Chapter 2

The Hausdorff metric

2.1 Definition

The Hausdorff metric actually measures the distance between non empty point sets of some metric space. Let X be a metric space and δ_X its metric. For a point $x \in X$ and a non empty set $A \subseteq X$ let us first define the ***distance of x to A*** by

$$\delta_H(x, A) := \inf_{a \in A} \delta_X(x, a).$$

Then the ***Hausdorff distance between A and B*** is defined for any non empty sets $A, B \subseteq X$ as

$$\begin{aligned} \delta_H(A, B) &:= \max(\delta_{\text{asym}}(A, B), \delta_{\text{asym}}(B, A)), \quad \text{where} \\ \delta_{\text{asym}}(A, B) &:= \sup_{a \in A} \delta_H(a, B) \end{aligned}$$

denotes the so called ***asymmetric Hausdorff distance from A to B***.

2.2 Drawbacks

As mentioned before the Hausdorff metric measures the distance between point sets of some metric space. Thus for example if we want to measure the distance between curves given by some parameterization we have to discard this parameterization first and consider only the image sets of the given curves. This may be, however, confusing, as the following picture demonstrates.

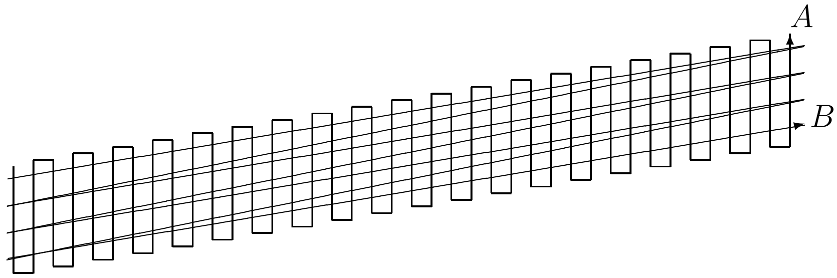

Intuitively we would say that these two curves shown above are completely different. But since both curves cover more or less the same area their corresponding image point sets have a relatively small Hausdorff distance.

2.3 Simplicial complexes

In order to describe the problems arising with respect to the Hausdorff *as well as* to the Fréchet metric (which will be discussed in the next chapter) in a unified framework we will use the concept of a simplicial complex. A ***simplicial complex*** $\mathcal{A}$ is defined to be a collection of simplices such that

1. every face of $\mathcal{A}$ is $\in \mathcal{A}$ and
2. every non empty intersection of any two simplices of $\mathcal{A}$ is a face of each of them.

The ***underlying space*** of a complex $\mathcal{A}$ is defined to be the union of the simplices in $\mathcal{A}$, hence denoted as $\cup\mathcal{A}$. Now we can describe for example polygons, unions of polyhedra, piecewise affine curves, piecewise affine triangulated surfaces and analogues in higher dimensions in a natural way by simplicial complexes. That is if we want to treat those objects as point sets (which is basically what we will do in this chapter) we consider them as the underlying space of some simplicial complex.

Note that insisting the input of an algorithm to be a simplicial complex instead of a set of simplices may hide an exponential factor in the input size where exponential is meant to be exponential in the dimension of the simplices. This is because a simplex of dimension j has exactly $2^j - 1$ sub simplices (namely the simplices spanned by every non empty set of the corners of the given simplex). Thus a simplicial complex containing simplices of dimension $\leqslant d$ may have a size which is $\sim 2^d$ times the size of a set of simplices describing the same underlying space.

2.4 Algorithms

We assume the objects to be described by a finite data structure. In fact we assume them to be simplicial complexes. So the problem is the following.

problem (h)

> We are given two non empty simplicial complexes $\mathcal{A}$ and $\mathcal{B}$ consisting of n and m simplices in $\mathbb{R}^k$, respectively. All simplices are of dimension $\leqslant d$. The task is to calculate $\delta_H(\cup\mathcal{A}, \cup\mathcal{B})$.

In this chapter we will show that there is an algorithm for that problem with a running time which is polynomial in k, n, m for any given fixed d. Assuming that n and m are asymptotically equal we will get the following run times

Solving problem (h)	
Algorithm	run time
brute force	$\mathcal{O}(n^{d+3} + k^3 n^2)$
lower envelope method	$\mathcal{O}(n^{d+2+\varepsilon} + k^3 n^2)$ randomized expected

It should be mentioned that the problem is easier for some special cases.

$d = 0$.

There is an obvious algorithm which does the job in $\mathcal{O}(n^2)$ time. But for $k \leqslant 2$ we can do it in $\mathcal{O}(n \log n)$ time the following way. Given sets A, B of at most n points in the plane we first calculate the Voronoi diagram of B, then we build a point location data structure. Then we calculate for any point $a \in A$ the Voronoi cell of a and the distance to the corresponding site. If we take the maximum of these distances we finally get $\delta_{\text{asym}}(A, B)$. Then calculating $\delta_H(A, B)$ is as easy as this is.

$d = 1$.

For $k \leqslant 2$ there is a $\mathcal{O}(n \log n)$ time algorithm given in section 2 of [ABB91].

$d \geqslant 2$.

No special improvement is known.

2.5 Framework

Instead of calculating $\delta_H(\cup\mathcal{A}, \cup\mathcal{B})$ directly we consider the following subproblem, namely

problem (h2)

> We are given a simplex T and a non empty simplicial complex $\mathcal{A}$ consisting of n simplices in $\mathbb{R}^k$, respectively. All simplices are of dimension $\leqslant d$. The task is to calculate $\delta_{\text{asym}}(T, \cup\mathcal{A})$.

By

$$\delta_H(\cup\mathcal{A}, \cup\mathcal{B}) = \max(\max_{A \in \mathcal{A}} \delta_{\text{asym}}(A, \cup\mathcal{B}), \max_{B \in \mathcal{B}} \delta_{\text{asym}}(B, \cup\mathcal{A}))$$

we can solve problem (h) in time $nt(m) + mt(n)$ assuming we have a running time of $t(n)$ for problem (h2). Consequently we will give algorithms for the problem (h2) with the following running times.

Solving problem (h2)	
Algorithm	run time
brute force	$\mathcal{O}(n^{d+2} + k^3 n)$
lower envelope method	$\mathcal{O}(n^{d+1+\varepsilon} + k^3 n)$ randomized expected

We refer to T in the following as the ***test simplex***. For each simplex $S \in \mathcal{A}$ let $\boldsymbol{A_S}$ now denote the affine span of S. Furthermore let $\boldsymbol{U_S}$ denote the linear subspace which is orthogonal to A_S. Next we define $\boldsymbol{D_S} := S + U_S$.

Then each element $x \in D_S$ has a unique representation as $a + u$ with $a \in S$ and $u \in U_S$. Then $\delta_H(x, S) = ||u||$ holds. Keeping that in mind let $\boldsymbol{d_S}\colon D_S \to \mathbb{R}_0^+$ be defined by $d_S(x) := \delta_H(x, S)$ for all $x \in D_S$. For the sake of simplicity define $\boldsymbol{\mathcal{A}(x)} := \{S \in \mathcal{A} : x \in D_S\}$ for all $x \in \mathbb{R}^k$.* We claim now that the following equality is true for all $x \in \mathbb{R}^k$.

---

* $\mathcal{A}(x)$ is a simplicial complex, by the way, but we never use this fact.

$$\delta_H(x, \cup\mathcal{A}) = \min_{S\in\mathcal{A}(x)} d_S(x) \qquad (\star)$$

Proof. Let $x \in \mathbb{R}^k$ be arbitrary. Then we have to show

- $\exists S \in \mathcal{A}(x) : d_S(x) = \delta_H(x, \cup\mathcal{A})$
  Proof. First there is certainly a simplex $S \in \mathcal{A}$ with $\delta_H(x, S) = \delta_H(x, \cup\mathcal{A})$. We choose now such a simplex with minimal dimension. First we show $x \in D_S$. We make the following case distinction:
  - S is zero dimensional. Then there is $D_S = \mathbb{R}^k$ and nothing to show.
  - S is not zero dimensional.
    Let a be the point of S closest to x. Now it is impossible that a is on the boundary of S since otherwise there would be a face S' of S on which a would be situated and then by $\delta_H(x, S) \leqslant \delta_H(x, S') \leqslant ||x - a|| = \delta_H(x, S)$ it would also be $\delta_H(x, S') = \delta_H(x, S) = \delta_H(x, \cup\mathcal{A})$ although $S' \in \mathcal{A}$ since $S \in \mathcal{A}$ and by the fact that $\mathcal{A}$ is a simplicial complex.* This would contradict to the minimality of the dimension of S. Thus $a \notin \partial S$. Then the line segment xa has to be perpendicular to S and that means $x \in D_S$.

  By $x \in D_S$ follows immediately $d_S(x) = \delta_H(x, S) = \delta_H(x, \cup\mathcal{A})$. Likewise follows $S \in \mathcal{A}(x)$ since $S \in \mathcal{A}$.
- $\forall S \in \mathcal{A}(x) : d_S(x) \geqslant \delta_H(x, \cup\mathcal{A})$
  Proof. By $S \in \mathcal{A}$ we have $S \subseteq \cup\mathcal{A}$ thus $d_S(x) = \delta_H(x, S) \geqslant \delta_H(x, \cup\mathcal{A})$.

Finally we define $\boldsymbol{\eta_S} : \mathbb{R}^k \to \mathbb{R}_0^+$ by $\eta_S(x) := (\delta_H(x, A_S))^2$. Then η_S is a quadratic form and by $(\star)$ the distance we are searching for in problem (h2) equals to

$$\sqrt{\max_{x\in T} \min_{S\in\mathcal{A}(x)} \eta_S(x)}.$$

2.6 Brute force

First we describe an algorithm for problem (h2) for which it is barely conceivable to find a more stupid one. But even that algorithm is polynomial in n for given, fixed k. And we will learn something about the structure of the problem by discussing it. In the next subsection we will show why it works. In subsection 2.6.2 we will see, how to turn this into an algorithm which is also polynomial in k. Thus in the following we will first assume that k is a constant.

D_S is a polyhedron for each $S \in \mathcal{A}$ which can be described as the intersection of at most $d + 1$ half-spaces of $\mathbb{R}^k$. We will call the hyperplanes limiting the half-spaces in the following the ***D_S* defining hyperplanes**. Let the affine span of the test simplex —we will call it ***test space*** in the following— be exactly j-dimensional. The test simplex can be described now as the intersection of exactly

*This is the only place in this chapter where we use that $\mathcal{A}$ is a simplicial complex.

$j+1$ half-spaces and $k-j$ hyperplanes. Let $\boldsymbol{\mathcal{H}}$ denote the set of the D_S defining hyperplanes *in addition to the hyperplanes limiting the half-spaces defining the test simplex* within the test space. Then we have $|\mathcal{H}| \leqslant n(d+1)+k+1 \in \mathcal{O}(n)$.

We now consider the arrangement induced by $\mathcal{H}$ within the test space. That is we subdivide the space into finitely many equivalence classes of points which are on or on the same side of the same hyperplanes of $\mathcal{H}$. Thus equivalence classes are polyhedra of various dimensions.

We now choose an equivalence class C which is contained in T with minimal dimension with the additional property that it contains a point c with $\delta_H(c, \cup\mathcal{A}) = \delta_{\text{asym}}(T, \cup\mathcal{A})$. Let X be the affine span of C and X' be the associated linear subspace. If we apply Lemma 9–10 which will be stated and proven at pages 11–16* we will see that there is a point $e \in X$, a set $M \subseteq \mathcal{A}(e)$ and a neighborhood U of e —the term neighborhood means here and in the following a neighborhood in the topological space X— such that

1. all $\eta_S(e)$ with $S \in M$ coincide,
2. e is the only point with this property within U,
3. the projections of the gradients of η_S at e onto X' for $S \in M$ establish an *affine* base of X',
4. $M \subseteq \mathcal{A}(e)$,
5. $e \in T$ and
6. $\delta_H(e, \cup\mathcal{A}) = \delta_{\text{asym}}(T, \cup\mathcal{A})$.

With this insight we search systematically for all affine subspaces X spanned by equivalence classes of the subdivision, for all $M \subseteq \mathcal{A}$ and for all $e \in X$ with the properties 1, 3, 4 and 5 and we take the maximum value of all $\delta_H(e, \cup\mathcal{A})$ obtained so far.

We perform the identification of the possible affine subspaces X in the simplest conceivable way. That is for each $i \in \{0, ..., j\}$ we intersect the test space with $j-i$ hyperplanes of $\mathcal{H}$ and we check whether we get —as we may expect in general— an i-dimensional affine subspace, which we use for X in that case. That are at most n^{j-i} possibilities for any i and we investigate each of them. Let i denote in the following the dimension of X, as before.

For the identification of the possible $M \subseteq \mathcal{A}$ we make use of the property 3 which asserts that M consist of exactly $i+1$ elements. Then we get at most n^{i+1} candidates for M and we investigate all of them.

Thus the property 1 can be expressed by $i+1$ quadratic equations of the form $\eta_S(e) = \xi$. The fact $e \in X$ leads to $k-i$ additional linear equations. Furthermore the property 3 can be expressed by the fact that an appropriate determinant of an $i \times i$-matrix of affine forms of e does not vanish. Let us call this determinant $\Delta(e)$. By introducing an additional variable this can also be written as an equation.**

*To be precise we have to make a case distinction. Either C is zero dimensional and then we have nearly nothing to prove or C is in the terminology of subsection 2.6.1 a relevant cell and then we are *indeed* able to apply Lemma 9–10.

**That is by $\Delta(e) \neq 0 \Leftrightarrow \exists \xi_2 : \xi_2 \Delta(e) = 1$.

Altogether we get a system of $k+2$ algebraic equations of a degree of at most i in $k+2$ variables.*

Lemma 10 asserts that the properties 1 and 3 imply the existence of a neighborhood U with the property 2. Thus it is clear that they are only zero dimensional solutions of the system (and therefore there are only finitely many of them (see [CLO92] for an introduction into dimensionality theorems of algebraic varieties)) and that we can determine them (see [Can87], chapter 3, algorithm 3.2 on page 56) within a time which is bounded by a function in k and that their number is bounded in the same way, say by $c(k)$.

If we do this for all $i \in \{0, ..., j\}$ then we get $(j+1) \cdot n^{j-i} \cdot n^{i+1} \cdot c(k)$ candidates for e for which we have to check whether $e \in T$ and for which if this is the case we calculate $\delta_H(e, \cup\mathcal{A})$. This can be done for each e in linear time. Altogether we calculate at most $\mathcal{O}(n^{j+1})$ many distances in time $\mathcal{O}(n^{j+2})$. Note that $j \leqslant d$ holds. We only have to take the maximum over these distances.

This completes the description of the brute force algorithm. So, if you want to see now how it can be improved, you may skip the proofs and continue reading at page 16.

---

*they are $e \in \mathbb{R}^k$ and $\xi, \xi_2 \in \mathbb{R}$.

2.6.1 Zero dimensionality lemma

For each $x \in \mathbb{R}^k$ we define $\eta(x) := \min_{S \in \mathcal{A}(x)} \eta_S(x)$. Furthermore $\varepsilon := \sqrt{\max_{x \in T} \eta(x)}$.

We were supposed to determine this ε.

In the following we refer to $\mathbb{R}^k$ as the ***base space***. Now, we are given an affine subspace of the base space, referred to as ***cell space*** and denoted with $\boldsymbol{X}$. All topological terms in this subsection refer to the topology restricted to X. Let the linear subspace associated to the cell space be referred to as the ***cell direction space*** and denoted with $\boldsymbol{X'}$. We refer to the unit sphere in the cell direction space as ***cell sphere***. If f is a real valued differentiable function on the base space then let the projection of the gradient of f onto the cell direction space be referred to as the ***cell gradient***.

For $e \in X$ and $M \subseteq \mathcal{A}$ define $\boldsymbol{L(M)} := \{x \in X : |\{\eta_S(x) \colon S \in M\}| = 1\}$. Furthermore let $\boldsymbol{G(e, M)}$ denote the set of the cell gradients of η_S at e for all $S \in M$. We describe e as being ***regular relating to*** $\boldsymbol{M}$, if and only if $e \in L(M)$ and $G(e, M)$ establishes an affine base of the cell direction space.

A set $C \subseteq X$ is said to be a ***cell***, if and only if it is compact and the set $\mathcal{A}(c)$ is the same for all $c \in C$. We say C to be ***relevant***, if and only if there is a $c \in C$ with $\eta(x) = \varepsilon^2$, but no c with this property on the boundary of the cell.

Lemma 9. If $C \subseteq X$ is a relevant cell then there is a $e \in C$ with $\eta(e) = \varepsilon^2$ and a $M \subseteq \mathcal{A}(e)$ such that e is regular relating to M.

Lemma 10. If $e \in X$ is regular relating to a $M \subseteq \mathcal{A}$ then there is a neighborhood U of e with $U \cap L(M) = \{e\}$.

The remaining part of this subsection is the proof for Lemma 9–10.

η_S is a quadratic form on the base space for each $S \in \mathcal{A}$. For each e of the base space let $\boldsymbol{\alpha_{S,e}}$ be the affine approximation* of η_S at the center of expansion e.

Lemma 1. For all e, x of the base space and each $S \in \mathcal{A}$ it holds $\alpha_{S,e}(x) \leqslant \eta_S(x) \leqslant \alpha_{S,e}(x) + ||x - e||^2$.

Proof. In the following let $\overline{e}$ denote the point of A_S closest to e and $\overline{x}$ the one closest to x. These points are (since A_S is an affine subspace) uniquely defined. Furthermore we define $\partial e := e - \overline{e}$ as well as $\partial x := x - \overline{x}$.

Now the vectors $\partial e, \partial x$ are perpendicular to A_S and thus perpendicular to $\overline{x} - \overline{e}$ as well. Let now $\beta(x) := \langle 2x - e - \overline{e}, \partial e \rangle$. Then we have at first

$$\begin{aligned}
\beta(x) &= \langle 2x - e - \overline{e}, \partial e \rangle \\
&= \langle 2x - 2\overline{x} - e + \overline{e} + 2\overline{x} - 2\overline{e}, \partial e \rangle \\
&= \langle 2\partial x - \partial e + 2\overline{x} - 2\overline{e}, \partial e \rangle \\
&= 2\langle \partial x, \partial e \rangle - \langle \partial e, \partial e \rangle + 2\langle \overline{x} - \overline{e}, \partial e \rangle \\
&= 2\langle \partial x, \partial e \rangle - \langle \partial e, \partial e \rangle \qquad \text{since } \partial e \perp \overline{x} - \overline{e}.
\end{aligned}$$

*meant to be the Taylor series up to degree 1

It follows

$$\begin{aligned}\eta_S(x) &= ||\partial x||^2 \\ &= \beta(x) + ||\partial x||^2 - \beta(x) \\ &= \beta(x) + \langle \partial x, \partial x\rangle - 2\langle \partial x, \partial e\rangle + \langle \partial e, \partial e\rangle \\ &= \beta(x) + \langle \partial x - \partial e, \partial x - \partial e\rangle \\ &= \beta(x) + ||\partial x - \partial e||^2.\end{aligned}$$

Thus we have $\eta_S \geqslant \beta$ as well as $\eta_S(e) = \beta(e)$ and therefore η_S osculates β at e from above and since β is furthermore affine we can state with good reasons that β *is* the affine approximation of η_S at the center of expansion e. Consequently we have $\beta = \alpha_{S,e}$. Thereby we have proven the first inequality of our lemma. Additionally we have

$$\begin{aligned}\eta_S(x) &= \beta(x) + ||\partial x - \partial e||^2 \\ &\leqslant \beta(x) + ||\partial x - \partial e||^2 + ||\overline{x} - \overline{e}||^2 \\ &= \beta(x) + ||(\partial x - \partial e) + (\overline{x} - \overline{e})||^2 \qquad \text{by } \partial x - \partial e \perp \overline{x} - \overline{e} \\ &= \beta(x) + ||x - \overline{x} - e + \overline{e} + \overline{x} - \overline{e}||^2 \\ &= \beta(x) + ||x - e||^2.\end{aligned}$$

□

We define $\boldsymbol{A(e, M)} := \{x \in X \colon |\{\alpha_{S,e}(x) \colon S \in M\}| = 1\}$ for each $e \in X$ and any $M \subseteq \mathcal{A}$. Since $\alpha_{S,e}$ and η_S do coincide at e (see Lemma 1) we immediately get the following

Observation 2. $e \in A(e, M) \Leftrightarrow e \in L(M)$

Lemma 3. Let $e \in X$ and $M \subseteq \mathcal{A}$ with $A(e, M) = \{e\}$. Then there is a neighborhood U of e with $U \cap L(M) = \{e\}$.

Proof. Let $\boldsymbol{\Delta}$ denote the diameter of a set of reals, i.e. the difference between the greatest and the smallest element of the set. Consider now $f(x) := \Delta(\{\alpha_{S,e}(x) \colon S \in M\})$. Then we have $A(e, M) = \{x \in X : f(x) = 0\}$. Consider now $W := \{x \in X \colon ||x - e|| = 1\}$. Then W is compact, f is continuous (since M is a finite set) and there is an element $x \in W$ with a minimal value of $f(x)$. Let x be such an element and $\mu := f(x)$ just this value.

Assume now it would be $\mu = 0$. Then it would be $x \in A(e, M) = \{e\}$ and therefore $x = e$. This would contradict $x \in W$. Thus we have $\mu \neq 0$. Since f is never negative anyway we get $\mu > 0$.

Now we define $U := \{x \in X : ||x - e|| < \mu\}$. By $\mu > 0$ it is U an open neighborhood of e. It remains to show $U \cap L(M) = \{e\}$. First we conclude $e \in U$ since U is a neighborhood of e and $e \in L(M)$ by $e \in A(e, M)$ and Observation 2. Altogether we get $e \in U \cap L(M)$. It remains to show $u = e$ for all $u \in U \cap L(M)$.

Assume now it would be $u \neq e$. Then we could find a $w \in W$ with $u = e + ||u - e||(w - e)$. Since f is actually the maximum of finitely many differences there would be $S, T \in M$ with $\alpha_{S,e}(w) = \alpha_{T,e}(w) + f(w)$. For the sake of simplicity we define furthermore $\beta := \alpha_{S,e}$ and $\gamma := \alpha_{T,e}$. Then it would be $\beta(w) = \gamma(w) + f(w)$

and as for the rest, since by $e \in A(e,M)$ we have $f(e) = 0$, it would also be $\beta(e) = \gamma(e)$. We would get

$$\begin{array}{ll} & \eta_S(u) \\ \geqslant & \beta(u) \quad \boxed{\text{by Lemma 1, first part}} \\ = & \beta(e) + ||u-e||(\beta(w) - \beta(e)) \quad \boxed{\text{since } \beta \text{ would be affine and since } u = e + ||u-e||(w-e)} \\ = & \gamma(e) + ||u-e||(\gamma(w) + f(w) - \gamma(e)) \quad \boxed{\text{since } \beta(w) = \gamma(w) + f(w) \text{ and } \beta(e) = \gamma(e)} \\ = & \gamma(e) + ||u-e||(\gamma(w) - \gamma(e)) + ||u-e||f(w) \\ = & \gamma(u) + ||u-e||f(w) \quad \boxed{\text{since } \gamma \text{ would be affine and since } u = e + ||u-e||(w-e)} \\ > & \gamma(u) + ||u-e||^2 \quad \boxed{\text{since } ||u-e|| > 0 \text{ and } f(w) \geqslant \mu > ||u-e||} \\ \geqslant & \eta_T(u), \quad \boxed{\text{by Lemma 1, second part}} \end{array}$$

which would imply $\eta_T(u) \neq \eta_S(u)$. By $S, T \in M$ this would be a contradiction to $u \in L(M)$. □

Let $e \in X$ and $M \subseteq \mathcal{A}$ be arbitrary. If the affine span of $G(e,M)$ is the whole cell direction space* then we say that M is ***complete for e***.

Lemma 4. Assume $e \in X$ and that $M \subseteq \mathcal{A}$ is not empty with $e \in L(M)$. Then $A(e,M) = \{e\}$ is true exactly if M is complete for e.

Proof. Let therefore $e \in X$ be arbitrary and $M \subseteq \mathcal{A}$ be non empty with $e \in L(M)$. Let us fix in the following one of the certainly existing elements out of M and let us denote it with R like the word "reference". To this define $\varrho := \eta_R(e)$, furthermore $Y := X \times I\!R$ as well as $Y' := X' \times I\!R$. For an arbitrary $S \in M$ we now define

$$\begin{aligned} \beta_S(x) &:= \alpha_{S,e}(x+e) - \varrho \qquad \forall x \in X' \\ H_S &:= \{(x, \beta_S(x)) \colon x \in X'\} \\ d_S &:= D\eta_S(e) \\ n_S &:= (d_S, -1), \end{aligned}$$

where D should denote the cell gradient operator.** Then H_S is a hyperplane in Y' through 0 perpendicular to $n_S \in Y'$, hence $H_S = \{y \in Y' \colon y \perp n_S\}$.***

$$\begin{array}{rcl} x \in A(e,M) & \Leftrightarrow & \exists \nu : (x-e, \nu) \in \bigcap_{S \in M} H_S \\ A(e,M) = \{e\} & \Leftrightarrow & \bigcap_{S \in M} H_S = \{0\} \\ & \Leftrightarrow^{*4} & \langle \{n_S \colon S \in M\} \rangle = Y' \\ & \Leftrightarrow & \text{affine span of} \{d_S \colon S \in M\} = X' \end{array}$$

□

---

*which is basically the same as if we would say that the linear span of the set $\{(x,-1) : x \in G(e)\}$ is the whole Cartesian product of the cell direction space with $I\!R$. But we discuss the meaning of this product space later on.

**It is $D\eta_S(e) = 2(e - \overline{e})$, where $\overline{e}$ denotes the point of S closest to e, by the way.

***To see this let $x \in X'$ and $z \in I\!R$ be arbitrary. Then we have

$$(x,z) \perp n_S \Leftrightarrow \langle (x,z), (d_S, -1) \rangle = 0 \Leftrightarrow \langle x, d_S \rangle + \langle z, -1 \rangle = 0 \Leftrightarrow \langle x, d_S \rangle = z.$$

*4 Actually this is an exercise in linear algebra. Since it is not obvious anyway, we want to verify this for short. "⇒": If $\langle \{n_S \colon S \in M\} \rangle \neq Y'$ then there would be a $v \in Y' \setminus \{0\}$ with $v \perp n_S \ \forall S \in M$. Thus we would get $v \in \bigcap_{S \in M} H_S \setminus \{0\}$. "⇐": If there would be a $v \in \bigcap_{S \in M} H_S \setminus \{0\}$ then $v \perp n_S \ \forall S \in M$ and by $v \neq 0$ we would get $v \notin \langle \{n_S \colon S \in M\} \rangle$.

Proof of Lemma 10. Let $e \in X$ be regular relating to $M \subseteq \mathcal{A}$. Then it is $e \in L(M)$ and M complete for e. By Lemma 4 follows $A(e, M) = \{e\}$. By Lemma 3 there is indeed a neighborhood U of e with $U \cap L(M) = \{e\}$. $\square$ **Lemma 10**

Proof of Lemma 9. Well then let $C \subseteq X$ be a relevant cell. We had to show that there is an $e \in C$ with $\eta(e) = \varepsilon^2$ and a $M \subseteq \mathcal{A}(e)$ such that e is regular relating to M.

For each $c \in C$ let the set of the ***sites of c***, denoted as $\boldsymbol{S(c)}$, defined as the set of all $S \in \mathcal{A}(c)$ with $\eta_S(c) = \eta(c)$. This set is never empty, since η is defined everywhere. Next, by definition of η follows immediately $\eta(c) \leqslant \eta_S(c)$ for all $S \in \mathcal{A}(c)$.

For each $e \in X$ we define $\boldsymbol{A(e)} := \{x \in X : \alpha_{S,e}(x) = \eta(e)\ \forall S \in S(e)\}$. Then by the way $A(e)$ is an affine subspace of X with $e \in A(e)$. For each c in the interior of the C we choose $\boldsymbol{U(c)}$ to be an open ball centered at c and small enough such that for all u in the closure of $U(c)$ firstly $u \in C$ holds and secondly every site of u is a site of c as well.

Let us define $\boldsymbol{E} := \{x \in C : \eta(x) = \varepsilon^2\}$. Then E is like C compact. Then the fact that C is a relevant cell can be reformulated as the fact that E is not empty and has no point on the boundary of C. Finally we have $\eta(c) \leqslant \varepsilon^2$ for all $c \in T$. From this we immediately get the following

Observation 5. Let $c \in C$. By $\eta(c) \geqslant \varepsilon^2$ follows $c \in E$.

Lemma 6. Let $e \in E$ with the property that $A(e)$ is not zero dimensional. Then there is a null set $N(e)$ of the cell sphere* with the property that an optimal element of E can not be inside of $U(e)$, where optimal is meant with respect to an objective function which is defined as the scalar product with some vector out of the cell sphere which is not contained in $N(e)$.

Proof. Let $e \in E$ and $A(e)$ be not zero dimensional. Then let $N(e)$ be the set of all vectors of the cell sphere which are perpendicular to $A(e)$. We can do this since $A(e)$ is not zero dimensional. For the same reason $N(e)$ is a null set.

Well then we consider an objective function v defined by $v(x) := \langle v, x \rangle$ for all $x \in X$ where v is an arbitrary vector out of the cell sphere which is not contained in $N(e)$. Consequently we will refer to $v(x)$ as the ***value of x***. Next x is said to be ***in M optimal*** if and only if $x \in M$ and $v(z) \leqslant v(x)$ holds for all $z \in M$.

Let e^+ be in E optimal. Then we have to show $e^+ \notin U(e)$.

Let V denote the closure of $U(e)$ and W denote its boundary.

Since $A(e)$ was not zero dimensional by assumption and an affine subspace anyway we have $A(e) \cap W \neq \{\}$. Then there is an in $A(e) \cap W$ optimal element, let us denote it with q. There is certainly a site of q. Let S_q be such a site. Then $q \in W \subseteq V$ and S_q is not only a site of q but also of e. We get

*i.e. a set of measure zero with respect to the Lebesgue measure on the cell sphere

$$\begin{array}{rcll} \eta(q) & = & \eta_{S_q}(q) & \text{since } S_q \text{ is a site of } q \\ & \geqslant & \alpha_{S_q,e}(q) & \text{by Lemma 1, first part} \\ & = & \eta(e) & \text{since } q \in A(e) \text{ and } S_q \in S(e) \\ & = & \varepsilon^2 & \text{since } e \in E. \end{array}$$

Hence $\eta(q) \geqslant \varepsilon^2$ and $q \in V \subseteq C$ and by Observation 5 we have $q \in E$. Since e^+ is in E optimal it follows $v(q) \leqslant v(e^+)$.

By $v \notin N(e)$ we have that v is not perpendicular to $A(e)$ thus $v(e) < v(q) \leqslant v(e^+)$ holds and therefore we get $e \neq e^+$. Now we can forget q. This was, by the way, the only place where we needed $v \notin N(e)$.

Assume now it would be $e^+ \in U(e)$. By $e \neq e^+$ the ray from e going through e^+ would be uniquely defined. Let e' denote the intersection point of this ray with W. Since e, e^+, e' are distinct and would appear in this ordering along the ray it would follow from $v(e) < v(e^+)$ that $v(e^+) < v(e')$ holds. Since e^+ was in E optimal we could conclude $e' \notin E$ and by $e' \in W \subseteq C$ with Observation 5 finally $\eta(e') < \varepsilon^2$.

Now there would be certainly a site of e'. Let S be such a site. Then it would be $\eta_S(e') = \eta(e') < \varepsilon^2$. Furthermore it would hold $\eta_S(e^+) \geqslant \eta(e^+) = \varepsilon^2$ and finally (since S would be by $e' \in V$ not only a site of e' but also of e) it would be $\eta_S(e) = \eta(e) = \varepsilon^2$. Let us summarize.

$$\begin{array}{rcl} \eta_S(e) & = & \varepsilon^2 \\ \eta_S(e^+) & \geqslant & \varepsilon^2 \\ \eta_S(e') & < & \varepsilon^2 \end{array}$$

Since e, e^+, e' would be distinct and would occur in this ordering along some line and since η_S is a quadratic form it would have η_S, restricted to this line, a local maximum.* But this can not be because η_S is the square of the distance to an affine subspace.** □

Lemma 7. There is an $e \in E$, such that $A(e)$ is zero dimensional.

Proof. First we can observe that the $U(e)$ for $e \in E$ establish an open covering of E. Since E is compact there is a finite subcovering $U(e_1), ..., U(e_j)$ of E.

If Lemma 7 would be false we would be able to apply Lemma 6 to all $e_1, ..., e_j \in E$ and we could define $N := N(e_1) \cup ... \cup N(e_j)$ which then would be a null set. Hence there would be a vector v out of the cell sphere with $v \notin N$. Since E is compact there would be an in E optimal element with respect to an objective function defined as the scalar product with v. Let e be such an element.

---

*That is between e and e', but this is not important for the proof.

**To see that we can use Lemma 1, like the following way: Assume the restricted function would have a local extremum at ξ. Then the graph of the restricted function would have a horizontal tangent at ξ. This tangent is a part of the tangential hyperplane of the graph on the whole function at ξ. By Lemma 1, first part, the graph would lie *above* the tangential hyperplane, hence above the tangent. Consequently the restricted function (but not necessarily the whole function) has to have a local *minimum* at ξ.

But by Lemma 6 this e could not lie inside of any $U(e_i)$. This would contradict to $e \in E$. $\square$

Lemma 8. We have $A(e) = A(e, S(e))$ for any $e \in E$.

Proof. "$\subseteq$" is trivial. So we only prove "$\supseteq$". Hence let $x \in A(e, S(e))$ be arbitrary. Now we show $\{a_{S,e}(x) : S \in S(e)\} \subseteq \{\eta(e)\}$ from which $x \in A(e)$ will follow. Let therefore $S \in S(e)$ be arbitrary. Define furthermore $\mu := \eta(e)$ and $\mu' := \alpha_{S,e}(x)$. Then we have to show $\mu' = \mu$. Assume not. Let σ the sign of $\mu' - \mu$. By assumption it would be $\sigma(\mu' - \mu) = |\mu' - \mu| > 0$.

Now we choose $\beta > 0$ small enough such that $e' := e + \beta\sigma(x - e)$ is inside of $U(e)$. Now there would be certainly a site of e'. Let S' be such a site. By $e' \in U(e)$ it would be S' a site of e as well. Since $\alpha_{S',e}$ would be the affine approximation of $\eta_{S'}$ at the center of expansion e its value at e would be $= \eta_{S'}(e)$ and this would be $= \eta(e) = \mu$ since S' would be a site of e. Altogether it would be $\alpha_{S',e}(e) = \mu$. By $S, S' \in S(e)$ and $x \in A(e, S(e))$ it would be furthermore $\alpha_{S',e}(x) = \alpha_{S,e}(x) = \mu'$. We would get

$$
\begin{array}{rcll}
\mu & < & \mu + \beta|\mu' - \mu| & \\
& = & \mu + \beta\sigma(\mu' - \mu) & \\
& = & \alpha_{S',e}(e) + \beta\sigma(\alpha_{S',e}(x) - \alpha_{S',e}(e)) & \\
& = & \alpha_{S',e}(e') & \text{since } \alpha_{S',e} \text{ would be affine} \\
& \leqslant & \eta_{S'}(e') & \text{since Lemma 1, first part} \\
& = & \eta(e') & \text{since } S' \text{ would be a site of } e' \\
& \leqslant & \varepsilon^2 & \text{since } e' \in U(e) \subseteq C \subseteq T \\
& = & \eta(e) & \text{since } e \in E \\
& = & \mu &
\end{array}
$$

$\square$

Now for the proof of Lemma 9. By Lemma 7 there is a $e \in E$ such that $A(e)$ is zero dimensional. Then we have* $A(e) = \{e\}$. By Lemma 8 we have also $A(e, S(e)) = A(e) = \{e\}$. By Observation 2 we have $e \in L(S(e))$ as well. By Lemma 4 we now get that $S(e)$ is complete for e. Now we choose an affine base $M \subseteq S(e)$. Then we have $e \in L(S(e))$ since $e \in A(e, S(e))$ and $L(S(e)) \subseteq L(M)$ since $M \subseteq S(e)$ and therefore $e \in L(M)$. Thus e is regular relating to M. $\square$ **Lemma 9**

2.6.2 Refinements

Here we will show how to turn the brute force algorithm into an algorithm which is polynomial in k. Consequently let j denote the dimension of the test simplex T, as before. Then let its corners are denoted by $c_0, ..., c_j \in \mathbb{R}^k$. Now let Φ be a function

*because $\alpha_{S,e}$ and η_S coincide at e (see Lemma 1).

which maps each $(\alpha_1, ..., \alpha_j) \in \mathbb{R}^j$ onto $c_0 + (c_1 - c_0)\alpha_1 + ... + (c_j - c_0)\alpha_j \in \mathbb{R}^k$. Especially $\underline{T} := \{(\alpha_1, ..., \alpha_j) : \alpha_1, ..., \alpha_j \geqslant 0;\ \alpha_1 + ... + \alpha_j \leqslant 1\}$ will be mapped bijectively onto T. For each $S \in \mathcal{A}$ we define now $\underline{D}_S := \Phi^{-1}[D_S]$ as well as $\overline{\eta}_S := \eta_S \circ \Phi$. For each $\underline{x} \in \mathbb{R}^j$ we define furthermore $\underline{\mathcal{A}}(\underline{x}) := \{S \in \mathcal{A} \colon \underline{x} \in \underline{D}_S\}$. Then we have

$$\max_{\underline{x} \in \underline{T}} \min_{S \in \underline{\mathcal{A}}(\underline{x})} \overline{\eta}_S(\underline{x}) \quad = \quad \max_{x \in T} \min_{S \in \mathcal{A}(x)} \eta_S(x) \quad = \quad (\delta_{\text{asym}}(T, \cup\mathcal{A}))^2$$

and we are able to perform the calculations within $\underline{T} \subseteq \mathbb{R}^j$ and with $\underline{D}_S, \underline{\mathcal{A}}(\underline{x}), \overline{\eta}_S$ quite in the same way as before, particularly since Lemmas 9–10 also apply analogously that way.* Thus all time consuming calculations can be done in $\mathbb{R}^j$ provided that we first determine representations of $\underline{D}_S, \overline{\eta}_S$ for all $S \in \mathcal{A}$. For each S this can certainly be done in $\mathcal{O}(k^3)$ time.** So the preprocessing consumes $\mathcal{O}(nk^3)$ time.

2.7 Using the structure of the lower envelope

What we actually have calculated so far is the maximal vertex of the lower envelope of the functions η_S when restricted to $D_S \cap T$. These partial functions are algebraic and their domains are semialgebraic sets of constant degree. Thus we are able to apply theorem 7.22 of [SA95]*** which states that we can determine all vertices, edges and 2-faces of this lower envelope in randomized expected time $\mathcal{O}(n^{d+1+\varepsilon})$,*4 hence solving problem (h2) in $\mathcal{O}(n^{d+1+\varepsilon} + k^3 n)$ randomized expected time.

*The reformulated Lemmas 9–10 actually follow immediately from the original ones.

**This is a generous upper bound for the time needed by a matrix multiplication, matrix inversion or orthogonalization procedure. And more than constantly many of those operations are not necessary.

***That is on page 213, at least on the cited edition.

*4Please note that in the cited book the letter d denotes the dimension of the domain of the functions *plus one*. This is because functions are identified with their graphs thereby allowing them to be defined implicitly. We do not use this generality anyway.