_{\geq h-\epsilon}$ is naturally combinatorially equivalent to $\text{lk}(v)_{\leq h} \cup \text{lk}(v)_{\geq h}$ by Lemma 2.35, and this combinatorial equivalence maps the subcomplex $\text{lk}(v)_{\leq h-\epsilon}$ to $\text{lk}(v)_{\leq h}$. Adding the apex v to this equivalence yields a combinatorial equivalence between $v(\text{lk}(v)_{\leq h-\epsilon} \cup \text{lk}(v)_{\geq h-\epsilon})$ and $v(\text{lk}(v)_{\leq h} \cup \text{lk}(v)_{\geq h})$ still mapping $\text{lk}(v)_{\leq h-\epsilon}$ to $\text{lk}(v)_{\leq h}$, in other words because both cones are subdivisions of $\text{st}(v)$, the pair $(\text{st}(v), \text{lk}(v)_{\leq h-\epsilon})$ is homeomorphic to the pair $(\text{st}(v), \text{lk}(v)_{\leq h})$.

The next homeomorphism is provided by the fact that v is a non-degenerate critical point of index λ . But instead of using the standard reference map f_λ^d from Definition 3.12 for characterising non-degenerate critical points by local equivalence, we use the map $h_\lambda^d = -\max_{i=1}^\lambda |X_i| + \max_{i=\lambda+1}^d |X_i|$ suggested by Eells and Kuiper [17] as reference map, which we showed to be equivalent to f_λ^d in Lemma 3.27. Therefore the triple (M, v, f) is locally equivalent to the triple $(\mathbb{R}^d, 0, h_\lambda^d)$. Recall from Theorem 3.11 that locally equivalent triples can be characterised by the property that any two link complexes contained in the stars of the considered points are homeomorphic in such a way that the shifted maps have the same sign at corresponding points. Choosing for our case $\text{lk}(v)$ as link complex for v and $\text{bd}([-1, 1]^d)$ as link complex for the origin in \mathbb{R}^d , we obtain a PL homeomorphism between the subdivision $\text{lk}(v)_{\leq h} \cup \text{lk}(v)_{\geq h}$ of $\text{lk}(v)$ and $\text{bd}([-1, 1]^d)$ such that $\text{lk}(v)_{\leq h}$ is mapped to $(\text{bd}([-1, 1]^d))_{h_\lambda^d \leq 0}$. Using the cone construction with apex v on the one hand, and the origin as apex on the other hand, this yields a PL homeomorphism from the cone with apex v and base $\text{lk}(v)_{\leq h} \cup \text{lk}(v)_{\geq h}$, which is a subdivision of $\text{st}(v)$ again, to the cone $0(\text{bd}([-1, 1]^d))$, which is a subdivision of $[-1, 1]^d$, such that $\text{lk}(v)_{\leq h}$ is mapped to $(\text{bd}([-1, 1]^d))_{h_\lambda^d \leq 0}$. In other words, the pair $(\text{st}(v), \text{lk}(v)_{\leq h})$ is homeomorphic to the pair $([-1, 1]^d, (\text{bd}([-1, 1]^d))_{h_\lambda^d \leq 0})$.

It remains to show that the level set $(\text{bd}([-1, 1]^d))_{h_\lambda^d \leq 0}$ that is homeomorphic with the attachment tube agrees with $\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda}$. Since the boundary of the hypercube $\text{bd}([-1, 1]^d)$ can be written as the union of $\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda}$ and $[-1, 1]^\lambda \times \text{bd}([-1, 1]^{d-\lambda})$, it suffices to check that the h_λ^d -values of points in $\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda}$ are non-positive and that the h_λ^d -values of points in $([-1, 1]^\lambda \times \text{bd}([-1, 1]^{d-\lambda})) \setminus (\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda})$ are positive. An element $(x_1, \dots, x_\lambda, x_{\lambda+1}, \dots, x_d)$ of $\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda}$ has the property that one of its first λ coordinates x_1, \dots, x_λ has absolute value 1. The remaining coordinates have absolute value at most one, which implies $h_\lambda^d(x_1, \dots, x_\lambda, x_{\lambda+1}, \dots, x_d) = -\max_{i=1}^\lambda |x_i| + \max_{i=\lambda+1}^d |x_i| \leq -1 + 1 = 0$. On the other hand, if $(x_1, \dots, x_\lambda, x_{\lambda+1}, \dots, x_d)$ does not belong to $\text{bd}([-1, 1]^\lambda) \times [-1, 1]^{d-\lambda}$ then none of its first λ coordinates has absolute

value 1, implying $\max_{i=1}^{\lambda} |x_i| < 1$. If moreover $(x_1, \dots, x_{\lambda}, x_{\lambda+1}, \dots, x_d)$ belongs to $[-1, 1]^{\lambda} \times \text{bd}([-1, 1]^{d-\lambda})$, then one of its last $d - \lambda$ coordinates has absolute value 1. Thus, an element $(x_1, \dots, x_{\lambda}, x_{\lambda+1}, \dots, x_d)$ lying in $([-1, 1]^{\lambda} \times \text{bd}([-1, 1]^{d-\lambda})) \setminus (\text{bd}([-1, 1]^{\lambda}) \times [-1, 1]^{d-\lambda})$ fulfils $h_{\lambda}^d(x_1, \dots, x_{\lambda}, x_{\lambda+1}, \dots, x_d) = -\max_{i=1}^{\lambda} |x_i| + \max_{i=\lambda+1}^d |x_i| > -1 + 1 = 0$. This completes the proof. \square

5.2 Consequences: Homotopy Type, Morse Inequalities, and Reeb Theorem

The remaining results of this chapter follow from Theorems 5.1 and 5.2 in much the same way as their smooth counterparts follow from Facts 1.2 and 1.4.

Homotopy Equivalence to CW-Complexes. We continue with an analogue of Fact 1.5, which states the result from [36, Theorem 3.5, p. 20] that a manifold with Morse function is homotopy equivalent to a CW-complex consisting of one λ -cell for each critical point of index λ .

Theorem 5.3. *Let M be a combinatorial manifold without boundary with a PL Morse function $f: |M| \rightarrow \mathbb{R}$ linear on cells of M . Assume general position for f , that is, all vertices of M have distinct f -values. Then M is homotopy equivalent to a CW-complex whose cells are in bijection with the critical points of f : one λ -cell for each non-degenerate critical point of index λ .*

Proof. The following proof is completely analogous to the one given in [36, Theorem 3.5, p. 20] for the smooth case, with the simplification that we assume that for each critical value there is only one critical vertex. We use the notion of a CW-complex in the sense that we briefly described in Section 1.1 following Forman [18, Section 1].

We order the vertices of M by increasing f -value to a sequence (v_1, \dots, v_n) with $f(v_1) < \dots < f(v_n)$. Let $h_i = f(v_i)$ for $1 \leq i \leq n$. We show by induction that for small enough $\epsilon > 0$, namely $\epsilon < \frac{1}{2} \cdot \min_{2 \leq i \leq n} (h_i - h_{i-1})$, the level set complex $M_{\leq h_i + \epsilon}$ is homotopy equivalent to a CW-complex W_i with one λ -cell for each non-degenerate critical vertex of index λ in (v_1, \dots, v_i) . Since $M_{\leq h_n + \epsilon} = M$, this gives the desired result.

By choosing a value $h_0 < h_1 - 2\epsilon$, we can start the induction with $i = 0$. Then $M_{\leq h_0 + \epsilon}$ is empty and the empty CW-complex W_0 has the desired properties.

Now assume for induction hypothesis that $M_{\leq h_{i-1} + \epsilon}$ is homotopy equivalent to a CW-complex W_{i-1} with one λ -cell for each non-degenerate critical vertex of index λ in (v_1, \dots, v_{i-1}) . If v_i is a regular vertex, then the interval $[h_{i-1} + \epsilon, h_i + \epsilon]$ contains one vertex, which is regular, and by Theorem 5.1, $M_{\leq h_i + \epsilon}$ is homeomorphic to $M_{\leq h_{i-1} + \epsilon}$, and hence homotopy equivalent to $W_i := W_{i-1}$.

It remains to consider the case that v_i is a critical vertex of index λ . We can again apply Theorem 5.1, this time for the interval $[h_{i-1} + \epsilon, h_i - \epsilon]$. It shows that $M_{\leq h_i - \epsilon}$ is homeomorphic to $M_{\leq h_{i-1} + \epsilon}$ and hence homotopy equivalent to W_{i-1} . By Theorem 5.2, the complex $M_{\leq h_i + \epsilon}$ is homotopy equivalent to $M_{\leq h_i - \epsilon}$ with a λ -cell attached. Now we use the following fact from [36, Lemma 3.7, p. 21]: Denote by $X \cup_g B^{\lambda}$ the space obtained from a space X by attaching the λ -cell B^{λ} along the attaching map $g: \text{bd}(B^{\lambda}) \rightarrow X$. Then any homotopy equivalence $\phi: X \rightarrow Y$ between spaces X and Y extends to a homotopy equivalence between $X \cup_g B^{\lambda}$ and $Y \cup_{\phi \circ g} B^{\lambda}$. For our situation, the fact

implies that the homotopy equivalence between $M_{\leq h_i - \epsilon}$ and W_{i-1} extends to a homotopy equivalence between $M_{\leq h_i - \epsilon}$ with a λ -cell attached and some CW-complex W_i that is obtained from W_{i-1} by a suitable attachment of a λ -cell. This CW-complex W_i consists of the cells from W_{i-1} and additionally one λ -cell corresponding to the critical point v_i . Composing the homotopy equivalences, we obtain that $M_{\leq h_i + \epsilon}$ is homotopy equivalent to W_i as desired. \square

Morse Inequalities and Reeb Theorem. The final two results of this chapter are the Morse inequalities and the Reeb theorem for PL Morse functions. They were stated previously by Kosinski [29, Theorems 2.3 and 2.4]. Analogous results for regular simplex-wise linear functions in the sense of Brehm and Kühnel [9] are also known: The Morse inequalities are stated in [9, Theorem 2] and [30, Lemma 7.3]; a variant of the Reeb theorem, where the conclusion for 4-dimensional manifolds asserts only a homotopy sphere, and not a PL sphere as asserted for the other dimensions, is proven in [30, Corollary 7.4].

From the representation of M by a homotopy equivalent CW-complex, we can now deduce the Morse inequalities. The smooth analogue was noted in Fact 1.6.

Theorem 5.4 (Morse Inequalities). *Let M be a combinatorial manifold without boundary with a PL Morse function $f: |M| \rightarrow \mathbb{R}$ linear on cells and in general position. Denote by c_λ the number of critical vertices of index λ in M , and by β_λ the λ -th Betti number of M . Then the following inequalities and equalities hold:*

1. *The Weak Morse Inequalities. For each λ :*

$$\beta_\lambda \leq c_\lambda$$

2. *An equality representing the Euler characteristic $\chi(M)$ of M in two ways:*

$$\chi(M) = \sum_{\lambda} (-1)^\lambda c_\lambda = \sum_{\lambda} (-1)^\lambda \beta_\lambda$$

3. *The Strong Morse Inequalities. For each λ :*

$$\sum_{i=0}^{\lambda} (-1)^i \beta_{\lambda-i} \leq \sum_{i=0}^{\lambda} (-1)^i c_{\lambda-i}$$

Proof. As Forman states in [18, Theorems 1.7 and 1.8], the equalities and inequalities hold for any CW-complex W , if c_λ denotes the number of λ -cells in W and β_λ the λ -th Betti number of W . By the previous Theorem 5.3, there is a CW-complex W whose number of λ -cells agrees with the number of critical points of index λ in M and which is homotopy equivalent to M , implying that the Betti numbers of W and M agree. From these observations the result follows immediately. \square

The final result of this chapter is the PL analogue of the Reeb theorem cited in Fact 1.7 from [36, Theorem 4.1, p. 25].

Theorem 5.5 (Reeb theorem). *Let M be a combinatorial manifold without boundary with a PL Morse function $f: |M| \rightarrow \mathbb{R}$ linear on cells and in general position. Assume that f has exactly two critical vertices. Then M is a PL sphere.*

Proof. The proof is taken almost literally from its smooth version in [36, Theorem 4.1, p. 25].

The two critical vertices must be the vertices where minimum and the maximum of f on M are attained. Let v be the vertex with minimum f -value and u be the vertex with maximum f -value. For sufficiently small ϵ the level set complex $M_{\leq f(v)+\epsilon}$ agrees with $\text{st}(v)_{\leq f(v)+\epsilon}$ and with the cone with apex v and base $\text{st}(v)_{=f(v)+\epsilon}$. By Lemma 2.38, the base $\text{st}(v)_{=f(v)+\epsilon}$ is combinatorially equivalent to $\text{lk}(v)_{\geq f(v)+\epsilon}$. But $\text{lk}(v)_{\geq f(v)+\epsilon}$ agrees with $\text{lk}(v)$ for sufficiently small ϵ , since $f(v)$ is the minimum value of f on M . Moreover, the cone with apex v and base $\text{lk}(v)$ is simply the PL ball $\text{st}(v)$, showing that $M_{\leq f(v)+\epsilon}$ is a PL ball. By a similar argument $M_{\geq f(u)-\epsilon}$ is a PL ball as well. Furthermore $M_{\leq f(v)+\epsilon}$ is PL-homeomorphic to $M_{\leq f(u)-\epsilon}$ by Theorem 5.1, since all remaining vertices with an f -value in $[f(v) + \epsilon, f(u) - \epsilon]$ are regular. Hence M is the union of two PL balls $M_{\leq f(v)+\epsilon}$ and $M_{\geq f(u)-\epsilon}$ that intersect in their common boundary $M_{=f(u)-\epsilon}$, showing that M is a PL sphere. \square

Chapter 6

Computational Aspects

The last chapter addresses two selected topics that arise when we consider PL Morse theory from a computational point of view. The first section deals with the question whether regular vertices can be recognised algorithmically. We obtain positive solutions for low dimensions, but for higher dimensions, the answer is more difficult. When arbitrary complexes are allowed as input, it is possible to show undecidability of the problem as stated in Corollary 6.6. The complexity remains open when the problem is restricted to inputs that are known to be combinatorial manifolds.

In the other section of this chapter, we use our results on the existence of f -level-preserving PL isotopies to sketch the design of algorithms in low dimensions that compute such isotopies. The aim is to compute combinatorially equivalent subdivisions of $M_{=a} \times [a, b]$ and $M_{\in[a, b]}$ that represent the isotopy $\phi: |M|_{=a} \times [a, b] \rightarrow |M|_{\in[a, b]}$ for an interval $[a, b]$ containing no critical values. We derive upper bounds on the number of cells in the combinatorially equivalent subdivisions produced by the outlined algorithms.

6.1 Recognising Regular Vertices

This section addresses the question whether regular points can be recognised algorithmically. The difficulty of this computational problem depends on the dimension of the considered domain. For combinatorial manifolds of dimension at most 3 and maps in general position, Chiang, Lenz, Lu, and Rote [13] give characterisations for regular points that lead to simple recognition algorithms which require only counting the number of connected components of the lower level link and the upper level link of the considered vertex. They also propose a characterisation for 4-dimensional manifolds, that is easy to check algorithmically. For higher dimensions they suspect NP-hardness or even undecidability of the problem. We shortly review the positive results in low dimensions, in particular, we verify the suggested characterisation for 4-dimensional domains, which is obtained by applying Lemma 6.1 to the 4-dimensional case. For complexes of dimension at least 6 we obtain an undecidability result in Corollary 6.6. This result however is somewhat unsatisfactory because it relies on a reduction with the property that all negative instances it produces consist of a vertex whose link is not a sphere. So the reduction rather re-establishes the undecidability

of the class of manifolds, which leads to the question whether regular and critical points become algorithmically distinguishable if the domain is asserted to be a manifold. This question remains open and is related to the problem that we already encountered in Subsection 3.2.5, when we compared the notion of criticality defined by Brehm and Kühnel [9] to our notion, namely the question whether there are homology $(d - 1)$ -balls different from the standard ball that embed into the standard $(d - 1)$ -sphere.

6.1.1 Algorithmic Solutions in Low Dimensions

Regular Vertex Recognition by Sphere Recognition Algorithms and Schoenflies Conjecture. In Chapter 4 we discussed several equivalent characterisations for critical points. The characterisations that are best suited for an algorithmic approach are those that we proved in the first part of that chapter: Lemma 4.1 characterised regular points x by requiring that $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is an unknotted sphere or ball pair and that the sets $|\text{lk}(x)|_{>f(x)}$ and $|\text{lk}(x)|_{<f(x)}$ are both non-empty. By Corollary 4.2, the second condition can be omitted for maps in general position except for 1-dimensional domains. Depending on the status of the Schoenflies conjecture, in the cases described in Lemma 4.3 a sphere pair $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is automatically unknotted, so that it is unnecessary to check this property explicitly. Consequently algorithms for sphere recognition allow to design algorithms for the recognition of regular points, whenever we are in a situation where the Schoenflies conjecture holds. Recall from Subsection 1.2.11 that for low dimensions $d \leq 3$, algorithms for d -sphere recognition exist. Combining this fact with the truth of the d -dimensional PL Schoenflies conjecture for $d \leq 3$ yields recognition algorithms for regular interior points in dimension at most 4.

Lemma 6.1. *Let $f: |M| \rightarrow \mathbb{R}$ be linear on cells of a simplicial complex M . For low dimensions $d \leq 4$ and each point $x \in |M|$, it can be tested algorithmically whether x is a regular interior point of f in dimension d .*

Proof. By Lemma 4.1, a regular interior point in dimension d is characterised by the requirements that $\text{lk}(x)$ contains vertices u_- and u_+ with $f(u_-) < f(x)$ and $f(u_+) > f(x)$ and that $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is an unknotted $(d - 1, d - 2)$ -sphere pair. Checking the first requirement is trivial. Then we can apply the existing recognition algorithms for $(d - 1)$ -spheres and $(d - 2)$ -spheres for $d \leq 4$ for verifying that $\text{lk}(x)$ is a $(d - 1)$ -sphere and that $\text{lk}(x)_{=f(x)}$ is a $(d - 2)$ -sphere. If both tests succeed, the truth of the $(d - 1)$ -dimensional PL Schoenflies conjecture for $d - 1 \leq 3$ asserts that the $(d - 1, d - 2)$ -sphere pair $(\text{lk}(x), \text{lk}(x)_{=f(x)})$ is unknotted. \square

As observed in Corollary 4.2, the test for the vertices u_- and u_+ can be omitted for $d > 1$ if f is in general position. If M is asserted to be a combinatorial d -manifold and x an interior point of $|M|$, the definition of a manifold confirms that $\text{lk}(x)$ is a $(d - 1)$ -sphere, so that checking this property becomes unnecessary as well. Furthermore we proved in Lemma 3.22 that non-vertices are always regular if M is a manifold and f is in general position, so that only vertices have to be checked for regularity under these assumptions. These observations verify the correctness of the characterisation suggested in [13] for regular interior vertices of 4-dimensional manifolds with a map in general position: The

suggested characterisation requires that the level set $\text{lk}(v)_{=f(v)}$ separates the link $\text{lk}(v)$ into two 3-balls intersecting in the 2-sphere $\text{lk}(v)_{=f(v)}$. The authors also remark that the validity of the 3-dimensional PL Schoenflies conjecture implies that $\text{lk}(v)_{=f(v)}$ being a 2-sphere is already sufficient to fulfil the requirements, and that this sufficient condition can be checked algorithmically. When we apply the preceding lemma and the above observations for the case $d = 4$, we can conclude that an interior vertex v of a combinatorial 4-manifold with a map in general position is regular if and only if $\text{lk}(v)_{=f(v)}$ is a 2-sphere, and that such a 2-sphere separates the link $\text{lk}(v)$ into two 3-balls because the 3-dimensional PL Schoenflies conjecture is true. In other words, we obtain exactly the regularity condition proposed in [13].

Let us briefly review some related observations by Chiang, Lenz, Lu, and Rote [13] concerning the classification of vertices in 2- and 3-dimensional domains. It is assumed that M is a combinatorial manifold with a map in general position. Therefore all non-vertices are regular and only vertices require further attention. The vertices are classified in [13] as follows: For a vertex v , consider the graph formed by the 1-skeleton of its link. The vertices u in the link with $f(u) < f(v)$ induce a sub-graph $N_-(v)$, a sub-graph $N_+(v)$ is induced analogously by the vertices u with $f(u) > f(v)$. One easily checks that $N_-(v)$ and $N_+(v)$ are the 1-skeleta of the subcomplexes $\text{scp}_{\leq f(v)}(\text{lk}(v))$ and $\text{scp}_{\geq f(v)}(\text{lk}(v))$ respectively. The vertices are classified according to the numbers $C_-(v) = \beta_0(N_-(v))$ and $C_+(v) = \beta_0(N_+(v))$, which denote the number of connected components of $N_-(v)$ and $N_+(v)$. Since a simplicial complex has the same 0-th Betti number as its 1-skeleton, these numbers can be expressed equivalently as $C_-(v) = \beta_0(\text{scp}_{\leq f(v)}(\text{lk}(v)))$ and $C_+(v) = \beta_0(\text{scp}_{\geq f(v)}(\text{lk}(v)))$. Moreover, the collapses from $\text{lk}(v)_{\leq f(v)}$ to $\text{scp}_{\leq f(v)}(\text{lk}(v))$ and from $\text{lk}(v)_{\geq f(v)}$ to $\text{scp}_{\geq f(v)}(\text{lk}(v))$ following from Lemma 2.46 and the homotopy equivalences induced by these collapses yield the alternative equalities $C_-(v) = \beta_0(\text{lk}(v)_{\leq f(v)})$ and $C_+(v) = \beta_0(\text{lk}(v)_{\geq f(v)})$.

Counting Connected Components of Upper and Lower Level Links.

The aim of the studies in [13] is the design of an algorithm for the construction of contour trees. Important for this purpose are only those critical vertices where the number of connected components of the level sets $|M|_{<h}$, $|M|_{=h}$, or $|M|_{>h}$ changes, when the level h passes across the f -value of that point. The authors show that a vertex v fulfilling $C_-(v) = C_+(v) = 1$ does not cause a change in the number of connected components of the level sets [13, Lemma 5]. Vertices not fulfilling this property are called *component-critical*, and this property is easy to check algorithmically. Clearly, the component-critical vertices form a subset of the critical vertices, and regular vertices are not component-critical. The case analysis in [13] for 2- and 3-dimensional manifolds establishes the following observations: In dimension 2, a vertex is regular if and only if it is not component-critical [13, Theorem 3]. This applies to interior and to boundary points. In dimension 3, there are critical points that are not component-critical, but such vertices are necessarily boundary points which are local extrema with respect to the boundary [13, Theorem 1]. In other words, interior points are regular if and only if they are not component-critical, whereas boundary points are regular if and only if they are not component-critical with respect to the whole manifold and with respect to the boundary. These observations imply

that in dimensions 2 and 3, regular points of manifolds with a map in general position, including regular boundary points, can be recognised by counting the number of connected components of the upper and lower level links $\text{lk}_M(v)_{\leq f(v)}$, $\text{lk}_M(v)_{\geq f(v)}$, $\text{lk}_{\text{bd}(M)}(v)_{\leq f(v)}$, and $\text{lk}_{\text{bd}(M)}(v)_{\geq f(v)}$. It is remarked in the appendix of [13] that a non-degenerate critical vertex of index 2 in a 4-dimensional manifold is not component-critical, so that counting connected components does not suffice for recognising regular vertices in dimension 4 and above.

Related Algorithmically Solvable Problems. In the remainder, we omit the treatment of boundary points and focus on combinatorial manifolds without boundary. Moreover we concentrate on the classification of vertices. When the considered point is not a vertex, we can turn it into a vertex by performing an elementary starring at the point. This operation might ruin a general position assumption, but under such an assumption non-vertices are already regular.

Recall that a simplicial complex is a combinatorial manifold if and only if the link of every vertex is a combinatorial sphere. For a single vertex, this requirement is also a necessary condition for the vertex being regular or non-degenerate critical. Consequently, we obtain as necessary condition for a map linear on cells of a simplicial complex being a PL Morse function that the complex is required to be a combinatorial manifold. The sphere recognition algorithms in low dimensions provide a method to recognise combinatorial manifolds of dimension at most 4. When the given complex is known to be a manifold for some reason, or even f is known to be a PL Morse function, some problems concerning the classification of vertices can be solved algorithmically even in higher dimensions. The results of this kind listed below are trivial or obvious consequences of facts we observed earlier.

Lemma 6.2. *Let $f: |M| \rightarrow \mathbb{R}$ be linear on cells of a simplicial complex M .*

If $\dim(M) \leq 4$, then it can be tested algorithmically whether M is a combinatorial manifold. Furthermore regular vertices, non-degenerate critical vertices of index 0, and non-degenerate critical vertices of index $\dim(M)$ can be recognised algorithmically.

If M is known to be a combinatorial manifold, then non-degenerate critical vertices of index 0, and non-degenerate critical vertices of index $\dim(M)$ can be recognised algorithmically for any dimension of M .

If $f: |M| \rightarrow \mathbb{R}$ is known to be a PL Morse function, then it can be computed whether a vertex v is regular or critical, and in the latter case, its index can be computed.

Proof. The complex M is a combinatorial d -manifold, if and only if the link of every vertex is a combinatorial $(d-1)$ -sphere. Algorithms for recognising combinatorial $(d-1)$ -spheres are known to exist for $d-1 \leq 3$. This shows that manifolds of dimension at most 4 can be recognised algorithmically.

We observed already in Lemma 6.1 that regular vertices of dimension d at most 4 are algorithmically recognisable: It suffices to check that $\text{lk}(v)$ is a $(d-1)$ -sphere and contains vertices u_+ and u_- with $f(u_+) > f(v)$ and $f(u_-) < f(v)$ and that $\text{lk}(v)_{=f(v)}$ is a $(d-2)$ -sphere. The existence of the required vertices can be checked easily, for the other two conditions we use again the existing algorithms for recognising low-dimensional spheres.

A non-degenerate critical vertex of index 0 is characterised by the condition that its link is a $(\dim(M) - 1)$ -sphere and that its value $f(v)$ is strictly smaller than all f -values of its neighbouring vertices in the link. Non-degenerate critical vertices of index $\dim(M)$ are characterised analogously with the word “smaller” replaced by “larger”. Comparison of the f -values is easy, and the condition that the link is a sphere can be checked for $\dim(M) \leq 4$, which completes the proof of the first claim, or follows from the assumption that M is a combinatorial manifold in the second claim.

If f is known to be a PL Morse function, then the triple (M, v, f) is locally equivalent to the triple $(\mathbb{R}^d, 0, \pi_1^d)$ in the case that v is regular or to one of the triples $(\Omega^d, 0, f_\lambda^d)$ with $0 \leq \lambda \leq d = \dim(M)$ in the case that v is a non-degenerate critical vertex of index λ . When we take S_1^{d-1} as link complex for the origin, we can conclude that $\text{lk}(v)_{\leq f(v)}$ is homeomorphic to one of the standard lower level links, namely to $(S_1^{d-1})_{\pi_1^d \leq 0}$ for regular v or to $(S_1^{d-1})_{f_\lambda^d \leq 0}$ for a critical vertex v of index λ . Recall that we verified in the proof of Lemma 3.17 that $(S_1^{d-1})_{\pi_1^d \leq 0}$ is a ball and that $(S_1^{d-1})_{f_\lambda^d \leq 0}$ is homotopy equivalent to a $(\lambda - 1)$ -sphere, with the empty set being treated as (-1) -sphere. Moreover we observed that the different types of standard lower level links can be distinguished from each other by their Betti numbers. Thus, for classifying the vertices under the assumption that f is a PL Morse function, it suffices to compute the Betti numbers of the lower level link $\text{lk}(v)_{\leq f(v)}$ or of the subcomplex $\text{scp}_{\leq f(v)}(\text{lk}(v))$, which yields the same results because the former collapses to the latter. If the Betti numbers agree with those of a single point, then v is regular, and if they agree with those of a $(\lambda - 1)$ -sphere for some $\lambda \geq 0$, then v is a non-degenerate critical vertex of index λ . \square

6.1.2 Undecidability in Higher Dimensions

Turing Equivalence of Regular Vertex Recognition and Sphere Recognition. Now we investigate the difficulties that we encounter when we want to recognise regular vertices in higher dimensions. Let us first argue why the recognition problem of regular vertices in dimension $d > 1$, at least in its most general form, and the problem of recognising $(d - 1)$ -spheres have a similar degree of computational difficulty. In fact, we will show that both problems are Turing reducible to each other.

Lemma 6.3. *The problem of deciding whether for a given simplicial complex M with a map f linear on cells in general position and a vertex v in M , the vertex v is a regular vertex in dimension d or not is Turing reducible to the $(d - 1)$ -sphere recognition problem in the class of simplicial complexes.*

Proof. Recall from Subsection 1.2.11 that sphere recognition and ball recognition in the same dimension are computationally equivalent. Thus, the following explanation on how to design an algorithm for recognising regular vertices in dimension d when an algorithm for recognising $(d - 1)$ -balls is available yields a Turing reduction from the former problem to the latter.

The correctness of the algorithm we want to design relies mainly on two facts. The first fact is the characterisation of regular points given in Lemma 4.1: A vertex v is regular if and only if the pair $(\text{lk}(v), \text{lk}(v)_{=f(v)})$ is an unknotted $(d - 1, d - 2)$ -sphere pair and $|\text{lk}(v)|_{>f(v)}$ and $|\text{lk}(v)|_{<f(v)}$ are both non-empty.

The second fact is the sufficient condition for unknotted sphere pairs that follows from the results cited from [31] in Subsection 1.2.12. If upper and lower level link $\text{lk}(v)_{\geq f(v)}$ and $\text{lk}(v)_{\leq f(v)}$ are both combinatorial $(d-1)$ -balls whose boundaries coincide with each other and with the common intersection $\text{lk}(v)_{=f(v)}$ of the balls, then $(\text{lk}(v), \text{lk}(v)_{=f(v)})$ is unknotted and $|\text{lk}(v)|_{>f(v)}$ and $|\text{lk}(v)|_{<f(v)}$ are both non-empty because a ball is always a proper superset of its boundary. Since the boundary of a combinatorial $(d-1)$ -ball is generated by those $(d-2)$ -faces that have only one instead of two $(d-1)$ -dimensional co-faces, the condition concerning the boundaries can be checked algorithmically and we are left with the problem of recognising $(d-1)$ -balls. \square

For a Turing reduction in the other direction, we show that the problem of recognising $(d-1)$ -spheres is many-one reducible to the problem of recognising regular vertices in dimension d . As briefly indicated earlier, the reduction has the property that it maps positive instances of the sphere problem to a triple (M, v, f) such that M is a combinatorial d -manifold without boundary with a map f in general position and a regular vertex v , whereas in a triple (M, v, f) resulting from a negative instance of the sphere problem, not only the vertex v is not a regular interior point, but also the complex M fails to be a combinatorial manifold without boundary because the link of v is not a sphere.

Lemma 6.4. *The problem of recognising combinatorial $(d-1)$ -spheres can be reduced to the problem whether a given vertex of a simplicial complex with a map linear on cells in general position is a regular interior vertex in dimension d or not.*

Proof. Let K be an instance of the sphere recognition problem. We want to construct from K an instance (M, v, f) of the regular vertex recognition problem consisting of a simplicial complex M , a vertex v in M , and a map f linear on cells of M in general position defined by distinct values at the vertices of M . The proof idea is to design the triple (M, v, f) in such a way that a subdivision of K is the link of the vertex v in question and the chosen values for the map f ensure that v is regular if K is a sphere.

As a first step of the construction, select a maximal simplex S of K and perform an elementary starring at a point w in the interior of the selected simplex. Let M be the suspension of the obtained subdivision L of K . One of the apices of the suspension is the vertex v , the other apex is denoted by v' . Define the values for the map f as follows: For the vertices v , v' , and w , the chosen values are $f(v) = 0$, $f(v') = -\epsilon$ for some small positive ϵ , and $f(w) = -1$. The remaining vertices of M are vertices of the original complex K . Assign distinct small positive values to the vertices of the selected maximal simplex S , and distinct large positive values for the remaining vertices of K . If the input instance K is a simplicial complex, the above construction clearly yields a simplicial complex M with a map f linear on cells in general position and a vertex v in M . For verifying that this construction reduces the sphere recognition problem to the regular vertex problem, we prove the following two implications: If the complex K is a combinatorial $(d-1)$ -sphere, then M is a combinatorial d -manifold without boundary and v is regular. If K is not a $(d-1)$ -sphere, then the link of v in M is not a $(d-1)$ -sphere, which means in particular that M is not a combinatorial manifold without boundary and that v is not a regular interior vertex.

The latter implication is easy to prove. The link of an apex in a suspension $\text{susp}(X)$ is the original complex X , in our case the link of v in M is the subdivision L of K . Hence $\text{lk}_M(v)$ is a $(d-1)$ -sphere if and only if K is a $(d-1)$ -sphere, yielding the desired implication as one direction of the equivalence. For the former implication, note first that a suspension can be described equivalently as a join with a 0-sphere, and recall from Fact 1.38 that a join of a $(d-1)$ -sphere and a 0-sphere is a d -sphere. Therefore the constructed complex M is a d -sphere if K is a $(d-1)$ -sphere. We saw already above that $\text{lk}_M(v)$ agrees with L and is a $(d-1)$ -sphere. For showing that v is a regular interior point, we verify that $\text{lk}_M(v)_{=0}$ is a combinatorial $(d-2)$ -sphere and that one of the components of the induced decomposition into $\text{lk}_M(v)_{\leq 0}$ and $\text{lk}_M(v)_{\geq 0}$, namely the component $\text{lk}_M(v)_{\leq 0}$, is a combinatorial $(d-1)$ -ball bounded by $\text{lk}_M(v)_{=0}$. By the facts cited in Subsection 1.2.12, this asserts that the pair $(\text{lk}_M(v), \text{lk}_M(v)_{=0})$ is an unknotted $(d-1, d-2)$ -sphere pair and consequently v is a regular vertex.

Consider the lower level link $\text{lk}_M(v)_{\leq 0}$. By the choice of values for f , the vertex w with $f(w) = -1$ is the only vertex in L with non-positive f -value. Hence $\text{lk}_M(v)_{\leq 0} = L_{\leq 0}$ agrees with $\text{st}_L(w)_{\leq 0}$, and $\text{lk}_M(v)_{=0} = L_{=0}$ agrees with $\text{st}_L(w)_{=0}$. Moreover, $\text{st}_L(w)_{\leq 0}$ is the cone with apex w and base $\text{st}_L(w)_{=0}$. By Lemma 2.38, the level set complex $\text{st}_L(w)_{=0}$ forming the base is combinatorially equivalent to the level set complex $\text{lk}_L(w)_{\geq 0}$. But $\text{lk}_L(w)$ is the boundary of the selected simplex S whose vertices all have positive f -values. Thus $\text{lk}_L(w)_{\geq 0}$ agrees with the link $\text{lk}_L(w)$ itself, and this link is a combinatorial $(d-2)$ -sphere because it is the boundary of a $(d-1)$ -simplex. We have obtained a combinatorial equivalence between $\text{st}_L(w)_{=0}$ and $\text{lk}_L(w)$ which extends to a combinatorial equivalence between the cones $w\text{st}_L(w)_{=0}$ and $w\text{lk}_L(w) = \text{st}_L(w)$. Since the latter cone is a combinatorial $(d-1)$ -ball bounded by the base $\text{lk}_L(w)$, we can conclude that the cone $w\text{st}_L(w)_{=0}$ is a combinatorial $(d-1)$ -ball bounded by its base $\text{st}_L(w)_{=0}$ as well and that the boundary $\text{st}_L(w)_{=0} = \text{lk}_M(v)_{=0}$ is a combinatorial $(d-2)$ -sphere. This completes the proof. \square

It is possible to modify the reduction by a preprocessing step in such a way that it avoids output instances where the link of v is a combinatorial ball: Since such output instances could only occur if K is a combinatorial ball, it suffices to check in the preprocessing whether K is a simplicial complex, whether all maximal simplices of K have dimension $d-1$, and whether all $(d-2)$ -simplices have exactly two $(d-1)$ -dimensional co-faces. Any combinatorial $(d-1)$ -sphere passes these tests, but for combinatorial $(d-1)$ -balls and balls and spheres of other dimensions at least one of the tests fails. By outputting a suitable negative instance if one of the tests fails, we can ensure that all negative instances of the $(d-1)$ -sphere recognition problem are reduced to instances where M is not a combinatorial manifold, neither with nor without boundary, and v is not regular, neither a regular interior point nor a regular boundary point.

Corollary 6.5. *Regular vertex recognition in dimension d for a given triple (M, v, f) consisting of a simplicial complex M , a map f linear on cells of M in general position, and a vertex v in M is Turing equivalent to $(d-1)$ -sphere recognition for a given simplicial complex.*

Proof. The Turing reductions are given in Lemmata 6.3 and 6.4. \square

Undecidability Result. Novikov's result cited in Fact 1.40 that d -sphere recognition in dimension $d \geq 5$ is undecidable together with the above reduction imply immediately the undecidability of the problem of recognising regular vertices in dimension $d \geq 6$ in simplicial complexes with maps in general position.

Corollary 6.6. *Given a simplicial complex with a map linear on cells in general position, a vertex v of the simplicial complex, and a dimension $d \geq 6$, it is undecidable whether v is a regular vertex in dimension d .*

Proof. It is undecidable whether a given simplicial complex is a $(d-1)$ -sphere for $d-1 \geq 5$ by Novikov's result. The reduction described in Lemma 6.4 reduces this problem to the question whether a given vertex in a simplicial complex with a map linear on cells in general position is a regular point in dimension d . Thus regular vertex recognition in dimension $d \geq 6$ is undecidable as well. \square

Since the reduction produces non-manifold instances, the hardness result applies only if the regular vertex recognition problem is considered on a class of inputs that contains such instances. It is still open whether this problem becomes decidable or remains hard in higher dimensions if it is restricted to a class of inputs where M is a combinatorial d -manifold or at least $\text{lk}(v)$ is asserted to be a $(d-1)$ -sphere. In the following, we will refer to this question as the regular vertex recognition problem for d -manifolds. The assertion that $\text{lk}(v)$ is a $(d-1)$ -sphere leaves as necessary and sufficient condition for v being regular that $\text{lk}(v)_{=f(v)}$ is a $(d-2)$ -sphere and its embedding in $\text{lk}(v)$ is unknotted. This means that the regular vertex recognition problem for d -manifolds becomes decidable if an algorithm for the following problem exists: Given a simplicial complex K that is asserted to be a $(d-1)$ -sphere and a subcomplex L of K , decide whether (K, L) is an unknotted $(d-1, d-2)$ -sphere pair. Note that the proof of Lemma 6.3 also shows that the existence of a $(d-1)$ -sphere recognition algorithm allows to decide whether a given pair (K, L) is an unknotted $(d-1, d-2)$ -sphere pair even without the assertion of K being a $(d-1)$ -sphere.

Is it also possible to reduce the decision problem whether a pair (K, L) , where K is asserted to be a $(d-1)$ -sphere, is an unknotted $(d-1, d-2)$ -sphere pair to the regular vertex recognition problem for d -manifolds? A simple approach for designing such a reduction would aim for the construction of an instance where the given pair (K, L) occurs as embedding of $L = \text{lk}(v)_{=f(v)}$ into $K = \text{lk}(v)$ for some vertex v in the constructed d -manifold M . A suitable candidate for M is the suspension of the given $(d-1)$ -sphere K with one of the apices playing the role of the vertex v . Then the remaining task is to define a map f , we can assume $f(v) = 0$, in such a way that $\text{lk}(v)_{=0}$ agrees with L and the map f has opposite signs on the two parts that $\text{lk}(v)$ is separated into by $\text{lk}(v)_{=0}$. As long as we are not requiring f to be in general position, it seems possible to achieve such a construction. But if we want to show that the problem of recognising unknotted $(d-2)$ -spheres in asserted $(d-1)$ -spheres is equivalent to the vertex recognition problem for d -manifolds not only for arbitrary maps f but also when restricted to maps in general position, then a suitable reduction seems to be harder to find.

The Status of the Regular Vertex Recognition Problem Depending on Dimension. Let us study the consequences of the above results depending

on the dimension of the complex M containing the vertex v for which we want to decide whether it is regular or not.

From Dimension 1 to Dimension 4. In dimension 1, it suffices to check whether v has exactly two adjacent vertices, one with smaller and one with larger f -value, for verifying that v is regular. In dimension 2, we can check at first whether the link of v is a 1-sphere, in other words, expressed in the language of graphs, whether $\text{lk}(v)$ is a simple cycle. This test is unnecessary if the given complex is already known to be a 2-manifold. If the test succeeds or is unnecessary, it remains to check that $\text{lk}(v)_{=f(v)}$ is a 0-sphere, which means that it consists of exactly two vertices. For 3-dimensional domains, we can check whether $\text{lk}(v)$ is a 2-sphere because algorithms for 2-sphere recognition can be derived from the classification of 2-dimensional manifolds. Again, this check can be omitted if M is known to be a 3-manifold. In either case, it remains to verify that $\text{lk}(v)_{=f(v)}$ is a 1-sphere for recognising v as regular vertex. Dimension 4 still follows the same pattern as the previous ones. If M is not already known to be a 4-manifold, we can choose one of the existing 3-sphere recognition algorithms and check whether $\text{lk}(v)$ is a 3-sphere. Afterwards we only have to check whether $\text{lk}(v)_{=f(v)}$ is a 2-sphere.

Dimension 5. Starting from dimension 5, the situation changes. According to Corollary 6.5, the hardness of the 5-dimensional case depends on the hardness of recognising 4-spheres. It is unknown whether an algorithm for recognising 4-spheres exists or the problem is undecidable. Since $\text{lk}(v)$ being a 4-sphere is a necessary condition for v being a regular vertex in dimension 5, it is open whether we can test this condition. For the same reason it is also unknown whether the class of 5-manifolds is decidable. We can check however the necessary condition requiring $\text{lk}(v)_{=f(v)}$ to be a 3-sphere using a suitable recognition algorithm for 3-spheres. But even if it is asserted that $\text{lk}(v)$ is a 4-sphere, for example because M is known to be a 5-manifold for some reason, checking this condition is not sufficient, because the Schoenflies conjecture for embeddings of the 3-sphere into the 4-sphere is open, even in the locally flat case that would apply if we assumed general position. If the Schoenflies conjecture in this dimension is true, in general or at least in the locally flat case, and $\text{lk}(v)$ is asserted to be a 4-sphere, then checking that $\text{lk}(v)_{=f(v)}$ is a 3-sphere is enough, either for arbitrary f or at least if f is in general position. Moreover, if an algorithm for 4-sphere recognition exists, then regular vertices in dimension 5 can be recognised using the reduction from Lemma 6.3, independently from the status of the Schoenflies conjecture: We only have to check that $\text{lk}(v)_{\leq f(v)}$ and $\text{lk}(v)_{\geq f(v)}$ are both combinatorial 4-balls and that $\text{lk}(v)_{=f(v)}$ is the boundary of both balls. On the other hand, if 4-sphere recognition is undecidable, the reduction from Lemma 6.4 shows that the problem of recognising regular vertices in dimension 5 is in general undecidable as well.

We may ask what happens if the Schoenflies conjecture for embeddings of the 3-sphere into the 4-sphere turns out to be false. The next question that arises naturally under this premise is whether we can recognise unknotted embeddings of the 3-sphere into the 4-sphere although knotted embeddings are known to exist. If the answer to that question is positive, we can solve the regular vertex recognition problem for 5-manifolds. This holds even if 4-sphere recognition and

consequently the general problem of recognising regular points in dimension 5 for arbitrary given simplicial complexes M is undecidable. If the answer is negative and unknotted (4,3)-sphere pairs are not algorithmically recognisable, we can apply the reduction sketched above which reduces this problem to the regular vertex recognition problem for 5-manifolds. This would yield undecidability of the latter problem, as long as we are not requiring f to be in general position. If the Schoenflies conjecture turned out to be false in general, but true for the locally flat case, such a result would be satisfactory, because then regular vertex recognition for 5-manifolds would be undecidable for arbitrary maps f but decidable for maps in general position. But if the Schoenflies conjecture were false even in the locally flat case, the question arises whether we can find a reduction where the constructed map is in general position, which would render vertex recognition for 5-manifolds undecidable even if the problem is restricted to maps in general position. A solution for this task however seems to be harder to achieve.

Dimension 6. In dimension 6, Novikov's result that 5-spheres are not algorithmically recognisable implies that it is undecidable whether a given complex is a 6-manifold and that the general problem of recognising regular vertices in dimension 6 is undecidable. If f is in general position and an algorithm for recognising 4-spheres exists, regular vertices in 6-manifolds can be recognised by checking whether $\text{lk}(v)_{=f(v)}$ is a 4-sphere because the Schoenflies conjecture for locally flat embeddings of 4-spheres into 5-spheres is true. If the Schoenflies conjecture holds for arbitrary embeddings, the recognition works as well for arbitrary f . In the case that recognising 4-spheres is undecidable, the situation in dimension 6 is similar to that of dimensions $d \geq 7$ considered next.

Dimension 7 and Above. For dimension $d \geq 7$ the general problem of recognising regular vertices is undecidable. When the problem is restricted to d -manifolds and we assume general position, then the remaining necessary and sufficient condition for v being regular is that $\text{lk}(v)_{=f(v)}$ is a $(d - 2)$ -sphere. The general problem of recognising $(d - 2)$ -spheres is undecidable for $d \geq 7$. Whether the problem becomes decidable or not if it is restricted to complexes that are defined as level set complexes of a $(d - 1)$ -sphere, as is the case for the level set complex $\text{lk}(v)_{=f(v)}$ of $\text{lk}(v)$, is open.

Variants of the Reduction from the Sphere Recognition Problem and the Problem of Recognising PL Morse Functions. Let us resume the discussion concerning reductions from the sphere recognition problem to the regular vertex recognition problem. We proposed in Lemma 6.4 a reduction from the recognition problem for $(d - 1)$ -spheres to the general recognition problem for regular vertices in dimension d . The reduction is designed in such a way that the complex given as instance for the sphere problem, modified by a single elementary starring, plays the role of $\text{lk}(v)$ in the instance of the regular vertex recognition problem. There is a conceptually simpler reduction where the given instance plays the role of $\text{lk}(v)_{=f(v)}$ instead. However, in this reduction all vertices of the given complex are mapped to the same f -value $f(v)$, so that the map is not in general position. Moreover it reduces $(d - 2)$ -sphere recognition to regular vertex recognition in dimension d , so that undecidability can only be

inferred from dimension 7 on and not for dimension 6. The reduction works as follows: Let K be the given instance of the $(d-2)$ -sphere recognition problem. First take the suspension L of K with apices a_+ and a_- , which will play the role of $\text{lk}(v)$ in our reduction. Then take again the suspension M of L with apices v and v' . Define f by the values $f(v) = 0$, $f(a_+) = 1$, $f(a_-) = -1$, and $f(w) = 0$ for all vertices w of K ; the value $f(v')$ can be chosen arbitrary because v' does not belong to $\text{st}(v)$. By construction, we have $\text{lk}(v) = L$ and $\text{lk}(v)_{=f(v)} = K$. If K is a combinatorial $(d-2)$ -sphere, then it is unknotted in its suspension $L = \text{susp}(K)$. Since additionally positive and negative values are attained in $\text{lk}(v)$, namely at the apices a_+ and a_- , the vertex v is a regular vertex in dimension d in that case. On the other hand, if K is not a combinatorial sphere, then L is not a combinatorial sphere either and consequently M is not a combinatorial manifold without boundary. As for the other reduction, we can sort out certain complexes K including combinatorial balls by a preprocessing to avoid instances M that are combinatorial manifolds with boundary.

Another variant of the reduction shows that it is undecidable whether a map f linear on cells of a simplicial complex M of dimension $d \geq 6$ is a PL Morse function. Recall from Lemma 6.2 that once f is known to be a Morse function, regular vertices and the different kinds of non-degenerate critical vertices are algorithmically distinguishable by computing the Betti numbers of the lower level link of the vertex. In this sense, the hardness of recognising regular vertices in higher dimensions lies in differentiating between a regular vertex and a degenerate critical vertex for which the homologies of the complexes $\text{lk}(v)$, $\text{lk}(v)_{=f(v)}$, $\text{lk}(v)_{\leq f(v)}$, and $\text{lk}(v)_{\geq f(v)}$ for the considered vertex v match with those required by a regular vertex.

Lemma 6.7. *The problem of recognising $(d-1)$ -spheres can be reduced to the decision problem whether a map $f: |M| \rightarrow \mathbb{R}$ that is linear on cells of a given d -dimensional complex M and in general position is a PL Morse function or not. In particular, the latter problem is undecidable for $d \geq 6$.*

Proof. Let K be an instance for the $(d-1)$ -sphere recognition problem. In a preprocessing, we check whether K is a simplicial complex in which all maximal cells have dimension $d-1$ and each $(d-2)$ -cell has exactly two proper co-faces. If one of the tests fails, the instance K is reduced to a fixed negative instance (M, f) where f is not a PL Morse function. Otherwise M is the suspension of K with apices v_+ and v_- , and for defining f we choose arbitrary distinct values for the vertices in K and a value $f(v_+)$ strictly greater and a value $f(v_-)$ strictly smaller than all these values. We claim that f is a PL Morse function if and only if K is a combinatorial $(d-1)$ -sphere.

By construction K is the link of both the apices v_+ and v_- . The preprocessing asserts that K is neither a combinatorial ball nor a combinatorial k -sphere for $k \neq d-1$, and if it is not a combinatorial $(d-1)$ -sphere either, then M is not a combinatorial manifold and f is not a PL Morse function.

It remains to consider the case that K is a combinatorial $(d-1)$ -sphere. We verify that all vertices in K are regular vertices of f , and the apices v_- and v_+ are non-degenerate critical vertices of index 0 and d respectively. Since $f(v_-) < f(u)$ holds for all vertices u of the $(d-1)$ -sphere $K = \text{lk}(v_-)$, the vertex v_- is a non-degenerate critical vertex in dimension d of index 0. Symmetrically, v_+ is a non-degenerate critical vertex in dimension d of index d . Let u be a vertex of K . Its link in M is the suspension of its link in K with the apices v_+

and v_- . Since K is a combinatorial $(d-1)$ -sphere, $\text{lk}_K(u)$ is a combinatorial $(d-2)$ -sphere and a suitable subdivision L of it is combinatorially equivalent to a $(d-2)$ -sphere that forms a link complex L' for the origin in the $(d-1)$ -dimensional subspace of \mathbb{R}^d consisting of those points whose first coordinate is zero. Joining this link complex L' with the 0-sphere E consisting of the points $e_1 = (1, 0, \dots, 0)$ and $-e_1 = (-1, 0, \dots, 0)$ yields a link complex for the origin in \mathbb{R}^d that is combinatorially equivalent to the subdivision VL of $\text{lk}_M(u)$, where V denotes the complex consisting of the two vertices v_+ and v_- , with v_+ corresponding to e_1 and v_- corresponding to $-e_1$ under the combinatorial equivalence. Then an ϵ -perturbation of EL' that moves each vertex of L' in positive or negative X_1 -direction depending on the sign of $f - f(u)$ at the corresponding vertex in L witnesses that f at u is locally equivalent to π_1^d at the origin of \mathbb{R}^d . Thus u is a regular vertex in dimension d .

The claimed undecidability for $d \geq 6$ follows immediately from the undecidability of the $(d-1)$ -sphere recognition problem for $d-1 \geq 5$ and the reduction outlined above. \square

Regular Vertex Recognition Problem for Manifolds. Let us continue with a few thoughts and remarks concerning the open question that we called the regular vertex recognition problem for d -manifolds. Our reduction from the sphere recognition problem to the general regular vertex recognition problem produced negative instances where the link of the considered vertex already failed to be a sphere, so that the vertex is not regular no matter which values for f we choose. Therefore we cannot exclude a priori that the problem might become decidable when $\text{lk}(v)$ is known to be a sphere. Clearly, there are some necessary conditions for v being regular that we can test algorithmically: We can compute the homologies of the level set complexes $\text{lk}(v)_{=f(v)}$, $\text{lk}(v)_{\leq f(v)}$, and $\text{lk}(v)_{\geq f(v)}$ and check whether they agree with those required by a regular vertex. If the homologies do not match we have verified that v is not regular. But how should we proceed, if $\text{lk}(v)_{\leq f(v)}$ and $\text{lk}(v)_{\geq f(v)}$ cannot be distinguished by means of homology from $(d-1)$ -balls and $\text{lk}(v)_{=f(v)}$ not from a $(d-2)$ -sphere?

If f is in general position, we can be sure that the three level sets are manifolds and $\text{lk}(v)_{=f(v)}$ is the boundary of the other two level set complexes. In this case we have decomposed the $(d-1)$ -sphere $\text{lk}(v)$ into two homology $(d-1)$ -balls $\text{lk}(v)_{\leq f(v)}$ and $\text{lk}(v)_{\geq f(v)}$ that intersect in a homology $(d-2)$ -sphere $\text{lk}(v)_{=f(v)}$ forming their boundaries. If it were true that the only homology $(d-1)$ -ball that embeds piecewise linearly into the $(d-1)$ -sphere is the standard $(d-1)$ -ball, such a decomposition could only occur if $(\text{lk}(v), \text{lk}(v)_{=f(v)})$ were an unknotted $(d-1, d-2)$ -sphere. This means that checking the homology would actually be sufficient for recognising regular vertices. We already encountered the question whether there are homology $(d-1)$ -balls different from the standard ball embeddable into the $(d-1)$ -sphere when we compared the definition of a critical vertex of Brehm and Kühnel with our definition.

Still assuming general position and additionally $d \neq 5$, the truth of the Schoenflies conjecture for locally flat embeddings suggests a similar question regarding the embedding of $\text{lk}(v)_{=f(v)}$: If the standard $(d-2)$ -sphere is the only homology $(d-2)$ -sphere with a locally flat piecewise linear embedding into the $(d-1)$ -sphere, then checking the homology of $\text{lk}(v)_{=f(v)}$ would suffice for solving the regular vertex recognition problem for d -manifolds with maps

in general position. Since for $d - 2 \leq 2$ the standard $(d - 2)$ -sphere is the only homology $(d - 2)$ -sphere, this observation re-establishes the decidability of the regular vertex recognition problem for d -manifolds for $d \leq 4$. For higher dimensions, the only results I found deal with the case $d = 5$, but only for topological manifolds on the one hand and for smooth manifolds on the other hand. Freedman [20] states that all homology 3-spheres embed topologically into the 4-sphere and notes that the Poincaré homology sphere and more generally homology 3-spheres with Rochlin invariant 1 do not embed smoothly into the 4-sphere. Budney and Burton [10] explore a census of 3-manifolds and ask which of them embed smoothly into the 4-sphere. It turns out that among the homology 3-spheres contained in the census both cases occur: Some of them embed smoothly, including examples different from the standard sphere, and some do not. However, the case $d = 5$ is only of minor interest for our purposes since the standard 3-sphere is algorithmically distinguishable from other homology 3-spheres and this still does not solve the problem because the Schoenflies conjecture for piecewise linear embeddings of the 3-sphere into the 4-sphere is open even for the locally flat case.

Strengthening our hardness result for general regular vertex recognition to the restricted problem for manifolds seems to be a hard task as well. With the characterisations of regular points we have at hand, other commonly known undecidable problems than sphere and ball recognition do not have an obvious connection to the regular vertex recognition problem. The reduction given in Lemma 6.4 can be understood as follows: From the given candidate K for being a sphere, a small ball is removed and the construction ensures that the rest K_- plays the role $\text{lk}(v)_{\geq f(v)}$ in the instance for the regular vertex problem. But in the construction, $\text{lk}(v)$ is only a sphere if K_- is a ball. For a reduction that produces only manifold instances, we would ideally aim for a construction that always produces a $(d - 1)$ -sphere as $\text{lk}(v)$ and embeds K_- as $\text{lk}(v)_{\geq f(v)}$ into that sphere. Clearly, this requires K_- to be embeddable into the $(d - 1)$ -sphere. Matoušek, Tancer, and Wagner [32] showed that the question whether a given $(d - 1)$ -dimensional complex embeds in \mathbb{R}^{d-1} is undecidable for $d - 1 \geq 5$. It seems likely that an analogous result holds regarding embeddability into the $(d - 1)$ -sphere. In this case, sorting out exactly those complexes that do not embed is impossible. Even if we manage to sort out some other subset of complexes that contains all those not embeddable but does not contain any ball, it is unclear how to construct an embedding algorithmically for the remaining complexes.

A Heuristic Approach and Structured Meshes. We conclude this section with two final remarks, that outline possible algorithmic approaches to solve the problem at least for certain instances. The first remark refers to an heuristic approach to the sphere recognition problem, the second remark considers the special properties of the problem, when it is restricted to structured meshes.

Experiments by Joswig, Lutz, and Tsuruga [27] suggest that an algorithm based on heuristics often succeeds in recognising spheres even in higher dimensions where the problem is undecidable in general. Thus we may hope that for most of the vertices that we want to test for regularity, our reductions to sphere recognition problems produce instances for which the heuristic approach performs quite well.

Among the scenarios proposed by Chiang, Lenz, Lu, and Rote [13] as examples for a map $f: |M| \rightarrow \mathbb{R}$ linear on cells of a simplicial complex M , the case of structured meshes deserves special attention when it comes to recognising regular vertices. A structured mesh is a triangulation of a set of grid points that follows a regular pattern. The grid points can be thought of as sample points where the values of a scalar function are measured and the simplex-wise linear interpolation f of these values serves as an approximation on the whole domain of the mesh for the actual function underlying the measurements. Observe that in a structured mesh, up to combinatorial equivalence, only finitely many different simplicial complexes occur as link of a vertex. In fact, the following considerations apply to any class of simplicial complexes where only finitely many combinatorially different vertex links occur. We can infer from Lemma 3.7, that if $\text{lk}_M(x)$ and $\text{lk}_{M'}(x')$ are combinatorially equivalent and $f - f(x)$ and $f' - f'(x')$ have the same sign at corresponding vertices, then the triples (M, x, f) and (M', x', f') are locally equivalent. This means that the local equivalence class of a triple (M, v, f) is uniquely determined by the combinatorial type of $\text{lk}_M(v)$ and for each of the vertices of $\text{lk}_M(v)$ the information whether its f -value is smaller than, greater than, or equal to $f(v)$. Thus, if the combinatorial type of $\text{lk}_M(v)$ is restricted to finitely many types, there are also only finitely many local equivalence classes that a triple (M, v, f) could belong to. This reduces the question whether (M, v, f) is locally equivalent to $(\mathbb{R}^d, 0, \pi_1^d)$, in other words whether v is regular, to a combinatorial problem with only finitely many instances.

For such a finite problem, an algorithm giving correct answers always exists. But there is a subtlety to keep in mind here: At this point, we can not exclude that among the finitely many combinatorially distinct instances, there are instances where the question whether they are locally equivalent to $(\mathbb{R}^d, 0, \pi_1^d)$ is independent of the axiomatic system we use for our proofs. For such a case, we might want to allow decision algorithms with a third possible result besides acceptance and rejection that indicates independence of the instance from the axiomatic system. In this sense, the problem of recognising regular vertices becomes decidable when it is restricted to a class of simplicial complexes with only finitely many combinatorially different vertex links, in particular for structured meshes of fixed kind and dimension.

6.2 Computing Isotopies

In Theorem 4.20 we proved for maps in general position that if $f^{-1}[a, b]$ contains no critical vertices, then an f -level-preserving PL isotopy $\phi: |M|_{=a} \times [a, b] \rightarrow |M|_{\in [a, b]}$ exists. Since the isotopy is a special kind of homeomorphism, it can be represented by suitable combinatorially equivalent simplicial subdivisions of the two complexes $M_{=a} \times [a, b]$ and $M_{\in [a, b]}$. The question arises whether these subdivisions and the combinatorial equivalence between them can be computed by an algorithm. Analysing the proofs leading to the existence result shows that there is only one non-constructive part, for which algorithmic solutions are only available in low dimensions or in form of heuristics that might fail for certain instances. All other steps can be performed in principle by algorithms when a solution for the non-constructive part is provided. In this section, we develop for the low dimensional cases upper bounds on the number of cells of

the combinatorially equivalent subdivisions that a straightforward algorithmic implementation of the theoretical results would produce. For our studies, we assume throughout the section that M is a combinatorial manifold without boundary and that the map $f: |M| \rightarrow \mathbb{R}$ is linear on cells of M and in general position.

6.2.1 General Remarks on the Computation of Level Preserving Isotopies and the Combinatorially Equivalent Subdivisions Representing Them

In the proof of Theorem 4.20, we decompose the problem into subintervals that are easier to deal with because their preimage contains at most one vertex, which is additionally one of the endpoints of the interval. The problem for those subintervals whose preimage contains a vertex is solved in Lemma 4.17 by considering star and deletion of the vertex separately while taking care that the isotopies match on the common intersection. This approach of decomposing the problem into smaller pieces has consequences for representing the isotopies by combinatorially equivalent subdivisions. When we do not have an assertion that the subdivisions coming from the isotopies of two sub-problems match on their common intersection, we have to subdivide both solutions even further to obtain subdivisions that match and provide a solution for the larger problem. For subintervals $[a, b]$ whose preimage do not contain vertices or when the deletion of the single vertex is considered, the complexes $M_{=a} \times [a, b]$ and $M_{\in[a,b]}$ or $(\text{del}_M(v))_{=a} \times [a, b]$ and $(\text{del}_M(v))_{\in[a,b]}$ are already combinatorially equivalent without subdivisions. However, since these complexes are usually polytopal, we have to subdivide the cells to obtain simplicial subdivisions. These tasks can be solved in principle algorithmically by taking intersection complexes and performing suitable stellar subdivisions like starring at every vertex or variants of the derived subdivision. But the number of cells produced by these kinds of subdivisions can easily grow large. For example, the subdivision of the hypercube $[0, 1]^d$ used in the proof of Lemma 3.27 can be obtained by starring at each vertex in a certain order and produces, starting from the $3^d + 1$ cells of the hypercube, a subdivision containing more than $d!$ simplices [16, Section 6.3.2, pp. 314–315].

The non-constructive part of the proof lies in finding an isotopy near a vertex, that is, an isotopy between $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$, or symmetrically between $\text{st}(v)_{=f(v)-\epsilon}$ and $\text{st}(v)_{=f(v)}$. Lemma 4.15 describes the construction when a homeomorphism between $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ with certain properties is given. The existence of such a homeomorphism is deduced in Lemma 4.16 from the existence of a homeomorphism between $\text{st}(v)_{=f(v)}$ and $\text{lk}(v)_{\geq f(v)}$. This homeomorphism in turn is an almost immediate consequence of the defining property of a regular point that yields homeomorphisms from $\text{lk}(v)_{\geq f(v)}$ and $\text{st}(v)_{=f(v)}$ to some standard combinatorial $(\dim(M) - 1)$ -balls. When a ball recognition algorithm is available that also provides for positive instances combinatorially equivalent subdivisions as certificate, the isotopy between $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ can be constructed by an algorithm from the certificate. With the general problem of recognising balls being undecidable in higher dimensions, it seems unlikely that the isotopy can be computed by an algorithm that works for arbitrary instances. Since under the given assumptions

the complex $\text{st}(v)_{=f(v)}$ for a regular vertex v is combinatorially equivalent to the cone on the boundary of $\text{st}(v)_{=f(v)+\epsilon}$, we could try to transform the complex $\text{st}(v)_{=f(v)+\epsilon}$ into the cone on its boundary by some random Pachner moves. If this heuristic attempt succeeds, we can also compute the isotopy between $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ for this instance.

On the other hand, if an algorithm existed that always succeeds in deciding whether $\text{st}(v)_{=f(v)+\epsilon}$ is PL-homeomorphic to the cone on its boundary, it would actually solve the regular vertex problem for manifolds without boundary and a map in general position because of the following facts: The complex $\text{st}(v)_{=f(v)+\epsilon}$ for sufficiently small ϵ and f in general position is a combinatorial manifold. Hence, for the cone on its boundary being PL-homeomorphic, it is necessary that the boundary is a combinatorial sphere because it forms the link of the apex in the cone. That means that $\text{st}(v)_{=f(v)+\epsilon}$ is a combinatorial ball. By general position, $\text{lk}(v)_{\geq f(v)+\epsilon}$ and $\text{lk}(v)_{\geq f(v)}$ are combinatorially equivalent to $\text{st}(v)_{=f(v)+\epsilon}$ according to Lemma 4.14. Thus $\text{lk}(v)_{\geq f(v)}$ is a ball which proves that $\text{lk}(v)_{=f(v)}$ is unknotted in $\text{lk}(v)$, provided that $\text{lk}(v)$ is a sphere, and consequently v is regular. Therefore $\text{st}(v)_{=f(v)+\epsilon}$ being PL-homeomorphic to the cone on its boundary implies v being regular. The reverse implication follows from the proof of Lemma 4.16, where we noted that $\text{st}(v)_{=f(v)+\epsilon}$ is a combinatorial ball for regular v , which implies that it is PL-homeomorphic to the cone on its boundary.

The constructions in the proofs of Lemmata 4.15, 4.16, and 4.17, which provide the isotopy in the interval $[f(v), f(v) + \epsilon]$, are mostly described in terms of PL homeomorphisms. For the purpose of this section, it is useful to consider how these constructions can be described in terms of combinatorially equivalent subdivisions. In the following, we sketch such a description.

The PL isotopy on $\text{del}(v)$ is a consequence of the combinatorial equivalence of $\text{del}(v)_{\in [f(v), f(v)+\epsilon]}$ and $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$. For obtaining a fixed PL isotopy, we only have to triangulate both complexes by starring at each vertex in corresponding order, so that the triangulations are still combinatorially equivalent. This induces combinatorially equivalent subdivisions of $\text{lk}(v)_{\in [f(v), f(v)+\epsilon]}$ and $\text{lk}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$, which represent a PL isotopy ψ on the link, and combinatorially equivalent subdivisions of $\text{lk}(v)_{=f(v)+\epsilon}$ and $\text{lk}(v)_{=f(v)} \times \{f(v) + \epsilon\}$, which represents a PL homeomorphism $\psi_{f(v)+\epsilon}$ between $\text{lk}(v)_{=f(v)}$ and $\text{lk}(v)_{=f(v)+\epsilon}$ via the identification of $\text{lk}(v)_{=f(v)} \times \{f(v) + \epsilon\}$ with $\text{lk}(v)_{=f(v)}$ by projection.

The next thing we need for the construction are combinatorially equivalent simplicial subdivisions of $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ such that the induced PL homeomorphism between $\text{lk}(v)_{=f(v)}$ and $\text{lk}(v)_{=f(v)+\epsilon}$ agrees with the homeomorphism $\psi_{f(v)+\epsilon}$ obtained above from the isotopy on $\text{del}(v)$. Ideally, the restriction of the two pairs of combinatorially equivalent subdivisions induce the same pair of combinatorially equivalent subdivisions for $\text{lk}(v)_{=f(v)}$ and $\text{lk}(v)_{=f(v)+\epsilon}$, because otherwise a further refinement of all the subdivisions is necessary in order to proceed with the construction. Fortunately, the algorithms for finding combinatorially equivalent simplicial subdivisions of $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ in the low dimensional cases considered here produce the same subdivisions of $\text{lk}(v)_{=f(v)}$ and $\text{lk}(v)_{=f(v)+\epsilon}$ as the triangulations of $\text{del}(v)_{\in [f(v), f(v)+\epsilon]}$ and $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$, either for trivial reasons or thanks to appropriate fine tuning of the algorithm. In this convenient situation the union of the obtained subdivisions for $\text{st}(v)_{=f(v)+\epsilon}$ and for $\text{lk}(v)_{\in [f(v), f(v)+\epsilon]}$ gives a subdivision

Θ_2 of $\text{lk}(v)_{\in [f(v), f(v)+\epsilon]} \cup \text{st}(v)_{=f(v)+\epsilon}$. Analogously, lifting the obtained subdivision for $\text{st}(v)_{=f(v)}$ in the obvious way to a subdivision of $\text{st}(v)_{=f(v)} \times \{f(v) + \epsilon\}$ and taking the union with the obtained subdivision $\text{lk}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ yields a subdivision Θ_1 of $\text{lk}(v)_{=f(v)} \times [f(v), f(v) + \epsilon] \cup \text{st}(v)_{=f(v)} \times \{f(v) + \epsilon\}$ combinatorially equivalent to Θ_2 .

The cone T_1 with apex $(v, f(v))$ and base Θ_1 is a subdivision of $\text{st}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$. It is combinatorially equivalent to the cone T_2 with apex v and base Θ_2 , which forms a subdivision of $\text{st}(v)_{\in [f(v), f(v)+\epsilon]}$. These combinatorially equivalent subdivisions T_1 and T_2 of $\text{st}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ and $\text{st}(v)_{\in [f(v), f(v)+\epsilon]}$ match on $\text{lk}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ and $\text{lk}(v)_{\in [f(v), f(v)+\epsilon]}$ with the restriction of the combinatorially equivalent triangulations of $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ and $\text{del}(v)_{\in [f(v), f(v)+\epsilon]}$ to the link. Therefore the union of T_1 and the triangulation of $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ on the one hand and the union of T_2 and the triangulation of $\text{del}(v)_{\in [f(v), f(v)+\epsilon]}$ on the other hand form the desired combinatorially equivalent subdivisions of $M_{\in [f(v), f(v)+\epsilon]}$ and $M_{=f(v)} \times [f(v), f(v) + \epsilon]$ representing an f -level-preserving PL isotopy.

6.2.2 Dimension 1

For giving upper bounds on the number of cells needed for combinatorially equivalent subdivisions of $M_{=a} \times [a, b]$ and $M_{\in [a, b]}$, we start with the 1-dimensional case. One easily checks that if $f^{-1}[a, b]$ contains no critical vertex, then $M_{=a}$ and $M_{=b}$ both consist of the same number V of vertices. In $M_{\in [a, b]}$, these vertices are paired by V disjoint paths with each path connecting one vertex in $M_{=a}$ to one vertex in $M_{=b}$. The intermediate vertices of the paths are the vertices lying in $f^{-1}(a, b)$. The complex $M_{=a} \times [a, b]$ consists of V disjoint edges, one for each vertex of $M_{=a}$. For constructing the combinatorial equivalence, it suffices to subdivide the edges in $M_{=a} \times [a, b]$ into paths by inserting vertices at the appropriate places. A subdivision of $M_{\in [a, b]}$ is unnecessary. So the number of cells needed for combinatorially equivalent subdivisions can be bounded by the number of cells in $M_{\in [a, b]}$.

For higher dimensional manifolds M , the interval is subdivided and the isotopy is obtained by combining the isotopies on the subintervals. Let us assume that the preimage of the interval $[a, b]$ contains k vertices and that the sequence $(h_i)_{i=1}^k$ denotes the f -values of these vertices in sorted order, that is $a \leq h_1 < \dots < h_k \leq b$. The proof of Theorem 4.20 uses subintervals that contain at most one value h_i , with the value being either the left or the right endpoint of the interval. Such a subdivision can be obtained by decomposing the interval at the values h_i and the intermediate values $(h_i + h_{i+1})/2$. We denote the interval complex formed by this subdivision of $[a, b]$ by \mathcal{I} .

6.2.3 Dimension 2

Isotopies across a Single Vertex. For a 2-dimensional manifold M , the 2-dimensional cells of $M_{\in \mathcal{I}}$ with \mathcal{I} defined as above are either triangles or quadrilaterals. For a regular vertex v , the isotopy for an interval $[f(v), f(v) + \epsilon]$ in the star of v is constructed from combinatorially equivalent subdivisions of $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$. The level set complex $\text{st}(v)_{=f(v)}$ is the cone on a 0-sphere consisting of two edges sharing the vertex v as apex of the cone, in other words a path formed by two edges. The level set complex $\text{st}(v)_{=f(v)+\epsilon}$ is

a path as well that consists of at least two edges. Combinatorially equivalent subdivisions for two paths can be found by inserting vertices on the shorter path until it has the same number of edges as the longer path. Note that we used two cells more than necessary for constructing such combinatorially equivalent subdivisions in the example depicted in Figure 4.1 on page 162, where we added one vertex to the longer path $\text{st}(v)_{=f(v)+\epsilon}$ as a counterpart to the vertex v of $\text{st}(v)_{=f(v)}$. Clearly we could have chosen any of the interior vertices of the path $\text{st}(v)_{=f(v)+\epsilon}$ as counterpart for v instead of a newly inserted vertex. The resulting upper bound on the number of cells of combinatorially equivalent subdivisions for $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ is the number of cells in $\text{st}(v)_{=f(v)+\epsilon}$.

The cone construction described in Lemma 4.15 extends this combinatorial equivalence between $\text{st}(v)_{=f(v)+\epsilon}$ and a subdivision of $\text{st}(v)_{=f(v)}$ to combinatorially equivalent subdivisions of $\text{st}(v)_{\in[f(v),f(v)+\epsilon]}$ and $\text{st}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$. From the illustration in Figure 4.1 on page 162 we see immediately that applying this cone construction to $\text{st}(v)_{\in[f(v),f(v)+\epsilon]}$ leaves all cells unchanged except for the two quadrilaterals present in the complex, which are each subdivided into two triangles and a separating diagonal edge. In particular, the two edges from $\text{lk}(v)_{\in[f(v),f(v)+\epsilon]}$ are unaffected and the subdivision matches with any subdivision of $\text{del}(v)_{\in[f(v),f(v)+\epsilon]}$ that only subdivides the quadrilaterals into two triangles by a diagonal. The polytopal complexes $\text{del}(v)_{\in[f(v),f(v)+\epsilon]}$ and $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ are already combinatorially equivalent without subdivision. It suffices to subdivide the quadrilaterals in both complexes in a consistent manner by diagonals to obtain combinatorially equivalent simplicial subdivisions. Choosing arbitrary subdivisions of that kind and combining them with the combinatorially equivalent subdivisions of $\text{st}(v)_{\in[f(v),f(v)+\epsilon]}$ and $\text{st}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ yields combinatorially equivalent subdivisions of $M_{\in[f(v),f(v)+\epsilon]}$ and $M_{=f(v)} \times [f(v), f(v) + \epsilon]$. The subdivision of $M_{\in[f(v),f(v)+\epsilon]}$ can be obtained from the original complex by only subdividing quadrilaterals by diagonals. Each cell S of M induces at most three cells in the complex $M_{\in[f(v),f(v)+\epsilon]}$, namely $S_{=f(v)}$, $S_{\in[f(v),f(v)+\epsilon]}$, and $S_{=f(v)+\epsilon}$. Since only cells of the form $S_{\in[f(v),f(v)+\epsilon]}$ can be quadrilaterals, the combinatorially equivalent subdivisions of $M_{\in[f(v),f(v)+\epsilon]}$ and $M_{=f(v)} \times [f(v), f(v) + \epsilon]$ have at most five times as many cells as M .

The observations made above for an interval $[f(v), f(v) + \epsilon]$ with $f(v)$ as left endpoint apply in a symmetric fashion to intervals $[f(v) - \epsilon', f(v)]$ with $f(v)$ as right endpoint. The complexes for the lower and upper interval meet in the common subcomplexes $M_{=f(v)}$ and $M_{=f(v)} \times \{f(v)\}$. The subdivisions constructed above leave these subcomplexes $M_{=f(v)}$ and $M_{=f(v)} \times \{f(v)\}$ unchanged, so that the union of the subdivisions $M_{\in[f(v)-\epsilon',f(v)]}$ and $M_{\in[f(v),f(v)+\epsilon]}$ on the one hand and the union of the subdivisions of $M_{=f(v)} \times [f(v) - \epsilon', f(v)]$ and $M_{=f(v)} \times [f(v), f(v) + \epsilon]$ on the other hand form combinatorially equivalent subdivisions of $M_{\in[f(v)-\epsilon',f(v)+\epsilon]}$ and $M_{=f(v)} \times [f(v) - \epsilon', f(v) + \epsilon]$. The number of cells of these subdivisions is at most $9n_M$, where n_M is the number of cells in M .

Isotopies across Several Vertices by Recursive Merging. Applying the above considerations to each vertex v whose f -value is one of the levels h_i yields combinatorially equivalent subdivisions of $M_{\in[a,(h_1+h_2)/2]}$ and $M_{=h_1} \times$

$[a, (h_1 + h_2)/2]$, of $M_{\in[(h_{k-1}+h_k)/2, b]}$ and $M_{=h_k} \times [(h_{k-1} + h_k)/2, b]$, and of $M_{\in[(h_{i-1}+h_i)/2, (h_i+h_{i+1})/2]}$ and $M_{=h_i} \times [(h_{i-1} + h_i)/2, (h_i + h_{i+1})/2]$ for each index i with $1 < i < k$, with each subdivision having at most $9n_M$ cells. The remaining task is to create subdivisions for the whole interval $[a, b]$ from the subdivisions of the subintervals. This can be achieved by combining subdivisions of pairs of neighbouring intervals in a divide and conquer fashion: Solve the problem for the first $\lceil k/2 \rceil$ subintervals and the last $\lfloor k/2 \rfloor$ subintervals recursively, and combine afterwards the solutions found for the interval covered by the first half of subintervals and for the interval covered by the second half of subintervals to a solution for the whole interval. For estimating the total number of cells needed, we first analyse a single recursive step and derive an upper bound on the number of cells produced by a merging operation of two neighbouring intervals with given combinatorially equivalent subdivisions of a certain size.

A Single Merging Step. Let us introduce some notation for the analysis of a general merging step. We consider a left hand interval $[l, m]$ and a right hand interval $[m, r]$. For the left hand interval, we have a subdivision L of $M_{\in[l, m]}$ combinatorially equivalent to a subdivision L^\times of $M_{=h^-} \times [l, m]$ for some $h^- \in [l, m]$. Analogously for the right hand interval, there is a subdivision R of $M_{\in[m, r]}$ combinatorially equivalent to a subdivision R^\times of $M_{=h^+} \times [m, r]$ for some $h^+ \in [m, r]$. Both combinatorial equivalences represent an f -level-preserving PL isotopy. They induce a such an isotopy on the interval $[l, r]$ according to Lemma 4.6. Our aim is to find an upper bound on the number of cells of combinatorially equivalent subdivisions that represent this induced isotopy.

In the proof of Lemma 4.6, we assume that both isotopies are the identity at the level at which they intersect, so that they can easily combined to an isotopy on the whole interval. Combinatorially equivalent subdivisions that represent an isotopy that is the identity at level m can be constructed from the given combinatorially equivalent subdivisions as explained in the following for the example of L and L^\times : The complexes L and L^\times contain as subcomplexes at level m combinatorially equivalent subdivisions L_m of $M_{=m}$ and L_m^\times of $M_{=h^-} \times \{m\}$, which can be identified with $M_{=h^-}$. Note that any simplicial subdivision of L_m induces a combinatorially equivalent subdivision of L_m^\times and vice versa via the PL homeomorphism between L_m and L_m^\times that is provided by their combinatorial equivalence. So let K be a simplicial subdivision of L_m and K^\times its combinatorially equivalent counterpart forming a subdivision of L_m^\times . Using the projection that identifies $M_{=h^-} \times \{m\}$ with $M_{=h^-}$, the subdivision K^\times of L_m^\times can be considered as a subdivision of $M_{=h^-}$ as well. In this sense, $K^\times \times [l, m]$ can be regarded as a subdivision of $M_{=h^-} \times [l, m]$. This subdivision is combinatorially equivalent to the subdivision $K \times [l, m]$ of $M_{=m} \times [l, m]$. Triangulating these subdivisions by inserting corresponding diagonals in the quadrilaterals of $K^\times \times [l, m]$ and $K \times [l, m]$ yields combinatorially equivalent simplicial subdivisions of $M_{=h^-} \times [l, m]$ and $M_{=m} \times [l, m]$.

Now we have two subdivisions of $M_{=h^-} \times [l, m]$, namely L^\times and the triangulation of $K^\times \times [l, m]$. Their intersection complex forms a common subdivision of L^\times and $K^\times \times [l, m]$, which we can triangulate by starring at each vertex. For this triangulated common subdivision N^\times , a combinatorially equivalent subdivi-

vision N of L and thus of $M_{\in[l,m]}$ can be found using the homeomorphism and isotopy represented by the combinatorial equivalence between L and L^\times . Similarly, we obtain a combinatorially equivalent subdivision N' of $M_{=m} \times [l, m]$ using the homeomorphism induced by the combinatorial equivalence between the triangulations of $K^\times \times [l, m]$ and $K \times [l, m]$. Since K and its combinatorially equivalent counterparts are already simplicial, the subdivisions leave the three subcomplexes at level m unaffected, namely the subcomplexes $N_m^\times = K^\times \times \{m\}$, $N_m = K$, and $N'_m = K \times \{m\}$. The combinatorial equivalence between N_m^\times and N_m , and the combinatorial equivalence between N_m^\times and N'_m are both given by the homeomorphism between $M_{=h^-}$ and $M_{=m}$ that results from the isotopy. Therefore the combinatorial equivalence between N'_m and N_m represents the projection that identifies $M_{=m} \times \{m\}$ with $M_{=m}$. In other words, N' and N are combinatorially equivalent simplicial subdivisions of $M_{=m} \times [l, m]$ and $M_{\in[l,m]}$ representing a PL isotopy that is the identity at level m as desired.

For finding an upper bound on the number of cells of N , N' , and N^\times , let n_L and n_K be the number of cells in L and K respectively. The complex $K \times [l, m]$ and its combinatorially equivalent counterpart $K^\times \times [l, m]$ have $3n_K - 2$ cells. Among them are $(n_K - 1)/2$ quadrilaterals because we assume that K triangulates a 1-dimensional manifold without boundary. Triangulating $K^\times \times [l, m]$ by inserting diagonals produces $4n_K - 3$ cells in total. The intersection complex of this triangulation with L^\times has at most $4n_K n_L$ cells. Triangulating this intersection complex yields N^\times which has the same number of cells as N and N' . For bounding the number of cells in N^\times , we observe that the polyhedra in the intersection complex are of bounded complexity, since they are obtained by intersecting simplices. Cells of dimension 0 and 1 do not require further subdivision, we only need to subdivide the 2-dimensional polytopes that arise from intersecting two triangles. Polytopes arising in that way are triangles, quadrilateral, pentagons, and hexagons. They can be triangulated by inserting suitable diagonals. Hence any polytope present in the intersection complex is replaced by at most 4 new triangles and 3 new diagonals in the triangulation. Leaving aside that actually only a bounded fraction of cells requires further subdividing we obtain $28n_K n_L$ as generous upper bound on the number of cells in N^\times .

For the recursion, it is convenient to have an upper bound on the number of those cells of N , N' , and N^\times that belong to level l , because if $[l, r]$ is used as a right hand interval in the next recursion step and treated symmetrically to the description for a left hand interval above, level l plays the role of m , the complex N plays the role of R , and some subdivision of N_l plays the role of K . Following the construction of N^\times with a focus on level l , we can observe that we start with $K^\times \times \{l\}$, which is simplicial and hence not affected by the triangulation of $K^\times \times [l, m]$. Forming the intersection complex of $K^\times \times [l, m]$ with L^\times yields at level l the intersection complex of $K^\times \times \{l\}$ with the subcomplex L_l^\times of L^\times that subdivides $M_{=h^-} \times \{l\}$. If n_l denotes the number of cells in L_l^\times , the usual bound on the number of cells in an intersection complex given by the product yields at most $n_l n_K$ cells in the intersection complex at level l . Since we are in the special case that we deal with two subdivisions of a 1-dimensional manifold without boundary, namely of $M_{=h^-} \times \{l\}$, the number of cells in the intersection complex is even bounded by the sum $n_l + n_K$ of the numbers of cells of the original complexes and this number does not increase by triangulating the intersection complex.

Summarising the construction so far, we have obtained from L and L^\times together with a chosen simplicial subdivision K_L of L_m combinatorially equivalent subdivisions N_L of $M_{\in[l,m]}$ and N'_L of $M_{=m} \times [l,m]$ representing an f -level-preserving PL isotopy that is the identity at level m . Symmetrically, we can construct from R , R^\times , and a chosen simplicial subdivision K_R of R_m combinatorially equivalent subdivisions N_R of $M_{\in[m,r]}$ and N'_R of $M_{=m} \times [m,r]$ representing an f -level-preserving PL isotopy that is the identity at level m as well. Moreover, in both subdivisions N_L and N_R , the subcomplex at level m is the chosen complex K_L or K_R respectively. Now when we choose for $K = K_L = K_R$ a simplicial subdivision of $M_{=m}$ that is a common subdivision of L_m and R_m , then the resulting subdivisions N_L and N_R match at level m , and so do the subdivisions N'_L and N'_R . Therefore the union of N_L and N_R defines a subdivision of $M_{\in[l,r]}$, just as the union of N'_L and N'_R defines a subdivision of $M_{=m} \times [l,r]$. These subdivisions are combinatorially equivalent and represent an f -level-preserving PL isotopy $\phi: |M|_{=m} \times [l,r] \rightarrow |M|_{\in[l,r]}$ that is the identity at level m . The number of cells in these subdivisions is at most $28n_K(n_L + n_R)$ and their subcomplexes at level l and r have at most $n_l + n_K$ or $n_r + n_K$ cells respectively, where the numbers n_L , n_R , n_K , n_l , and n_r denote the numbers of cells of the complexes L , R , K , $L_{=l}$, and $R_{=r}$ in that order.

Setting Up and Solving the Recurrence Relation. Now let us track how the size of the subdivisions grows during the recursive steps. We use an index j for differentiating the levels of recursion, where $j = 0$ denotes the deepest level, where we have k pairs of complexes, each covering one of the intervals $[a, (h_1 + h_2)/2]$, $[(h_{k-1} + h_k)/2, b]$, and $[(h_{i-1} + h_i)/2, (h_i + h_{i+1})/2]$ for an index i with $1 < i < k$. With growing level j , the constructed subdivisions cover up to 2^j subsequent subintervals while the number of pairs roughly halves in each step, until we obtain finally at level $j = \lceil \log k \rceil$ a single pair of subdivisions covering the whole interval $[a, b]$. The term $n_I(j)$ is used in the following to denote an upper bound on the number of cells in a subdivision that occurs at level j in a pair that represents an isotopy covering up to 2^j subsequent subintervals. In particular, we have $n_I(0) \leq 9n_M$, and $n_I(\lceil \log k \rceil)$ is the desired upper bound on the number of cells of combinatorially equivalent subdivisions representing an f -level-preserving PL isotopy from level a to level b . For a merging step at recursion level j , the value $n_I(j)$ can be used as upper bound for n_L and n_R , and an upper bound on the size of the complex N produced by this merging step yields an upper bound for $n_I(j + 1)$.

In the analysis of a merging step, we also needed the size of those subcomplexes of the subdivisions L and R that are defined by the f -levels l , m , and r , which are among the levels $(h_{i-1} + h_i)/2$. For the maximum size of such a subcomplex at some f -level $(h_{i-1} + h_i)/2$ that separates two subintervals which are not merged yet at recursion level j , we use the upper bound $n_h(j)$. This bound can be used as an upper estimate on the size n_l of $L_{=l}$, the size n_r of $R_{=r}$ and of the sizes of $L_{=m}$ and $R_{=m}$ in each merging step occurring at recursion level j . The maximum size of a complex K used in a merging step at recursion level j is the maximum size of a common subdivision of complexes $L_{=m}$ and $R_{=m}$ occurring in such a merging step. Since $L_{=m}$ and $R_{=m}$ are 1-dimensional and their sizes are bounded by $n_h(j)$, the maximum size $n_K(j)$ of a complex K used in a merging step at recursion level j is at most $2n_h(j)$. For the case $j = 0$, we

even have $n_K(0) \leq n_h(0) \leq n_M$ because we start the recursion with complexes L and R that fulfil $L_{=m} = R_{=m} = M_{=m}$, so that we can choose $K = M_{=m}$.

Our analysis of a merging step allows to compute the values $n_I(j)$, $n_h(j)$, and $n_K(j)$ recursively. The fact that the constructed subdivision N has at most $28n_K(n_L + n_R)$ cells yields the recursive bound $n_I(j+1) \leq 56n_K(j)n_I(j)$. From the bounds $n_l + n_K$ and $n_r + n_K$ for the subcomplexes $N_{=l}$ and $N_{=r}$, we obtain the inequality $n_h(j+1) \leq n_h(j) + n_K(j)$. For $n_K(j)$, we have already mentioned the relation $n_K(j) \leq 2n_h(j)$. Resolving first the recursion for the $n_h(j)$ and the $n_K(j)$ with initial bounds $n_K(0) \leq n_h(0) \leq n_M$ yields the bounds $n_h(j) \leq 2 \cdot 3^{j-1}n_M$ and $n_K(j) \leq 4 \cdot 3^{j-1}n_M$ for $j \geq 1$. Replacing $n_K(j)$ by this upper bound in the recursion for $n_I(j)$ produces the inequality $n_I(j+1) \leq 224 \cdot 3^{j-1}n_M n_I(j)$. With the initial bound $n_I(0) \leq 9n_M$, we obtain the bounds $n_I(j) \leq 224^j 3^{\frac{1}{2}j(j-1)} 9n_M^{j+1}$ for $j \geq 1$. Thus, the number of cells in combinatorially equivalent subdivisions representing an f -level-preserving PL isotopy in an interval $[a, b]$ without critical values for a combinatorially 2-manifold M without boundary and a map f in general position can be bounded by $224^{\lceil \log k \rceil} 3^{\frac{1}{2} \lceil \log k \rceil (\lceil \log k \rceil - 1)} 9n_M^{\lceil \log k \rceil + 1}$, where k is the number of vertices that have an f -value in $[a, b]$ and n_M the number of cells in M . This bound holds for an PL isotopy from $M_{=h} \times [a, b]$ to $M_{\in[a,b]}$ for a median level h among the levels $(h_{i-1} + h_i)/2$. When we insist on representing an isotopy from $M_{=a} \times [a, b]$ to $M_{\in[a,b]}$, we have to treat the obtained subdivisions as R^\times and R and construct N_R and N'_R just like we did for the other merging steps. Then N_R and N'_R represent an isotopy from $M_{=a} \times [a, b]$ to $M_{\in[a,b]}$, and their size is bounded by $28n_K(\lceil \log k \rceil)n_I(\lceil \log k \rceil) \leq 224^{\lceil \log k \rceil} 3^{\frac{1}{2}(\lceil \log k \rceil + 2)(\lceil \log k \rceil - 1)} 1008n_M^{\lceil \log k \rceil + 2}$.

6.2.4 Merging Isotopies

The analysis of the 2-dimensional case is finished with the result above. Before we continue with the 3-dimensional case, we generalise the construction for merging the subdivisions representing isotopies on subintervals to subdivisions representing an isotopy on the whole interval to higher dimensions. The principal ideas that we outlined for the 2-dimensional case remain the same. Let us repeat the steps for merging the subdivisions for two subsequent subintervals in a short summary while indicating the necessary changes for higher dimensions.

The merging step is performed on two pairs of combinatorially equivalent subdivisions, the pair L and L^\times for the left hand interval and the pair R and R^\times for the right hand interval. The first step is constructing a common simplicial subdivision K of the subcomplexes L_m and R_m at level m , where L and R intersect. The complex K is obtained by first taking the intersection complex of L_m and R_m and then triangulating it by starring at each vertex. Note that triangulating was not necessary in the case of a 2-dimensional complex M because L_m and R_m were 1-dimensional in that case. Now L and L^\times on the one hand and R and R^\times on the other hand are treated symmetrically in the way explained for the left hand interval as follows: The subdivision K of L_m is mapped bijectively to a subdivision K^\times of L_m^\times . Then $K^\times \times [l, m]$ is triangulated, then the intersection complex of this triangulation with L^\times is constructed, and this intersection complex is finally triangulated to form a complex N_L^\times . The resulting complex is mapped by homeomorphisms to combinatorially equivalent subdivisions N_L of $M_{\in[l,m]}$ and N'_L of $M_{=m} \times [l, m]$. For symmetrically obtained complexes N_R

of $M_{\in[m,r]}$ and N'_R of $M_{=m} \times [m,r]$, the unions $N_L \cup N_R$ and $N'_L \cup N'_R$ form the desired combinatorially equivalent subdivisions for the interval $[l,r]$.

For bounding the size of the constructed subdivisions, we have to deal in particular with the effect that triangulating by starring at each vertex has on the number of cells. The first such triangulation is applied to the intersection complex of two $(d-1)$ -dimensional simplicial complexes, the second one to the Cartesian product of a $(d-1)$ -dimensional simplicial complex with an interval, and the last one to an intersection complex of two d -dimensional simplicial complexes. For fixed dimension d , the cells in an intersection complex of two d -dimensional simplicial complexes and the cells in the product complex of a $(d-1)$ -dimensional simplicial complex with an interval are of bounded complexity. Therefore triangulating such a cell by starring at each vertex replaces each cell by at most a constant number of new cells. It suffices to count only those new cells that intersect the interior of the cell because the other cells subdivide the boundary and each one can be attributed to the smallest boundary face containing it.

So let c_d be the maximum number of cells produced by triangulating a face of an intersection of two d -simplices by starring at each vertex, and let c_d^\times be the maximum number of cells produced by triangulating a face of a Cartesian product of a $(d-1)$ -simplex and an interval by starring at each vertex. Then a triangulation of an intersection complex of two d -dimensional simplicial complexes with n_1 and n_2 cells constructed by triangulating at each vertex has at most $c_d n_1 n_2$ cells. Similarly, triangulating a Cartesian product of an interval with a $(d-1)$ -dimensional simplicial complex having at most n cells by starring at each vertex produces at most $(c_d^\times + 2)n$ cells, because those cells that belong to a subcomplex formed by the Cartesian product with an endpoint of the interval are already simplicial and will not be subdivided further.

With these preliminaries, we analyse again a single merging step, now for arbitrary dimension d of M . We reuse the notation from the 2-dimensional case and bound the number of cells of the produced complex N for given sizes n_L of L and n_K of K .

The triangulation of $K^\times \times [l,m]$ constructed first has at most $(c_d^\times + 2)n_K$ cells. Forming the intersection complex of this triangulation with L^\times and triangulating again yields at most $c_d(c_d^\times + 2)n_K n_L$ cells for the complex N^\times and its combinatorially equivalent counterpart N . Focusing on level l , the construction forms an intersection complex of $K^\times \times \{l\}$ with L_l^\times and triangulates it. Therefore the number of cells of N at level l is bounded by $c_{d-1} n_K n_l$.

Now we can derive recursive formulas for $n_I(j)$, $n_h(j)$, and $n_K(j)$ just as in the 2-dimensional case. The starting values $n_I(0)$ and $n_h(0)$ are treated as given. For low dimension, as we did already for $d=2$, they can be determined by analysing the number of cells needed to represent an isotopy for an interval $[f(v) - \epsilon', f(v) + \epsilon]$ containing only a single vertex in its preimage.

The recursion for $n_I(j)$ follows from a bound on the size of the union $N_L \cup N_R$. An obvious upper bound is provided by taking twice an upper bound for the complex N , yielding $n_I(j+1) \leq 2c_d(c_d^\times + 2)n_K(j)n_I(j)$. The formula for $n_h(j)$ is $n_h(j+1) \leq c_{d-1}n_K(j)n_h(j)$. For bounding the size of a complex K at recursion level j by some value $n_K(j)$, note that we construct it by first intersecting the $(d-1)$ -dimensional simplicial complexes $L_{=m}$ and $R_{=m}$, which have size at most $n_h(j)$ and then triangulating the result. Thus the number $n_K(j)$ of cells in K is bounded by $c_{d-1}n_h(j)^2$. Combining this with the recursion formula for

$n_h(j+1)$ gives $n_h(j+1) \leq c_{d-1}^2 n_h(j)^3$, which resolves to $n_h(j) \leq c_{d-1}^{3^j-1} n_h(0)^{3^j}$ and $n_K(j) \leq c_{d-1}^{3^j} n_h(0)^{2 \cdot 3^j}$. Plugging in the term for $n_K(j)$ into the recursion formula for $n_I(j+1)$ yields $n_I(j+1) \leq 2c_d(c_d^\times + 2)c_{d-1}^{3^j} n_h(0)^{2 \cdot 3^j} n_I(j)$. This recursion in turn resolves to $n_I(j) \leq 2^j c_d^j (c_d^\times + 2)^j c_{d-1}^{\frac{3}{2}(3^j-1)} n_h(0)^{3(3^j-1)} n_I(0)$.

Choosing recursion level $j = \lceil \log k \rceil$, where k is the number of subintervals, yields the desired upper bound on the number of cells of combinatorially equivalent subdivisions representing an isotopy on the whole interval. When we insist on representing an isotopy mapping $M_{=a} \times [a, b]$ to $M_{\in[a, b]}$, the additional subdivisions lead to the term $c_d(c_d^\times + 2)n_K(\lceil \log k \rceil)n_I(\lceil \log k \rceil)$ as upper bound for the number of cells. Replacing $n_K(\lceil \log k \rceil)$ and $n_I(\lceil \log k \rceil)$ by the bounds derived above yields as upper bound for the number of cells the term $2^{\lceil \log k \rceil+1} c_d^{\lceil \log k \rceil+1} (c_d^\times + 2)^{\lceil \log k \rceil+1} c_{d-1}^{\frac{3}{2}(5 \cdot 3^{\lceil \log k \rceil-1}-1)} n_h(0)^{5 \cdot 3^{\lceil \log k \rceil}-3} n_I(0)$.

6.2.5 Dimension 3

Now let us discuss the 3-dimensional case. The crucial part is finding combinatorially equivalent subdivisions of $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ for a regular vertex v and sufficiently small ϵ . We analyse two approaches for solving this task. The first approach is based on an idea suggested in the appendix of the paper on constructing contour trees by Chiang, Lenz, Lu, and Rote [13]. They represent a triangulation of $\text{st}(v)_{=f(v)+\epsilon}$ by the vertices, edges, and faces of a planar straight line drawing of its 1-skeleton such that the outer face is a regular polygon and bounded by those edges and vertices that represent the boundary $\text{lk}(v)_{=f(v)+\epsilon}$ of $\text{st}(v)_{=f(v)+\epsilon}$. The complex $\text{st}(v)_{=f(v)}$ can be represented analogously by a wheel graph with the same regular polygon as outer face. This yields two simplicial complexes with the same domain, namely the convex region defined by the regular polygon. Hence combinatorially equivalent subdivisions can be found by triangulating the intersection complex of these two triangulations of the regular polygon. The other approach is based on Pachner moves that transform $\text{st}(v)_{=f(v)+\epsilon}$ into the cone on its boundary.

Approach by Planar Straight Line Drawings. For the first approach we proceed as outlined above. First we triangulate $\text{st}(v)_{=f(v)+\epsilon}$, which requires only subdividing quadrilaterals. Thus the number of cells in the triangulation is at most three times the number of cells in $\text{st}(v)_{=f(v)+\epsilon}$ and the boundary is unaffected by the triangulation. Then we represent this triangulation by a combinatorially equivalent complex formed by the vertices, edges, and interior faces of a planar straight-line drawing of its 1-skeleton. The drawing is chosen in such a way that its outer face is bounded by the edges and vertices representing the boundary $\text{lk}(v)_{=f(v)+\epsilon}$ and forms a regular polygon. Taking the cone on this bounding polygon with an apex in its interior, for example the centre point, yields a complex with the same domain that is combinatorially equivalent to $\text{st}(v)_{=f(v)}$. The combinatorial equivalence comes from the fact that $\text{lk}(v)_{=f(v)}$ and $\text{lk}(v)_{=f(v)+\epsilon}$ are combinatorially equivalent and $\text{st}(v)_{=f(v)}$ is a cone on $\text{lk}(v)_{=f(v)}$. Now we have two triangulations of the convex region bounded by the regular polygon. Triangulating their intersection complex yields a common subdivision of these triangulations. The common subdivision induces corresponding subdivisions on $\text{st}(v)_{=f(v)}$ and on $\text{st}(v)_{=f(v)+\epsilon}$, which form the

desired pair of combinatorially equivalent subdivisions. A bound on the number of cells of these subdivisions is given by first multiplying the sizes of the triangulation of $\text{st}(v)_{=f(v)+\epsilon}$ and the size of $\text{st}(v)_{=f(v)}$, which gives an upper bound on the size of the intersection complex, and multiplying the result by the factor $c_2 = 7$ for taking the final triangulation into account. Since the cells in the 1-dimensional subcomplexes $\text{lk}(v)_{=f(v)+\epsilon}$ and $\text{lk}(v)$ are not affected by the subdivisions, the combinatorial equivalence between $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ restricts to the natural one between $\text{lk}(v)_{=f(v)}$ and $\text{lk}(v)_{=f(v)+\epsilon}$.

Approach by Pachner Moves. The other approach is based on a fact that we cite from Mijatović [34, Lemma 4.2]: Any combinatorial 2-ball with t triangles can be transformed into the cone on its boundary by at most t Pachner moves. Since a triangulation of $\text{st}(v)_{=f(v)+\epsilon}$ is a combinatorial 2-ball for a regular vertex of a combinatorial 3-manifold M and $\text{st}(v)_{=f(v)}$ is combinatorially equivalent to the cone on the boundary of $\text{st}(v)_{=f(v)+\epsilon}$, the cited fact yields for any triangulation of $\text{st}(v)_{=f(v)+\epsilon}$ obtained from subdividing quadrilaterals by diagonals a sequence of Pachner moves transforming this triangulation into $\text{st}(v)_{=f(v)}$. The number of necessary steps in the sequence is bounded by the number of triangles in the triangulation $\text{st}(v)_{=f(v)+\epsilon}$, hence at most twice the number of 2-dimensional cells in $\text{st}(v)_{=f(v)+\epsilon}$ itself. How can we find combinatorially equivalent subdivisions of two complexes and bound their sizes when a sequence of Pachner moves transforming one complex into the other is given? For a single Pachner move the combinatorially equivalent subdivisions of the initial and the resulting complex are quite obvious. For two Pachner moves we have to perform the following steps: First intersect the subdivisions for the first and the second move in the intermediate complex, then triangulate the result, and finally map the common subdivision in the intermediate complex to the initial and the final complex for obtaining combinatorially equivalent subdivisions of all three complexes. With each additional Pachner move, we have to intersect and triangulate iteratively the subdivisions induced by the additional move with the subdivisions obtained from the previous moves.

Let us start by considering the subdivisions induced by single moves. There are three 2-dimensional Pachner moves. A (1 - 3) move replaces a triangle of the initial complex by a subcomplex formed by three triangles, three edges, and a vertex, where the triangles and edges are arranged cyclically around the vertex and the faces of the triangles and edges opposite to the vertex are identified with the boundary of the original triangle in cyclic order. Hence the resulting complex is combinatorially equivalent to a subdivision of the initial complex obtained by an elementary starring at an interior point of the original triangle. A (3 - 1) move is inverse to a (1 - 3) move. It replaces three triangles arranged cyclically around a vertex, their common vertex, and the three edges adjacent to the vertex by a single triangle filling in the whole arising from removing the seven cells. Here the roles of initial and resulting complex switch. The initial complex is combinatorially equivalent to a subdivision of the resulting complex obtained from starring at an interior point of the inserted triangle. Finally, a (2 - 2) move removes two triangles sharing an edge and inserts in the resulting quadrilateral hole two other triangles sharing an edge in such a way that the common edge of the original triangles and the common edge of the new triangles connect opposite pairs of vertices on the quadrilateral boundary. When the two triangles lie in

a common plane and form a convex quadrilateral, this move amounts to a flip of the diagonal subdividing the quadrilateral. These two triangulations of the quadrilateral have a common subdivision, which is visualised by drawing both diagonals: It consists of the intersection point of the diagonals as vertex and four triangles and edges arranged in cyclic order around the vertex, so that their faces opposite to the vertex form the boundary of the quadrilateral. From this example we can easily read off the way to obtain combinatorially equivalent subdivisions in the general case. When we perform in the initial complex an elementary starring at an interior point of the common edge of the removed triangles, and in the resulting complex an elementary starring at an interior point of the common edge of the inserted triangles, we obtain combinatorially equivalent subdivisions of the initial and the resulting complex. In the common subdivisions the pairs of triangles in question together with their common edges are each subdivided into four triangles and edges arranged cyclically around a common new vertex, with one adjacent pair of triangles with their common edge subdividing one of the triangles and the other adjacent pair with their common edge subdividing the other triangle.

Now assume that we have inductively constructed combinatorially equivalent subdivisions of the first and the κ -th complex in the sequence of Pachner moves. We can construct combinatorially equivalent subdivisions of the first and the $(\kappa + 1)$ -th complex in the sequence as follows: Consider the κ -th complex in the sequence and two of its subdivisions, namely the assumed subdivision combinatorially equivalent to some subdivision of the first complex, and the subdivision associated as outlined above with the next Pachner producing the $(\kappa + 1)$ -th complex. Take the intersection complex of these two subdivisions and triangulate it by starring at each vertex. Map this refined subdivision of the κ -th complex via the given combinatorial equivalence to a combinatorially equivalent refined subdivision of the first complex and by the combinatorial equivalence associated with the Pachner move as described above to a refined subdivision of the $(\kappa + 1)$ -th complex. Iterating this step finally yields combinatorially equivalent subdivisions of the initial and the final complex in the sequence of Pachner moves.

For obtaining a bound on the number of cells in the resulting combinatorially equivalent subdivisions, we bound the increase in the number of cells caused by an inductive refinement step. The size of the resulting combinatorially equivalent subdivisions of the first and the $(\kappa + 1)$ -th complex agrees with the size of the refined subdivision of the κ -th complex, which is the triangulation of the intersection complex of the inductively given subdivision and the subdivision associated with the Pachner move. We distinguish three cases depending on the type of the Pachner move that is performed when going from the κ -th to the $(\kappa + 1)$ -th complex. For a $(3 - 1)$ move, no further subdivision of the κ -th complex is necessary, hence the size of combinatorially equivalent subdivisions for the first and the $(\kappa + 1)$ -th complex can be bounded by the size of combinatorially equivalent subdivisions for the first and the κ -th complex. A $(2 - 2)$ move is performed on two triangles with a common edge in the κ -th complex. In the given subdivision of the κ -th complex combinatorially equivalent to the first complex, there is a subcomplex that forms a subdivision of the two triangles and their common edge. Intersecting this subdivision with the subdivision associated with the $(2 - 2)$ move, which can be visualised as cutting each of the triangles by an edge into two triangles, has the effect of cutting some of the

triangles and edges into two parts, namely all those triangles and edges whose interior intersect one of the two cutting edges associated with the (2 - 2) move. This produces up to three new cells for each cell that is cut: Edges are subdivided into two sub-edges and their common vertex, triangles are cut either into two triangles or into a triangle and a quadrilateral, and an edge separating the two parts. Since the quadrilaterals are triangulated afterwards, each cell of the given subdivision corresponds to at most five cells in the refined subdivision of the κ -th complex. In a (3 - 1) move, a triangle of the κ -th complex is subdivided into seven new cells, namely one vertex, three edges, and three triangles. Therefore the cell of the given subdivision containing the vertex is subdivided into up to seven cells in the intersection complex, all other cells correspond to at most six cells. At most three of those cells are 2-dimensional. They are of bounded complexity because they are intersections of 2-dimensional simplices. Hence the triangulation produces at most seven cells for each 2-dimensional cell in the intersection complex. This implies that, leaving aside the additional vertex from the subdivision associated with (3 - 1) move, each cell of the given subdivision contributes at most 24 cells to the refined subdivision. This is the highest possible increase among all Pachner moves.

Summing up, if the given subdivision of the κ -th complex has at most n_κ cells, the size of the refined subdivision is bounded by $24n_\kappa + 1$. Iterating this argument, the size of combinatorially equivalent subdivisions for the first and the last complex in a sequence of 2-dimensional Pachner moves that we obtain from the construction above is bounded by $24^t n + t$ where n is the number of cells in the initial complex and t is the number of Pachner moves. Note that this bound takes into account that the additional cell that causes the addition of 1 in the term $24n_\kappa + 1$ is a vertex and is not subdivided in the subsequent steps. Moreover, since the reverse sequence of the inverse Pachner moves transforms the last complex into the first, we can choose for n the minimum size of the first and the last complex. For our application this implies a bound of $24^t n + t$ on the size of combinatorially equivalent subdivisions of $\text{st}(v)_{=f(v)+\epsilon}$ and $\text{st}(v)_{=f(v)}$ where n is the minimum of the sizes of $\text{st}(v)_{=f(v)+\epsilon}$ and $\text{st}(v)_{=f(v)}$ and t is twice the number of quadrilaterals plus the number of triangles in $\text{st}(v)_{=f(v)+\epsilon}$.

Completing the Construction. The remaining steps for constructing combinatorially equivalent simplicial subdivisions of $M_{=f(v)} \times [f(v), f(v) + \epsilon]$ and $M_{\in[f(v), f(v)+\epsilon]}$ from the combinatorially equivalent subdivisions of $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ are performed as outlined earlier in this section. The complexes $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ and $\text{del}(v)_{\in[f(v), f(v)+\epsilon]}$ are combinatorially equivalent polytopal complexes and are subdivided into combinatorially equivalent simplicial complexes by starring at each vertex. The isotopy represented by these subdivisions restricts to an isotopy on $\text{lk}(v)$, and the induced homeomorphism between $\text{lk}(v)_{=f(v)}$ and $\text{lk}(v)_{=f(v)+\epsilon}$ is the natural one that matches with the restriction of the homeomorphism between $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$. In fact, the naturally combinatorially equivalent 1-dimensional complexes $\text{lk}(v)_{=f(v)}$ and $\text{lk}(v)_{=f(v)+\epsilon}$ remain unaffected by all construction steps so far. Hence we can take the union Θ_1 of the triangulation of $\text{lk}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ and the subdivision of $\text{st}(v)_{=f(v)} \times \{f(v) + \epsilon\}$. Analogously, we form the union Θ_2 of the triangulation of $\text{lk}(v)_{\in[f(v), f(v)+\epsilon]}$ and the subdivision of $\text{st}(v)_{=f(v)+\epsilon}$. Then we take the cones $T_1 = (v, f(v))\Theta_1$ and $T_2 = v\Theta_2$. The

union of T_1 and the triangulation of $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ is the desired subdivision of $M_{=f(v)} \times [f(v), f(v) + \epsilon]$, and the union of T_2 and the triangulation of $\text{del}(v)_{\in[f(v), f(v) + \epsilon]}$ is the desired subdivision of $M_{\in[f(v), f(v) + \epsilon]}$.

Before we can bound the total number of cells in the constructed subdivision, we have to analyse the effect of triangulating $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ and its combinatorially equivalent counterpart $\text{del}(v)_{\in[f(v), f(v) + \epsilon]}$ by starring at each vertex. The 2-dimensional cells in these complexes are either triangles or quadrilaterals. Hence each one produces at most three cells in the subdivision. The 3-dimensional cells are combinatorially equivalent to a Cartesian product of an interval with a triangle or quadrilateral, in other words to triangular or quadrilateral prisms. The first starring at a vertex of a triangular prism subdivides the interior into a triangular and a quadrilateral pyramid and a triangle common to both. Among the remaining starrings, only the first one that is performed at the base of the quadrilateral pyramid changes the subdivision. It replaces the quadrilateral pyramid by two triangular pyramids, in other words tetrahedra, separated by another new triangle. Thus we obtain in total five cells replacing a triangular prism, namely three tetrahedra and two triangles. A similar analysis applies to quadrilateral prisms. The first starring produces three quadrilateral pyramids and three triangles arranged cyclically around a common edge. During the subsequent starrings the quadrilateral pyramids are each subdivided into two tetrahedra and a triangle. Thus we obtain six tetrahedra, six triangles, and an edge as replacement for a quadrilateral prism, contributing 13 cells in total.

For bounding the total number of cells in the subdivision of $M_{\in[f(v), f(v) + \epsilon]}$, we count as follows: The size of combinatorially equivalent subdivisions of $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v) + \epsilon}$ was bounded by the first approach by 7 times the product of their sizes. In other words, the number of cells that each cell of $\text{st}(v)_{=f(v) + \epsilon}$ contributes to these combinatorially equivalent subdivisions is at most 7 times the number of cells in $\text{st}(v)_{=f(v)}$. Bounding the number of cells in $\text{st}(v)_{=f(v)}$ by n_M , we obtain that each cell of $\text{st}(v)_{=f(v) + \epsilon}$ contributes at most $7n_M$ cells to the subdivision of $\text{st}(v)_{=f(v) + \epsilon}$ forming one part of Θ_2 . The other part of Θ_2 is obtained from triangulating $\text{lk}(v)_{\in[f(v), f(v) + \epsilon]}$. Only the 2-dimensional cells, which are quadrilaterals, are subdivided. Thus each cell in $\text{lk}(v)_{\in[f(v), f(v) + \epsilon]}$ induces at most 3 cells in Θ_2 . Using $3 \leq 7n_M$, we can bound the number of cells in Θ_2 by $7n_M$ times the number of cells in $\text{st}(v)_{=f(v) + \epsilon} \cup \text{lk}(v)_{\in[f(v), f(v) + \epsilon]}$. The number of cells in a cone is twice the number of cells in its base. Thus, the number of cells in T_2 can be bounded by $14n_M$ times the number of cells in the complex $\text{st}(v)_{=f(v) + \epsilon} \cup \text{lk}(v)_{\in[f(v), f(v) + \epsilon]}$, which is a subcomplex of $\text{st}(v)_{\in[f(v), f(v) + \epsilon]}$. The cells of the subdivision of $M_{\in[f(v), f(v) + \epsilon]}$ lying outside of T_2 come from triangulating cells of $\text{del}(v)_{\in[f(v), f(v) + \epsilon]}$ lying outside of $\text{st}(v)_{\in[f(v), f(v) + \epsilon]}$. The triangulation produces at most 13 cells for each original cell. The factor 13 for cells outside of $\text{st}(v)_{\in[f(v), f(v) + \epsilon]}$ is dominated by the factor $14n_M$ for cells inside of $\text{st}(v)_{\in[f(v), f(v) + \epsilon]}$. Therefore the number of cells in the subdivision of $M_{\in[f(v), f(v) + \epsilon]}$ can be bounded by $14n_M$ times the number of cells in $M_{\in[f(v), f(v) + \epsilon]}$. This original complex $M_{\in[f(v), f(v) + \epsilon]}$ in turn has at most $3n_M$ cells, yielding the upper bound $42n_M^2$ for the number of cells in combinatorially equivalent subdivisions of $M_{=f(v)} \times [f(v), f(v) + \epsilon]$ and $M_{\in[f(v), f(v) + \epsilon]}$.

In the recursion for merging isotopies on subintervals, we assumed that we have combinatorially equivalent subdivisions for $M_{\in[f(v) - \epsilon', f(v) + \epsilon]}$ and $M_{=f(v)} \times$

$[f(v) - \epsilon', f(v) + \epsilon]$. We can easily arrange that the subdivisions for the interval $[f(v), f(v) + \epsilon]$ and for the interval $[f(v) - \epsilon', f(v)]$ match at the common level $f(v)$ by a careful choice of the triangulations of the level set cells in the deletion. It suffices that the triangulation of $\text{del}(v)_{\in [f(v) - \epsilon', f(v)]}$ and the triangulation of $\text{del}(v)_{\in [f(v), f(v) + \epsilon]}$ both induce the same triangulation of $\text{del}(v)_{=f(v)}$. This can be achieved by choosing the orders of the vertices, in which the starrings at the vertices are performed, in such a way that the restrictions of the sequences to the vertices of $\text{del}(v)_{=f(v)}$ yields the same order for both sequences. In that case we can take the unions of the subdivisions for the intervals $[f(v) - \epsilon', f(v)]$ and $[f(v), f(v) + \epsilon]$ as subdivisions for the interval $[f(v) - \epsilon', f(v) + \epsilon]$, so that the number of necessary cells can be bounded by $84n_M^2$. This bound can be used as value for $n_I(0)$ in the formula bounding the number of cells produced by the recursive merging of isotopies on subintervals. For the other required value $n_h(0)$ we have to find a suitable upper bound on the number of cells of the induced subdivision on $M_{=f(v)+\epsilon}$. As observed above, each cell in $\text{st}(v)_{=f(v)+\epsilon}$ contributes at most $7n_M$ cells to the subdivision of $\text{st}(v)_{=f(v)+\epsilon}$. The cells in the subdivision of $M_{=f(v)+\epsilon}$ lying outside of $\text{st}(v)_{=f(v)+\epsilon}$ come from triangulating the cells in $\text{del}(v)_{=f(v)+\epsilon}$, which requires only subdividing quadrilaterals. So each cell in $\text{del}(v)_{=f(v)+\epsilon}$ contributes at most three cells to the subdivision. Hence the number of cells in the subdivision of $M_{=f(v)+\epsilon}$ is at most $7n_M$ times the original number of cells in $M_{=f(v)+\epsilon}$. This allows to use $7n_M^2$ as upper bound on the value $n_h(0)$.

6.2.6 Dimension 4

The last remaining dimension where we have an algorithm at our disposal for finding combinatorially equivalent subdivisions of $\text{st}(v)_{=f(v)}$ and $\text{st}(v)_{=f(v)+\epsilon}$ is dimension 4. Just like the second approach in the 3-dimensional case, the algorithm considered here uses Pachner moves to construct the subdivisions. The following result of Mijatović [35, Theorem 5.2] that bounds the number of necessary Pachner moves allows to derive a bound on the number of cells of corresponding combinatorially equivalent subdivisions: Any triangulation of a 3-ball with t tetrahedra can be transformed into the cone on its boundary without affecting the triangulation of the boundary by at most $6 \cdot 10^6 t^2 2^{6 \cdot 10^6} t^2$ Pachner moves. The combinatorially equivalent subdivisions and the number of their cells can be deduced from the Pachner moves in essentially the same way as outlined for the 3-dimensional case. For a single 3-dimensional Pachner move there are simple combinatorially equivalent subdivisions of the two involved complexes that add only a constant number of new cells to them. For longer sequences, we iteratively intersect the subdivisions obtained in the previous steps with the subdivision representing the next step and triangulating the resulting intersection complex. Constructing the intersection complex has the effect of subdividing each cell of the previously obtained subdivisions into at most a constant number of new polytopal cells of bounded complexity. The bounded complexity ensures that a subsequent triangulation by starring at each vertex produces only a constant number of new cells for each polytopal cell in the intersection complex. Therefore each refinement step produces at most a constant number of cells for each cell of the previously obtained subdivisions. Hence a sequence of t' three-dimensional Pachner moves starting with a complex of size n produces combinatorially equivalent subdivisions of the initial and the

final complex with at most $c^t n$ cells for some suitable constant c . Note that, when the Pachner moves do not affect the triangulation of the boundary, then forming the intersection complexes and triangulating them by starring at the vertices in the iterative steps will never affect the boundary either.

We construct the desired subdivision of $M_{\in[f(v), f(v)+\epsilon]}$ as follows: Triangulate the subcomplex $\text{del}(v)_{\in[f(v), f(v)+\epsilon]} \cup \text{st}(v)_{=f(v)+\epsilon}$ of $M_{\in[f(v), f(v)+\epsilon]}$ by starring at each vertex. Since the cells in $M_{\in[f(v), f(v)+\epsilon]}$ are of bounded complexity, each cell produces at most a constant number c' of cells in the triangulation. Find a sequence of Pachner moves according to Mijatović's result that transforms the obtained triangulation of $\text{st}(v)_{=f(v)+\epsilon}$ into the cone on its boundary. Construct from the Pachner moves a subdivision of the triangulation of $\text{st}(v)_{=f(v)+\epsilon}$ combinatorially equivalent to a subdivision of $\text{st}(v)_{=f(v)}$. Since the constructed subdivision does not change the boundary, it matches with the original triangulation of $\text{del}(v)_{\in[f(v), f(v)+\epsilon]}$ on $\text{lk}(v)_{=f(v)+\epsilon}$. So choose for Θ_2 the union of the constructed subdivision for $\text{st}(v)_{=f(v)+\epsilon}$ and the triangulation of $\text{lk}(v)_{\in[f(v), f(v)+\epsilon]}$. Then the union of the cone $T_2 = v\Theta_2$ and the triangulation of $\text{del}(v)_{\in[f(v), f(v)+\epsilon]}$ is the desired subdivision of $M_{\in[f(v), f(v)+\epsilon]}$. The combinatorially equivalent subdivision of $M_{=f(v)} \times [f(v), f(v) + \epsilon]$ is obtained analogously: Triangulate $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon] \cup \text{st}(v)_{=f(v)} \times \{f(v) + \epsilon\}$, replace $\text{st}(v)_{=f(v)} \times \{f(v) + \epsilon\}$ by the constructed subdivision combinatorially equivalent to $\text{st}(v)_{=f(v)+\epsilon}$, and take the union of the triangulation of $\text{del}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$ with the cone $T_1 = (v, f(v))\Theta_1$ on the union Θ_1 of the constructed subdivision of $\text{st}(v)_{=f(v)} \times \{f(v) + \epsilon\}$ and the triangulation of $\text{lk}(v)_{=f(v)} \times [f(v), f(v) + \epsilon]$.

For counting the cells in the constructed subdivision of $M_{\in[f(v), f(v)+\epsilon]}$ we proceed as in the 3-dimensional case. The triangulation of $\text{del}(v)_{\in[f(v), f(v)+\epsilon]} \cup \text{st}(v)_{=f(v)+\epsilon}$ of $M_{\in[f(v), f(v)+\epsilon]}$ produces at most c' cells for each of its original cells. Therefore the number of tetrahedra in the triangulation of $\text{st}(v)_{=f(v)+\epsilon}$ is at most $c'n_M$. This implies that the number of necessary Pachner moves is bounded by $6 \cdot 10^6 (c'n_M)^2 2^{6 \cdot 10^6 (c'n_M)^2}$. Hence, each cell of $\text{st}(v)_{=f(v)+\epsilon}$ induces at most $c^{6 \cdot 10^6 (c'n_M)^2} 2^{6 \cdot 10^6 (c'n_M)^2} c'$ cells in the constructed subdivision of $\text{st}(v)_{=f(v)+\epsilon}$ that is combinatorially equivalent to a subdivision of $\text{st}(v)_{=f(v)}$. For the complex Θ_2 , we can infer that each cell from $\text{st}(v)_{=f(v)+\epsilon}$ contributes at most $c^{6 \cdot 10^6 (c'n_M)^2} 2^{6 \cdot 10^6 (c'n_M)^2} c'$ cells and each cell from $\text{lk}(v)_{\in[f(v), f(v)+\epsilon]}$ contributes at most c' cells. Since each cell of Θ_2 contributes at most two cells to the cone T_2 , each cell of $\text{st}(v)_{\in[f(v), f(v) + \epsilon]}$ contributes at most $2c^{6 \cdot 10^6 (c'n_M)^2} 2^{6 \cdot 10^6 (c'n_M)^2} c'$ cells to T_2 . The cells outside of $\text{st}(v)_{\in[f(v), f(v) + \epsilon]}$ still contribute at most c' cells to T_2 . The triangulation outside of the star. Using the dominant factor for all cells, we conclude that the constructed subdivision of $M_{\in[f(v), f(v)+\epsilon]}$ has at most $2c^{6 \cdot 10^6 (c'n_M)^2} 2^{6 \cdot 10^6 (c'n_M)^2} c'$ times as many cells as $M_{\in[f(v), f(v)+\epsilon]}$ itself. Just like in the 3-dimensional case, we can arrange that the subdivisions for the intervals $[f(v) - \epsilon', f(v)]$ and $[f(v), f(v) + \epsilon]$ match at level $f(v)$ by choosing the order of the vertices for the starrings performed when the deletion is triangulated in such a way that they match for the vertices at the common level $f(v)$. Thus we can use the term $12c^{6 \cdot 10^6 (c'n_M)^2} 2^{6 \cdot 10^6 (c'n_M)^2} c'n_M$ as upper bound for $n_I(0)$ in the recursion formula for merging isotopies on subintervals. The same arguments show that $n_h(0)$ can be bounded by $c^{6 \cdot 10^6 (c'n_M)^2} 2^{6 \cdot 10^6 (c'n_M)^2} c'n_M$.

Due to the lack of suitable algorithms in higher dimensions, we cannot con-

tinue our examination regarding the number of cells necessary to represent f -level-preserving PL isotopies. If there were an upper bound on the number of 4-dimensional Pachner moves that transform a 4-ball into the cone on its boundary, it would not only allow bounding the number of cells for isotopies in 5-dimensional manifolds, but also imply a positive solution to the long outstanding problem whether the 4-sphere can be recognised algorithmically. An analogous approach by Pachner moves for even higher dimensions $d \geq 6$ is impossible because it would contradict Novikov's result that $(d - 1)$ -spheres for $d - 1 \geq 5$ are not algorithmically recognisable.

Bibliography

- [1] Karim Adiprasito and Bruno Benedetti. *Metric Geometry, Convexity and Collapsibility*. Version 4. Apr. 9, 2013. arXiv: 1107.5789v4 [math.MG].
- [2] Karim Adiprasito and Bruno Benedetti. “Tight complexes in 3-space admit perfect discrete Morse functions”. In: *European Journal of Combinatorics* 45 (2015), pp. 71–84. DOI: 10.1016/j.ejc.2014.10.002.
- [3] A. A. Agrachev, D. Pallaschke, and S. Scholtes. “On Morse theory for piecewise smooth functions”. In: *Journal of Dynamical and Control Systems* 3.4 (1997), pp. 449–469. DOI: 10.1007/BF02463278.
- [4] Thomas F. Banchoff. “Critical points and curvature for embedded polyhedra”. In: *Journal of Differential Geometry* 1.3 & 4 (1967), pp. 245–256. URL: <http://projecteuclid.org/euclid.jdg/1214428092>.
- [5] Thomas F. Banchoff. “Critical Points and Curvature for Embedded Polyhedral Surfaces”. In: *The American Mathematical Monthly* 77.5 (1970), pp. 475–485. DOI: 10.2307/2317380.
- [6] Ulrich Bauer. “Persistence in discrete Morse theory”. PhD thesis. Georg-August-Universität Göttingen, 2011.
- [7] Mladen Bestvina. “PL Morse theory”. In: *Mathematical Communications* 13.2 (2008), pp. 149–162. URL: http://hrcak.srce.hr/index.php?show=clanak&id_clanak_jezik=48930.
- [8] Ethan D. Bloch. *Polyhedral Representation of Discrete Morse Functions on Regular CW Complexes and Posets*. Version 1. Aug. 22, 2010. arXiv: 1008.3724v1 [math.CO].
- [9] Ulrich Brehm and Wolfgang Kühnel. “Combinatorial manifolds with few vertices”. In: *Topology* 26.4 (1987), pp. 465–473. DOI: 10.1016/0040-9383(87)90042-5.
- [10] Ryan Budney and Benjamin A. Burton. *Embeddings of 3-manifolds in S^4 from the point of view of the 11-tetrahedron census*. Version 5. Sept. 26, 2012. arXiv: 0810.2346v5 [math.GT].
- [11] Manoj K. Chari. “On discrete Morse functions and combinatorial decompositions”. In: *Discrete Mathematics* 217.1–3 (2000), pp. 101–113. DOI: 10.1016/S0012-365X(99)00258-7.
- [12] A. V. Chernavsky and V. P. Leksine. “Unrecognizability of manifolds”. In: *Annals of Pure and Applied Logic* 141.3 (2006), pp. 325–335. DOI: 10.1016/j.apal.2005.12.011.

- [13] Yi-Jen Chiang, Tobias Lenz, Xiang Lu, and Günter Rote. “Simple and Optimal Output-Sensitive Construction of Contour Trees Using Monotone Paths”. In: *Computational Geometry. Theory and Applications* 30.2 (2005), pp. 165–195. DOI: 10.1016/j.comgeo.2004.05.002.
- [14] Marshall M. Cohen. *A Course in Simple-Homotopy Theory*. Graduate Texts in Mathematics 10. New York Heidelberg Berlin: Springer-Verlag, 1973. x, 114 pp. ISBN: 0-387-90055-1.
- [15] John B. Conway. *A Course in Functional Analysis*. Graduate Texts in Mathematics 96. New York Heidelberg Berlin Tokyo: Springer-Verlag, 1985. xiv, 404 pp. ISBN: 0-387-96042-2.
- [16] Jesús A. De Loera, Jörg Rambau, and Francisco Santos. *Triangulations. Structures for Algorithms and Applications*. Algorithms and Computation in Mathematics 25. Berlin Heidelberg Dordrecht London New York: Springer-Verlag, 2010. xiii, 535 pp. ISBN: 978-3-642-12970-4. DOI: 10.1007/978-3-642-12971-1.
- [17] James Eells and Nicolaas H. Kuiper. “Manifolds which are like projective planes”. In: *Publications Mathématiques de l’Institut des Hautes Études Scientifiques* 14 (1962), pp. 5–46; *Introduction*. DOI: 10.1007/BF02684323; *I. On the topological structure of X*. DOI: 10.1007/BF02684324; *II. Consequences of additional structure on X*. DOI: 10.1007/BF02684325.
- [18] Robin Forman. “A User’s Guide to Discrete Morse Theory”. In: *Séminaire Lotharingien de Combinatoire* 48 (2002), Article B48c, 35 pp. URL: <http://www.mat.univie.ac.at/~slc/wpapers/s48forman.html>.
- [19] Robin Forman. “Morse Theory for Cell Complexes”. In: *Advances in Mathematics* 134.1 (1998), pp. 90–145. DOI: 10.1006/aima.1997.1650.
- [20] Michael Hartley Freedman. “The topology of four-dimensional manifolds”. In: *Journal of Differential Geometry* 17.3 (1982), pp. 357–453. URL: <http://projecteuclid.org/euclid.jdg/1214437136>.
- [21] Leslie C. Glaser. *Geometrical Combinatorial Topology*. Vol. I. Van Nostrand Reinhold Mathematical Studies 27. New York Cincinnati Toronto London Melbourne: Van Nostrand Reinhold, 1970. viii, 161 pp.
- [22] Mark Goresky and Robert MacPherson. *Stratified Morse Theory*. Ergebnisse der Mathematik und ihrer Grenzgebiete, 3. Folge 14. Berlin Heidelberg New York London Paris Tokyo: Springer-Verlag, 1988. XIV, 272 pp. ISBN: 978-3-642-71716-1.
- [23] Branko Grünbaum. *Convex Polytopes*. Pure and Applied Mathematics XVI. London New York Sydney: Interscience Publishers (John Wiley & Sons), 1967. xiv, 456 pp.
- [24] Joel Hass. *What is an Almost Normal Surface?* Version 2. Oct. 16, 2012. arXiv: 1208.0568v2 [math.GT].
- [25] Allen Hatcher. *Algebraic Topology*. Cambridge: Cambridge University Press, 2002. xii, 544 pp. ISBN: 0-521-79540-0.
- [26] John F. P. Hudson. *Piecewise Linear Topology*. Mathematics Lecture Note Series. New York Amsterdam: W. A. Benjamin, 1969. ix, 282 pp.

- [27] Michael Joswig, Frank H. Lutz, and Mimi Tsuruga. *Sphere Recognition: Heuristics and Examples*. Version 2. Feb. 3, 2015. arXiv: 1405.3848 [math.GT].
- [28] Tomasz Kaczynski, Konstantin Mischaikow, and Marian Mrozek. *Computational Homology*. Applied Mathematical Sciences 157. New York Berlin Heidelberg Hong Kong London Milan Paris Tokyo: Springer-Verlag, 2004. XVII, 480 pp. ISBN: 0-387-40853-3.
- [29] A. Kosinski. “Singularities of piecewise linear mappings. I Mappings into the real line”. In: *Bulletin of the American Mathematical Society* 68.2 (1962), pp. 110–114. DOI: 10.1090/S0002-9904-1962-10744-7.
- [30] Wolfgang Kühnel. “Triangulations of Manifolds with Few Vertices”. In: *Advances in differential geometry and topology*. Ed. by Franco Tricerri. Singapore New Jersey London Hong Kong: World Scientific, 1990, pp. 59–114. ISBN: 981-02-0494-9.
- [31] Erhard Luft. “On the combinatorial Schoenflies conjecture”. In: *Proceedings of the American Mathematical Society* 16.5 (1965), pp. 1008–1011. DOI: 10.1090/S0002-9939-1965-0184238-6.
- [32] Jiří Matoušek, Martin Tancer, and Uli Wagner. *Hardness of embedding simplicial complexes in \mathbb{R}^d* . Version 2. Apr. 22, 2009. arXiv: 0807.0336v2 [cs.CG].
- [33] Yukio Matsumoto. *An Introduction to Morse Theory*. Trans. from the Japanese by Kiki Hudson and Masahico Saito. Translations of Mathematical Monographs 208. Providence, Rhode Island: American Mathematical Society, 2002. xiv, 219 pp. ISBN: 0-8218-1022-7.
- [34] Aleksandar Mijatović. *Simplifying triangulations of S^3* . Version 1. Aug. 15, 2000. arXiv: math/0008107v1 [math.GT].
- [35] Aleksandar Mijatović. “Triangulations of Seifert fibred manifolds”. In: *Mathematische Annalen* 330.2 (2004), pp. 235–273. DOI: 10.1007/s00208-004-0547-9.
- [36] John Willard Milnor. *Morse Theory*. Fifth Printing, 1973. Annals of Mathematical Studies 51. Princeton, New Jersey: Princeton University Press, 1963. viii, 153 pp. ISBN: 0-691-08008-9.
- [37] Dmitriy Morozov. “Homological Illusions of Persistence and Stability”. PhD thesis. Duke University, 2008.
- [38] Marston Morse. “Topologically non-degenerate functions on a compact n -manifold M ”. In: *Journal d’Analyse Mathématique* 7 (1959/1960), pp. 189–208. DOI: 10.1007/BF02787685.
- [39] Ralph Tyrrell Rockafellar. *Convex Analysis*. Second Printing, 1972. Princeton Mathematical Series 28. Princeton, New Jersey: Princeton University Press, 1970. xviii, 451 pp. ISBN: 0-691-08069-0.
- [40] Dale Rolfsen. *Knots and Links*. Mathematics Lecture Series 7. Berkeley, CA: Publish or Perish, Inc., 1976. 439 pp. ISBN: 0-914098-16-0.
- [41] Günter Rote. personal communication. Sept. 2015.

- [42] Günter Rote. “Piecewise linear Morse theory”. In: *Report No. 12/2006*. Discrete Differential Geometry. (Oberwolfach, Mar. 5–11, 2006). Mathematisches Forschungsinstitut Oberwolfach, pp. 696–698; Alexander I. Bobenko, Richard W. Kenyon, John M. Sullivan, and Günter M. Ziegler. “Discrete Differential Geometry”. In: *Oberwolfach Reports* 3.1 (2006), pp. 653–728. DOI: 10.4171/OWR/2006/12.
- [43] Günter Rote and Gert Vegter. “Computational Topology: An Introduction”. In: *Effective Computational Geometry for Curves and Surfaces*. Ed. by Jean-Daniel Boissonnat and Monique Teillaud. Mathematics and Visualization. Berlin Heidelberg: Springer-Verlag, 2006. Chap. 7, pp. 277–312. ISBN: 978-3-540-33258-9. DOI: 10.1007/978-3-540-33259-6_7.
- [44] Colin P. Rourke and Brian J. Sanderson. *Introduction to Piecewise-Linear Topology*. Ergebnisse der Mathematik und ihrer Grenzgebiete 69. Berlin Heidelberg New York: Springer-Verlag, 1972. VIII, 123 pp. ISBN: 3-540-05800-1.
- [45] Joachim H. Rubinstein. “An algorithm to Recognize the 3-Sphere”. In: *Proceedings of the International Congress of Mathematicians, August 3–11, 1994, Zürich, Switzerland*. International Congress of Mathematicians 1994. (Zürich, Aug. 3–11, 1994). Ed. by S. D. Chatterji. Vol. 1. 2 vols. Basel Boston Berlin: Birkhäuser Verlag, 1995, pp. 601–611. ISBN: 3-7643-5153-5.
- [46] Abigail Thompson. “Thin Position and the Recognition Problem for S^3 ”. In: *Mathematical Research Letters* 1.5 (1994), pp. 613–630. DOI: 10.4310/MRL.1994.v1.n5.a9.
- [47] I. A. Volodin, V. E. Kuznetsov, and A. T. Fomenko. “The problem of discriminating algorithmically the standard three-dimensional sphere”. Trans. from the Russian by A. West. In: *Russian Mathematical Surveys* 29.5 (1974), pp. 71–172. DOI: 10.1070/RM1974v029n05ABEH001296.
- [48] Afra J. Zomorodian. *Topology for Computing*. Cambridge Monographs on Applied and Computational Mathematics 16. Cambridge: Cambridge University Press, 2005. xiii, 243 pp. ISBN: 0-521-83666-2.

Zusammenfassung

Klassische Morsetheorie betrachtet glatte Mannigfaltigkeiten und auf ihnen definierte Morsefunktionen, also glatte reellwertige Funktionen, deren kritische Punkte sämtlich nicht ausgeartet sind. Dazu werden abhängig von einem reellen Funktionswert die topologischen Änderungen der zugehörigen Niveaumengen untersucht. Die Ergebnisse der Theorie erlauben es, aus den lokalen Änderungen an kritischen Punkten Rückschlüsse über globale topologische Eigenschaften der Mannigfaltigkeit zu ziehen. In der vorliegenden Arbeit wird eine analoge Theorie für kombinatorische Mannigfaltigkeiten und auf ihnen definierten stückweise linearen Funktionen vorgestellt. Dabei stehen drei Themen im Vordergrund:

Zunächst ist es das Ziel der Arbeit, die Herleitung der grundlegenden Ergebnisse der klassischen Morsetheorie aus der Betrachtung von Niveaumengen Schritt für Schritt auf den stückweise linearen Fall zu übertragen und sorgfältig ausgearbeitete Beweise für alle Schritte vorzulegen. Nützliches Hilfsmittel dafür ist die zunächst vorgenommene eingehende Untersuchung, wie ein polyedrischer Komplex mit einer Abbildung, die linear auf seinen Zellen ist, für jede als Urbild eines abgeschlossenen Intervalls definierte Niveaumenge auf natürliche Weise einen polyedrischen Komplex definiert, dessen zugrundeliegender Raum mit der Niveaumenge übereinstimmt.

Als weiteres zentrales Thema der Arbeit werden verschiedene mögliche Charakterisierungen für reguläre und nicht ausgeartete kritische Punkte verglichen. Dabei erweisen sich eine Reihe der Definitionen für solche Punkte als äquivalent, aber zwei der in der Literatur vorgeschlagenen Charakterisierungen erheben graduell schwächere Forderungen an solche Punkte. In diesem Zusammenhang wird auch eine Methode vorgestellt, wie eine diskrete Morsefunktion auf einer kombinatorischen Mannigfaltigkeit in eine stückweise lineare Morsefunktion umgewandelt werden kann, deren kritische Punkte den kritischen Zellen der diskreten Morsefunktion entsprechen.

Das dritte Thema betrifft Isotopien zwischen den Niveaumengen, wie sie auch in der klassischen Morsetheorie betrachtet werden. Zumindest für ausreichend generische stückweise lineare Abbildungen auf kombinatorischen Mannigfaltigkeiten wird die Existenz von Isotopien über alle einem Intervall zugehörigen Niveaumengen hinweg nachgewiesen, wenn das Intervall keine kritischen Werte enthält.

Die Arbeit schließt mit Betrachtungen zu ausgewählten berechnungstheoretischen und algorithmischen Gesichtspunkten ab. Zum einen wird das Entscheidungsproblem erörtert, ob ein gegebener Punkt regulär ist. Zum anderen wird die algorithmische Konstruktion der Isotopie zwischen den Niveaumengen mit dem Ziel analysiert, eine obere Schranke für die Anzahl der Zellen in den kombinatorisch äquivalenten Komplexen zu finden, die die Isotopie repräsentieren.

Erklärung

Hiermit versichere ich, dass ich die Arbeit auf Grundlage der angegebenen Hilfsmittel und Hilfen selbständig verfasst habe. Ebenso versichere ich, dass die Arbeit nicht schon einmal in einem früheren Promotionsverfahren eingereicht worden ist.

Lebenslauf

Der Lebenslauf ist in der Online-Version aus Gründen des Datenschutzes nicht enthalten.