Enhanced sampling methods for molecular systems: multiscale and data-driven techniques

Dissertation

zur Erlangung des Grades eines Doktors der Naturwissenschaften

am Fachbereich Physik der Freien Universität Berlin

vorgelegt von

Manuel Dibak
Berlin, 2021

Erstgutachter: Prof. Dr. Frank Noé

Zweitgutachter: Prof. Dr. Jens Eisert

Tag der Disputation: 10.12.2021

Selbstständigkeitserkläung

Ich erkläre gegenüber der Freien Universität Berlin, dass ich die vorliegende Dissertation selbstständig und ohne Benutzung anderer als der angegebenen Quellen und Hilfsmittel angefertigt habe. Die vorliegende Arbeit ist frei von Plagiaten. Alle Ausführungen, die wörtlich oder inhaltlich aus anderen Schriften entnommen sind, habe ich als solche kenntlich gemacht. Diese Dissertation wurde in gleicher oder ähnlicher Form noch in keinem früheren Promotionsverfahren eingereicht. Mit einer Prüfung meiner Arbeit durch ein Plagiatsprüfungsprogramm erkläre ich mich einverstanden.

Datum: Unterschrift:

Abstract

Simulations of molecular systems have led to significant discoveries in molecular biology. The high accuracy of these simulations enables us to understand biological functions on a molecular scale. In connection with experimental results, they have proved to be a powerful tool to investigate biological functions. While the applications for such simulations are countless, in practice it is only possible to simulate small systems due to computational limitations; reaching biologically relevant time- and length-scales is still beyond feasibility, even for the most powerful computers. This constraint is commonly known as the sampling problem. With the progress in hardware development slowing down, demand for new methods that enable reaching relevant scales is high. This thesis aims to provide new tools that help molecular simulations reach biologically relevant scales. It is split into two parts:

The first part provides new methods for rate computations in reactive systems, which can consist e.g. of a protein-ligand binding, oligomerization, or protein-protein association. The first method combines Markov state models of molecular kinetics with particle-based reaction-diffusion (PBRD) to generate a coarse-grained simulation of interacting molecules. This method conserves the characteristic kinetics of the interactions – at atomistic detail – observed in molecular dynamics simulations of the interacting molecules in close proximity. Furthermore, a method is introduced to provide realistic parameters for PBRD simulations. In particular, it enables for tuning the microscopic parameters of PBRD simulations such that experimentally obtained rates are reproduced in the dilute limit. This provides a well-defined starting point to study effects such as crowding, which are common at the cellular scale.

The second part provides new methods based on Markov chain Monte Carlo. These can be utilized to speed up the generation of equilibrium samples from the Boltzmann distribution and thus enabling faster computation of stationary observables. In biological systems, it is often observed that high barriers in the free energy landscape dramatically slow down the sampling process. To speed up computations, a whole range of methods has been developed. The latest advancements are facilitated by the recent rise of machine learning research, which provides new promising tools to approach the sampling problem from completely different angles. In this spirit a new method is introduced that aims for directly proposing transitions between regions of high populations in phase space, thus directly jumping over energetic barriers. These long-range moves are proposed by a neural network trained to generate high-efficiency moves, allowing for circumventing the slow transitions across energy barriers altogether. A second proposed method is based on the recently developed Boltzmann Generators and aims to combine these with parallel tempering in order to speed up sampling significantly. To this end, a machine learning technique is employed which generates samples close to the Boltzmann distribution at different temperatures. In both of these methods, the convergence to the correct distribution is ensured by enforcing detailed balance.

Zusammenfassung

Simulationen molekularer Systeme haben zu bedeutenden Entdeckungen in der Molekularbiologie geführt. Die hohe Genauigkeit dieser Simulationen ermöglicht es, biologische Prozesse auf molekularer Ebene zu verstehen. In Verbindung mit Experimenten haben sie sich als leistungsfähiges Werkzeug zur Untersuchung biologischer Funktionen erwiesen. Während die Anwendungen für solche Simulationen zahllos sind, ist es in der Praxis aufgrund von beschränkter Rechenleistung nur möglich, kleine Systeme zu simulieren. Das Erreichen biologisch relevanter Zeit- und Längenskalen ist selbst für die leistungsstärksten Computer noch nicht möglich. Diese Einschränkung wird allgemein als Samplingproblem bezeichnet. Da sich die Fortschritte in der Hardwareentwicklung verlangsamen, ist die Nachfrage nach neuen Methoden, die es ermöglichen, relevante Größenordnungen zu erreichen, groß. Diese Dissertation zielt darauf ab, neue Werkzeuge bereitzustellen, die molekularen Simulationen helfen, biologisch relevante Größenordnungen zu erreichen. Sie ist in zwei Teile aufgeteilt:

Der erste Teil stellt neue Methoden zur Berechnung von Raten in reaktiven Systemen vor, in diesem Kontext bestehen diese z.B. aus Protein-Ligand-Bindung, Oligomerisierung oder Protein-Protein-Assoziation. Die erste Methode kombiniert Markov-Modelle von molekularer Kinetik mit partikelbasierter Reaktionsdiffusion (PBRD), um die wechselwirkenden Moleküle auf gröberen Skalen zu simulieren. Diese Methode bewahrt die charakteristische Kinetik der Wechselwirkungen im atomaren Detail, die in Molekulardynamiksimulationen der Moleküle in unmittelbarer Nähe beobachtet wird. Darüber hinaus wird eine Methode vorgestellt, um realistische Parameter für PBRD-Simulationen zu berechnen. Insbesondere ermöglicht dies, die mikroskopischen Parameter von PBRD-Simulationen so abzustimmen, dass experimentell ermittelte Raten im verdünnten Limit reproduziert werden. Dies bietet einen wohldefinierten Startpunkt, um Effekte wie Crowding zu untersuchen, die auf zellulärer Ebene üblich sind.

Der zweite Teil bietet neue Methoden basierend auf Monte-Carlo Methoden. Diese ermöglichen es, das Erzeugen von Gleichgewichtsproben aus der Boltzmann-Verteilung zu beschleunigen und somit stationäre Observablen effizienter zu berechnen. In biologischen Systemen wird oft beobachtet, dass hohe Barrieren in der freien Energie das Erzeugen von Stichproben dramatisch verlangsamt. Um dies zu beschleunigen, wurden eine ganze Reihe von Methoden entwickelt. Die jüngsten Entwicklungen in der Forschung zum maschinellen Lernen bietet neue vielversprechende Ansätze, um das Sampling von stationären Observablen aus ganz anderen Blickwinkeln zu betrachten. In diesem Sinne wird eine neue Methode eingeführt, die darauf abzielt, direkt Übergänge zwischen Regionen mit hoher Population im Phasenraum vorzuschlagen und damit energetische Barrieren direkt zu überspringen. Diese weitreichenden Vorschläge werden von einem neuronalen Netzwerk erzeugt, das darauf trainiert ist, hocheffiziente Vorschläge zu erzeugen. Ein zweites Verfahren basiert auf den kürzlich entwickelten Boltzmann-Generatoren und zielt darauf ab, diese mit Parallel Tempering zu kombinieren. Dazu wird maschinelles Lernen verwendet, um Proben nahe der Boltzmann-Verteilung bei verschiedenen Temperaturen zu erzeugen. Bei beiden Verfahren wird die Konvergenz zur korrekten Verteilung durch die Einhaltung des detaillierten Gleichgewichts sichergestellt.

Acknowledgements

I would like to thank Prof. Dr. Frank Noé for his support throughout the last five years, especially for proposing exciting and challenging research topics and providing the supervision and expertise to finalize them. I thank all collaborators, colleagues, and friends in the group that were very supportive in all kind of matters and made the lab more than just a place to work. In particular, I would like to thank Tom and Mauricio for patiently proofreading drafts of this thesis and providing me with valuable comments. I am very grateful for the support I have received from my family, especially my parents and grandparents who were adamant in supporting me throughout my studies.

List of publications

Following publications contain the results of this work

- Manuel Dibak, Mauricio J. del Razo, David De Sancho, Christof Schütte, and Frank Noé. “MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations”. In: *The Journal of Chemical Physics* 148.21 (2018), p. 214107. DOI: `10.1063/1.5020294`
- Manuel Dibak, Christoph Fröhner, Frank Noé, and Felix Höfling. “Diffusion-influenced reaction rates in the presence of pair interactions”. In: *The Journal of Chemical Physics* 151.16 (2019), p. 164105. DOI: `10.1063/1.5124728`
- Luigi Sbailò, Manuel Dibak, and Frank Noé. “Neural mode jump Monte Carlo”. In: *The Journal of Chemical Physics* 154.7 (2021), p. 074101. DOI: `10.1063/5.0032346`
- Mauricio J. del Razo, Manuel Dibak, Christof Schütte, and Frank Noé. “Multiscale molecular kinetics by coupling Markov state models and reaction-diffusion dynamics”. In: *The Journal of Chemical Physics* 155.12 (2021), p. 124109. DOI: `10.1063/5.0060314`
- Manuel Dibak, Leon Klein, and Frank Noé. “Temperature Steerable Flows and Boltzmann Generators”. 2021. arXiv: `2108.01590`, *submitted*

The individual contribution of the authors is clarified in the respective chapters.

Contents

Part I.

Multiscale dynamics of molecular systems

1. Introduction

Understanding the interaction between molecules plays an important role in studying biological processes. In many cases, a biological function is only evident by the way molecules interact with each other and in the way they transform. Examples of such include ligand binding [1, 2], complex formation [3, 4], and oligomerization [5, 6]. These interactions at molecular resolution are fundamental since they dictate the behavior at much larger scales. Often a small change in these interactions has a large impact on the whole system; in short, these processes span and couple multiple length- and time-scales. One might think of a mutation in a crucial part of a protein that affects the living being as a whole [7].

In many cases, the impact of small changes in the interaction to the behavior of the whole system is not a priori evident, and it has to be looked at the whole picture level. Thus, the interaction at small length- and time-scales plays a crucial role in scales far beyond the molecular scale. This behavior makes simulations on these scales difficult, as they have to preserve a large level of detail, and in order to reach time- and length-scales of the process being studied, an enormous amount of computational resources has to be invested, often to the point that even with the largest of computers, the relevant scales are far beyond reach. To approach this problem, simulation methods have been developed that mimic the behavior of short timescales but operate on much larger ones, commonly approximating the behavior of the system at different levels of coarseness.

In simulations of biological systems, such as proteins, it is often observed that the configuration space Γ of a molecule consists of regions where trajectories tend to concentrate, named *metastable states* $\Omega_i \subset \Gamma$. These regions are separated by energetic barriers that are rarely crossed. This gives rise to different relaxation timescales, namely the timescales at which barrier crossings occur τ_b, and the relaxation timescale of the dynamics within the metastable state τ_m. In these cases, the timescale at which transitions between the states happen is much larger than the timescale describing the relaxation within the metastable state, namely $\tau_m \ll \tau_b$. The transitions between the metastable states then no longer exhibit a memory, thus becoming Markovian [8–10]. One interpretation of this is that, if the position within the metastable state decorrelates much faster than the transition between the metastable states, it does no longer matter *where* the metastable state is entered, as equilibrium within the metastable state is reached before any escape from it can happen thus rendering the transitions between the metastable states a memoryless stochastic process. The remaining task is then to identify the metastable state that corresponds to a certain function of the molecule under consideration. Fig. 1.2 (left) shows a trajectory of the φ-angle of the alanine dipeptide molecule. This system has two metastable regions, which can roughly be characterized by $\varphi > 0$ and $\varphi < 0$. Its dynamics exhibits a typical behavior of spending long times within metastable states

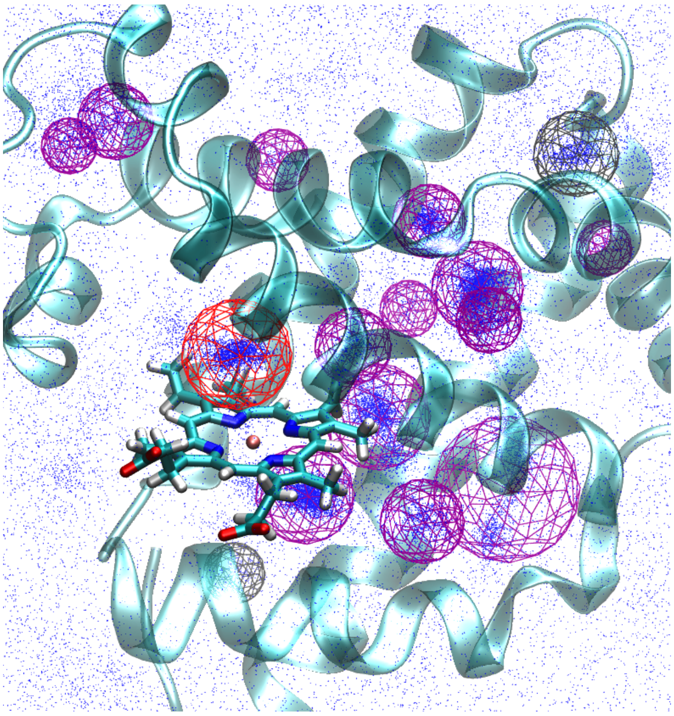

Figure 1.1.: Simulation of the binding of carbon monoxide(CO) to moyglobin. The positions at several points in time of the CO molecules are depicted in blue. The wireframe spheres show regions in space in which the CO molecules reside for a long time, thus are identified as metastable states. The *red* sphere shows the state in which the CO molecule is considered to be bound to the heme group of myoglobin, which prohibits the molecule from performing its proper function: the transport of oxygen in muscle cells.

and randomly switching between the two on a slower timescale.

Especially of interest are the timescales at which the transitions between the states happen as they give rise to a quantitative analysis of the systems at hand. These are typically expressed in terms of rates. An example of such a system is the binding of carbon monoxide (CO) to myoglobin, a molecule responsible for intra-muscular oxygen transport. Figure 1.1 shows the locations of high density areas of the CO molecules in the frame of reference of the myoglobin molecule. The binding of the CO to the heme site inhibits the transport of oxygen in the muscle, hindering it from performing its function. Understanding the rate at which these competing events occur allows to better understand the mechanism of CO-related poisoning.

The following section will discuss how binding and unbinding rates can be obtained formally.

1.1. Projection into a reaction coordinate

In a classical setup, a *reaction* can be understood as a molecular system overcoming an energy barrier in the free energy landscape. It is thus of great interest to study rates at which certain transitions occur. In this context of reactions one often defines a reaction coordinate $\mathbf{y}$ which gives a clear distinction between a system being in the educt- Ω_e or the product-space Ω_p, i.e. the reaction having occurred or not. The description of the dynamics of the system in terms of the reaction coordinate constitutes of a projection of the high-dimensional phase or configuration space $\mathbf{x} \in \Gamma$, which describes positions and momenta of all atoms involved, onto a low, possibly one dimensional coordinate on a manifold $\mathbf{y} = \int_\Gamma \delta(\mathbf{y} - f(\mathbf{x}))\mathrm{d}\mathbf{x}$. The dynamics of the reaction coordinate then give a reduced description of the dynamics of the whole system. As a result, this significant dimensionality reduction introduces memory in the friction terms of the system's dynamics and entropy terms in the effective potential $U(\mathbf{y})$ (potential of mean force) of the reaction coordinate. An exact formulation of this procedure is given by [11, 12]. While this procedure makes the formulation of the problem more tractable, the solution remains difficult due to the non-Markovian equations of motion.

In cases where noise correlations decay on very fast timescales, a Markovian approximation can be applied, i.e. a memory-less description of the dynamics of the reaction coordinate can be recovered [13], which can be described by the Langevin equation

$$\dot{\mathbf{y}} = \mathbf{v}, \quad m\dot{\mathbf{v}} = -U'(\mathbf{y}) - \gamma\mathbf{v} + \xi(t), \tag{1.1}$$

with Gaussian distributed noise with vanishing mean $\mathbb{E}\left[\xi(t)\right] = 0$ and Dirac delta distributed noise correlations $\mathbb{E}\left[\xi(t)\xi(s)\right] = 2k_\mathrm{B}T\gamma\delta(t-s)$, where m denotes the mass, k_B the Boltzmann constant, T the temperature and γ the damping coefficient. In the cellular environment, friction is usually high such that momenta of the system decay on a very fast timescale. This regime is called the overdamped limit and allows for setting the momenta terms in Eq. (1.1) to zero, greatly simplifying any calculations

$$\gamma\dot{\mathbf{y}} = -U'(\mathbf{y}) + \xi(t). \tag{1.2}$$

An equivalent description, based on the evolution of densities rather than stochastic trajectories is given by the Fokker-Planck equation. In the overdamped limit it is given by the Smoluchowski equation [14]

$$\partial_t p(\mathbf{y},t) = \partial_\mathbf{y} D\left[\beta U'(\mathbf{y})p(\mathbf{y},t) + \partial_\mathbf{y} p(\mathbf{y},t)\right], \tag{1.3}$$

with the inverse temperature $\beta = 1/k_\mathrm{B}T$ and the diffusion constant $D = \beta/\gamma$. It describes the time evolution of the probability distribution $p(\mathbf{y},t)$.

These equations are the starting point to compute the rates at which transitions across the barriers in U occur.

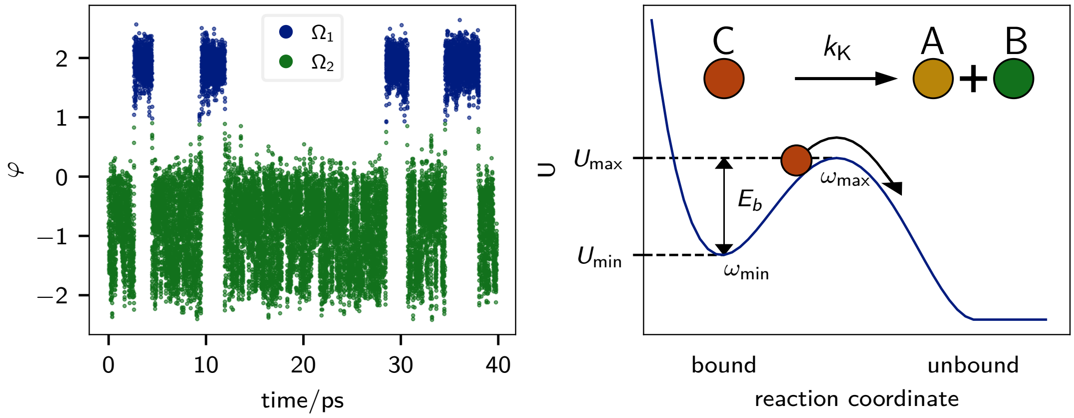

Figure 1.2.: **Left:** Transitions between metastable states from a simulation of the alanine dipeptide molecule. The time spent in the metastable states $\Omega_{1/2}$ is long compared to the relaxation times within the states. Transitions between the states only occur rarely in this situation $\tau_b \gg \tau_m$. **Right:** Setup of the Kramers rate calculations. An energetic barrier traps the system in the bound state. The system is able to escape that state at a rate depending on the temperature. Figure adapted from [15].

1.2. The Kramers problem: from potentials to rates

Given the assumptions provided above, i.e. that a reaction coordinate can be found, the friction behaves memoryless on the relevant timescales and a generalized potential U can be found which describes the system in the reaction coordinate space, a formalism to compute reaction rates was introduced by Kramers [15]. The Kramers setup studies reactions that are motivated by crossing an energy barrier in the free energy landscape. In the example before one might think of the unbinding of the CO of the myoglobin molecule as such a dissociation reaction $\mathrm{C} \to \mathrm{A} + \mathrm{B}$, where C is the bound complex and A and B are the CO and myoglobin respectively.

The Kramers approach studies mean first passage times (MFPT) $\tau = \mathbb{E}\left[t_\mathrm{FP}\right]$ across a free energy barrier, where t_FP is the time it takes a trajectory starting on one side of the barrier to cross once, and the average runs over all trajectories crossing the barrier. It is thus the average time it takes a particle to surpass an energy barrier for the first time. The MFPT is the inverse of the transition rates $k = \tau^{-1}$. Differential equations of the MFPT can be derived from the Smoluchowski equation Eq. (1.3) (see e.g. [12] for details) and for a one-dimensional setup reduces to

$$D \exp\left(\beta U(\mathbf{y})\right) \partial_\mathbf{y} \exp\left(-\beta U(\mathbf{y})\right) \partial_\mathbf{y} \tau(\mathbf{y}) = -1. \tag{1.4}$$

The problem that was studied in Kramers' original paper [15] is the case when the potential is either in the shape of a double well and the transition from one well to the other is observed or when the escape rate from one of the wells is considered, see Fig. 1.2 (right). In both cases a universal result is that if the free energy barrier U_{max} is considerably higher than the thermal energy $k_{\mathrm{B}}T$, the solution can be found by a second-order approximation of the potential using the method of steepest descent [16] to solve the integrals containing the general expressions of the potential functions.

The resulting rate in the overdamped regime is given by [15]

$$k_{\mathrm{K}} = \frac{\omega_{\mathrm{min}}\omega_{\mathrm{max}}}{2\pi\gamma}\mathrm{e}^{-\beta E_b}, \tag{1.5}$$

where $\omega_{\mathrm{min/max}}$ are the curvatures of the potential at the minimum and maximum respectively, and $E_b = (U_{\mathrm{max}} - U_{\mathrm{min}})$ is the height of the barrier. This result shows that the rate of a reaction of this type follows the Van't Hoff-Arrhenius law [17, 18] which states that the rate exhibits an exponential scaling with respect to temperature. This has significant practical implications when it comes to enhanced sampling algorithms for stationary observables, which will be discussed in Part II.

Note that finding a reaction coordinate is a highly nontrivial task and has severe limitations in practical problems. It is thus in many cases preferable to find estimates of rates from simulations of the dynamics of the system at hand. A common method are Markov state models (MSMs).

1.3. A data-driven approach to rate calculations

Given trajectories from a simulation of a dynamical system, such as molecular dynamics, rate calculations can be carried out by means of approximating the *transfer operator*. To define this operator, one has to ask: what is the probability of finding a system in phase space point $\mathbf{x}$, given that the dynamical system was started in phase space point $\mathbf{y}$ and propagated for time τ? This probability is called the *transfer probability* $p(\mathbf{y} \to \mathbf{x};\, \tau)$, and here it is used to describe the action of the underlying dynamics [19]. The propagation of probability densities $p(\mathbf{x}, t)$ in time is formally described by the *propagator* $\mathcal{P}_\tau$:

$$p(\mathbf{x}, t+\tau) = \mathcal{P}_\tau p(\mathbf{x};\, t) = \int p(\mathbf{y} \to \mathbf{x};\, \tau) p(\mathbf{y};\, t)\,\mathrm{d}\mathbf{y} \tag{1.6}$$

The goal is to approximate the propagator in a way that preserves certain aspects of the underlying dynamics. As the long-time behavior of the system is of special interest, this is the quantity that is important to be conserved.

It is often useful to consider densities relative to the stationary density $\pi(\mathbf{x})$ given by

$$u(\mathbf{x}, t) = \frac{p(\mathbf{x}, t)}{\pi(\mathbf{x})}, \tag{1.7}$$

which defines the propagator relative to the stationary density. This is the *transfer* *operator* [9]

$$u(\mathbf{x}, t+\tau) = \mathcal{T}_\tau u(\mathbf{x}, t) \tag{1.8}$$

The definition of densities relative to the stationary density allows for making use of detailed balance $\pi(\mathbf{x})p\left(\mathbf{x} \to \mathbf{y}\right) = \pi(\mathbf{y})p(\mathbf{y} \to \mathbf{x})$ to find

$$u(\mathbf{x}, t+\tau) = \mathcal{T}_\tau u(\mathbf{x}, t) = \int \frac{\pi(\mathbf{y})}{\pi(\mathbf{x})} p(\mathbf{y} \to \mathbf{x};\, \tau) u(\mathbf{y}, t)\, \mathrm{d}\mathbf{y} = \int p(\mathbf{x} \to \mathbf{y};\, \tau) u(\mathbf{y}, t)\, \mathrm{d}\mathbf{y}. \tag{1.9}$$

For reversible systems $\mathcal{T}_\tau$ is often called *backward propagator*, as it appears to evolve densities backward in time.

A propagation of multiple steps $l\tau$, with some scalar l, can be achieved by sequentially applying the transfer operator which yields the Chapman-Komlmogorov equation

$$u\left(\mathbf{x}, t+l\tau\right) = \left(\mathcal{T}_\tau\right)^l u\left(\mathbf{x}, t\right). \tag{1.10}$$

This operator has eigenfunctions $\psi_i(\mathbf{x})$ and eigenvalues λ_i: $\mathcal{T}_\tau \psi_i(\mathbf{x}) = \lambda_i \psi_i(\mathbf{x})$, with $1 \leqslant \lambda_i < -1$. Furthermore, since this operator is continuous, it posseses a continuous spectrum of eigenvalues, and sorting them in a descending order $\lambda_1 > \lambda_2 \ldots$. The largest eigenvalue is unity $\lambda_1 = 1$, with the corresponding eigenfunction being the identity $\psi_1(\mathbf{x}) = \mathbf{x}$.

Utilizing the spectral decomposition of the transfer operator, the Chapman-Kolmogorov equation can be written as

$$u(\mathbf{x}, t+l\tau) = \sum_{i=1}^{\infty} \lambda_i^l \left\langle \psi_i(\mathbf{x}), u(\mathbf{x}, t) \right\rangle_\pi \psi_i(\mathbf{x}), \tag{1.11}$$

where $\langle \cdot, \cdot \rangle_\pi$ denotes the scalar product with respect to the stationary density π. From this equation, the eigenvalues can be interpreted as timescales

$$\lambda_i(\tau) = \mathrm{e}^{-\tau/\tau_i}, \tag{1.12}$$

where τ_i denotes the relaxation timescale of the process ψ_i. With the assumption that there exists a separation between the fast and slow processes of the system under investigation, Eq. (1.11) can be truncated after the m slowest timescales thus only getting rid of the fast processes which are commonly not the subject that is studied. Note, that for this separation to hold, the system under consideration must be reversible [9]. The Chapman-Kolmogorov equation can be rewritten in terms of of the implied timescales of the m fastest processes

$$u(\mathbf{x}, t+l\tau) \approx 1 + \sum_{i=2}^{m} \mathrm{e}^{-l\tau/\tau_i} \left\langle \psi_i(\mathbf{x}), u(\mathbf{x}, t) \right\rangle_\pi \psi_i(\mathbf{x}). \tag{1.13}$$

For the truncation after m terms in Eq. (1.13) to be a good approximation, a scale separation between τ_m and τ_{m+1} is assumed, s.t. $\tau_m \gg \tau_{m+1}$. Furthermore, τ needs to be large compared to τ_{m+1}, commonly $3\tau_{m+1} < \tau$ is sufficient for practical purposes.

One way of approximating this operator is by a projection of characteristic functions that span the phase space. This is commonly denoted as an MSM that is introduced in the following section.

1.3.1. Approximating the slow eigenfunctions of the transfer operator

An approximation to the slow kinetics of the system under investigation can be found by approximating the m slowest eigenfunctions of the transfer operator. A way to do so is by Galerkin projection, which descretizes the configuration space using basis functions $\chi_i(\mathbf{x})$, $i = 1, ..., m$. In the context of MSMs these are commonly chosen to be characteristic functions

$$\chi_i(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in \Omega_i \\ 0 & \mathbf{x} \notin \Omega_i. \end{cases} \tag{1.14}$$

where the Ω_i form a complete partition of configuration space, i.e. $\Gamma = \bigcup_{i=1}^{n} \Omega_i$, and have no overlap, i.e. $\Omega_i \cap \Omega_j = \varnothing$ for $i \neq j$. The configuration space has now been discretized into a finite state space. The local densities become vectors simply given by projection onto the basis functions

$$\pi_j = \langle \pi, \chi_j \rangle = \int_{x \in \Omega_j} \pi(\mathbf{x}) d\mathbf{x}, \qquad p_j(t) = \int_{x \in \Omega_j} p(\mathbf{x}, t) d\mathbf{x}. \tag{1.15}$$

The transfer operator corresponding to this discretized phase space is a transition matrix, given by the Galerkin projection of the transfer operator to the basis functions χ_i [19]:

$$(\mathbf{T}_\tau)_{ij} = \frac{\langle \chi_j, (\mathcal{T}_\tau \chi_j) \rangle_\pi}{\langle \chi_i, \chi_i \rangle_\pi}. \tag{1.16}$$

It describes the probability of finding the system in state Ω_j at time $t + \tau$, given that it was observed in state Ω_i at time t.

The transition matrix can be estimated from trajectories of the system $\mathbf{x}(t)$. The most simple way to do so is by generating a count matrix $\mathbf{C_{ij}}$ that counts the number of times the system was found in state Ω_i at time t and in state Ω_j at time $t + \tau$. Normalizing this matrix yields an estimate of the transition matrix

$$\mathbf{T}_{ij} \approx \mathbf{C}_{ij} \sum_{i=1}^{n} C_{ij}. \tag{1.17}$$

More elaborate estimators of the transition matrix have been developed, that e.g. enforce detailed balance [20].

Using these definitions, the Chapman-Kolmogorov equation Eq. (1.13) can be written in matrix-vector notation as

$$\boldsymbol{p}(t + \tau) = \mathbf{T}_\tau^\top \boldsymbol{p}(t), \tag{1.18}$$

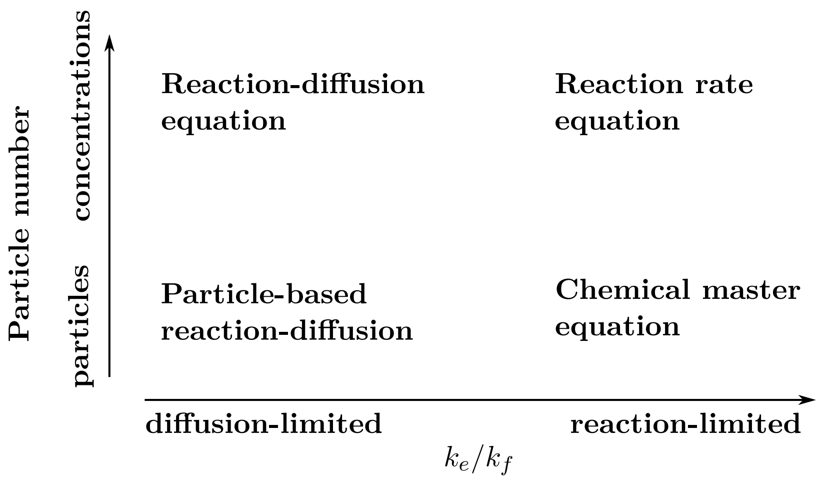

Figure 1.3.: Summary of different methods describing reaction and diffusion schemes. Four of the most relevant models for reaction-diffusion systems find their application in different regimes of the system of interest. Significant simplifications can be made if particle numbers are high and the system can be considered *well-mixed*. Figure adapted from [22]

with the probability mass vectors $\boldsymbol{p}(t) = (p_1(t), \dots, p_n(t))$. The combination of a discretization and the resulting transition matrix is denoted as an MSM.

This approach allows extracting the long-time dynamics of the system under investigation. Furthermore, it enables extracting information from a multitude of *short* trajectories, as the global equilibrium can be reconstructed from counting transitions of trajectories that have locally equilibrated [9]. This enables estimations of dynamic quantities from many short trajectories, whose length is only a couple of τ, a significant advantage when scaling to more data.

A similar description can be found in terms of rate matrices $\mathbf{K}$, with the relation between transition and rate matrices given by

$$\mathbf{T} = \exp\left(\tau \mathbf{K}\right). \tag{1.19}$$

Markov state models are therefore a suitable method for studying the transition rates between the metastable states of such systems and have in the past years become a wide-spread tool for analyzing molecular data [9, 21].

1.4. Macroscopic descriptions of reactive systems

The approach above assumes that simulations can be performed on the molecular level. For large systems, this is no longer the case and a more coarse description is required. Furthermore, the Kramers rate discussed above explains reactions on a very detailed level, namely that reactants are in close proximity. In real-world systems such as cells, another crucial part of a reaction is that reactants are able to find one another. This is commonly denoted as *reaction-diffusion* and different methods have been developed depending on the coarseness of the system under investigation. Considering for example a bi-molecular reaction of type $A + B \to C$. Then the rate k at which the product C is generated is a two-step process. Reacting particles have to *find* each other and then undergo a reaction. The total rate depends on the rate at which the molecules encounter, namely the *encounter rate* k_e, and the rate at which the transformation to molecule C occurs, the *formation rate* k_f. The encounter rate is determined by spatial diffusion of the reactants and effects such as crowding, confining geometries or interaction potentials are common factors that influence its magnitude [23]. The formation rate on the other hand is determined by the detailed chemistry of the interacting particles. Again one might think of overcoming an energy barrier once the particles are in close proximity, as studied in the Kramers problem. Under the assumption of statistical independence of the duration of both steps, the total rate is given as the harmonic mean [24–26] of the two

$$k^{-1} = k_e^{-1} + k_f^{-1}. \tag{1.20}$$

Methods modeling reactions typically distinguish along two dimensions which allow for certain simplifications

- the number of particles N
- the ratio of encounter to formation rate k_e/k_f

A brief overview of the methods is given in Fig. 1.3 and the main four pillars are discussed below.

1.4.1. Reaction-diffusion equation

In the case of many reactant particles, the fluctuations around the average number of particles vanish according to the law of large numbers and, thus, the time evolution of the number of particles of a species i in a volume V_j can be described by its particle number $N_{ij}(t) = \langle N_{ij} \rangle(t)$. In the limit of infinitely small volumes V_j this leads to a description in terms of local concentrations $c_i(r, t)$. The reactive system can be described entirely by the evolution of the local concentrations of all M particle types $C(\mathbf{r}, t) = [c_i(\mathbf{r}, t), \dots, c_M(\mathbf{r}, t)]$. This transition from particles to concentrations has recently been studied in [27]. The evolution of concentrations is governed by two operators

$$\partial_t C(\mathbf{r}, t) = \mathcal{D} C(\mathbf{r}, t) + \mathcal{R} C(\mathbf{r}, t), \tag{1.21}$$

namely the diffusion operator $\mathcal{D}$ and the reaction operator $\mathcal{R}$. The specific choice of these operators depends on the problem at hand.

This description is denoted as *reaction-diffusion equations* and is the general approach that is required especially in the case of *diffusion-influenced reactions*, which are characterized by $k_e \approx k_f$, and *diffusion-limited reactions*, where $k_e \ll k_f$.

The study of bi-molecular reactions in this regime dates back to Smoluchowski in 1917, who proposed and analyzed a model for coagulation of particles [14] which he described as a reaction happening instantaneously upon contact. This model was later extended by Debye [28] by including pair potentials between the reactants which would significantly influence the encounter rate. His study specifically aimed at including electrostatic interactions between the particles. An extension of Smoluchowski's model into a different direction was performed by Collins and Kimball [24, 29], who introduced a finite rate at which particles would undergo a reaction upon contact enabling a finite formation rate. As this model has certain drawbacks in computer simulations, i.e. it requires event-base algorithms, an alternative microscopic reaction scheme was proposed by Teramoto and Shigesada [30]. This model, which was further characterized by Doi [31–33] permits reactions of molecules with a microscopic rate λ, called *prospensity* [34], as long as they are in close proximity. This proximity is characterized as the distance between the particles being lower than a reaction radius R. This model is commonly referred to as the *volume reaction model* or simply the *Doi model.* It is subject to a detailed study in Chapter 3, where historic results, as well as the general case of volume reactions with interaction, will be discussed.

1.4.2. Reaction rate equation

In the case of slow reactions $k_e \gg k_f$, the diffusion of particles can be neglected as the formation of the new species happens on a much slower timescale compared to the encounter rate. In this *reaction-limited* regime, particles can encounter many times before any reaction is taking place. Thus the position of the particles relaxes on a timescale that allows for a description of the system solely in terms of the total concentration of each species $c_i(t) = \int c_i(\mathbf{r}, t) dV$. As a result, the evolution of the concentration can solely be described by a reaction operator $\mathcal{R}$

$$\frac{\mathrm{d}}{\mathrm{d}t} C(t) = \mathcal{R} C(t). \tag{1.22}$$

The rates might still be influenced by the diffusion of the particles. This contribution is readily incorporated into the diffusion operator. The reactive part in the *reaction rate equations* is often modeled by the law of mass action, which states, that the rate of a reaction is proportional to the product of the concentrations of the reactants [35]. One interpretation of this law is that the chance of uniformly distributed particles of reactants A and B being simultaneously in an infinitesimal volume dV is proportional to the product of the concentrations of the reactants $c_A c_B$ and furthermore the reactions of particles at close proximity occur with a fixed rate k. With the law of mass action, Eq. (1.22) takes the form of a set of ordinary differential equations. As in the example of the bimolecular reaction yielding

$$\frac{\mathrm{d}}{\mathrm{d}t} c_C(t) = k c_A(t) c_B(t). \tag{1.23}$$

This setup is often chosen when describing chemical reactions in a *well-stirred* setup which is common under laboratory conditions. In the context of biological systems, a notable application of these equations is the Michaelis-Menten kinetics [36] which describes the formation rate of an enzymatic reaction.

1.4.3. Chemical master equation

In the case of fast diffusion and low particle number, the fluctuations in the particle numbers can no longer be ignored, and instead the instantaneous particle numbers have to be tracked. Here, the description of the system in terms of concentrations breaks down, and instead, the time evolution of the number of particles $N(t) = [N_0(t), \dots, N_M(t)]$ needs to be considered. This approach is commonly denoted as the *chemical master equation*.

1.4.4. Particle-based reaction-diffusion

The fourth case, an interesting one from a biological perspective, is when both diffusion and reaction happen on a similar timescale, with a small number of particles and potentially the diffusion no longer being trivial due to diffusive barriers or crowding effects. This regime requires explicit resolution of the positions of the particles and is denoted as *particle-based reaction-diffusion* (PBRD).

For example, one might think of a simulation of a cell, a highly crowded environment, where the rate of a reaction significantly depends on the rate at which particles encounter while at the same time reactants might spend longer times in close proximity once these encounters happen. Such complex systems call for computational models to study their reactive behavior. At the same time such systems are composed of a multitude of atoms such that the computational model has to be chosen at a level of coarseness that corresponds to the length and timescales that are relevant to the system.

A computational model suitable to study such systems is interacting particle reaction dynamics (iPRD). It models the reacting molecules as point particles. This reaction-diffusion model commonly utilizes the Doi model to describe reactions on a microscopic level. In the framework of iPRD these particles are subject to interaction- and external-forces which are modeled through a potential U. The spatial motion of the reactants is resolved by overdamped Langevin dynamics. Reactions are usually distinguished between uni- and bi-molecular. Uni-molecular reactions are of the type $A \to B$, with a constant rate λ, essentially modeling a Poisson process. This type of reaction conserves the total particle number and can e.g. model spontaneous conformational changes of the molecules. Bi-molecular reaction of type $A + B \to C$ can occur with a fixed rate λ if the distance between two reacting molecules A and B is less than the *reaction radius* $R > |\mathbf{x}_A - \mathbf{x}_B|$. This type of reaction can e.g. model the binding of a ligand, or the association of proteins. Note that when the study of the equilibrium distribution is of interest, these reactions also have to be reversible, i.e. the reactions $C \to A + B$ and $B \to A$ have to have a non-zero chance and their rates have to fulfill detailed balance [37]. The spatial motion of the point particles is modeled as overdamped Langevin dynamics

$$\mu^{-1}\dot{\mathbf{x}}_i = -\nabla_{\mathbf{x}_i} U(\mathbf{x}) + \xi, \tag{1.24}$$

where μ denotes the mobility tensor, U the potential and ξ a stochastic force with zero mean and its variance following the fluctuation-dissipation theorem [12]

$$\mathbb{E}\left[\xi(t) \otimes \xi(t')\right] = 2k_{\mathrm{B}}T\mu^{-1}\delta(t-t'), \quad \mathbb{E}\left[\xi(t)\right] = 0. \tag{1.25}$$

The potential U is used in order to model both crowding effects that occur when many particles are in a relatively small volume and diffusive barriers that might occur in the system of interest, e.g. a cell wall that constrains the movement of particles.

This setup to model reactive systems has recently seen more attention as a means to simulate biological processes that span multiple length- and time-scales [6, 38].

While in principle a powerful tool, care has to be taken when interpreting results arising from this type of simulations, as they crucially depend on the parameters that are set for the reactions and interactions, namely the reaction radius R, the rate λ, and the type of interactions U.

It further has to be mentioned that in between the scenarios discussed here, there exist a multitude of methods at intermediate resolution, that e.g. compartmentalize the diffusion into boxes to efficiently describe the reaction kinetics at intermediate timescales. An exhaustive overview of current methods in reaction-diffusion can be found in Ref. [22].

In this part of the thesis, two novel methods for improving on this reaction-diffusion model are introduced:

In Chapter 2 a coupling between molecular dynamics simulations and PBRD simulations is introduced. To this end the relevant processes of the interaction kinetics of the molecules are extracted by an MSM estimated from short MD simulations. This MSM is used in the hybrid simulation to model the interaction between the molecules when they are in proximity, resulting in a jump process in the space of metastable states of the associated complex. If the molecules are far apart, they are treated as independent particles and propagated via a PBRD scheme. This novel scheme allows for the study of systems on longer time- and length-scales than MD alone while preserving the kinetics of the interaction on the molecular scale.

In Chapter 3 the relationship between the microscopic model parameters of the Doi model, extended by interaction potentials, to the macroscopic reaction rate, is derived. Relating these parameters is crucial when parametrizing quantitative iPRD simulations, as it allows for tuning the microscopic parameters of the simulation model to match e.g. experimental values in the dilute limit. This provides a well-defined starting point to study systems on the cellular scale where crowding effects are present.

2. Coupling Markov state models of molecular kinetics with reaction-diffusion simulations

The results of this chapter have been published in:

> Manuel Dibak, Mauricio J. del Razo, David De Sancho, Christof Schütte, and Frank Noé. "MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations". In: *The Journal of Chemical Physics* 148.21 (2018), p. 214107. DOI: `10.1063/1.5020294`

Parts of text and figures have been adopted unchanged in this chapter. Reproduced from *The Journal of Chemical Physics* "MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations", Dibak et al., 2018, with the permission AIP Publishing. Manuel Dibak (MD) and Mauricio J del Razo (MR) contributed equally to this work. In particular, the contributions of the authors were as follows: Frank Noé (FN), Christof Schütte, MD and MR conceived the project. MD and MR worked out the theory and the coupling scheme, performed experiments and visualized the data. David De Sancho provided molecular dynamics simulations of the myoglobin system. MD, MR and FN wrote the manuscript.

Summary

Molecular dynamics (MD) simulations are a suitable model for the interactions between macromolecules with high spatiotemporal resolution. However, their high computational cost hinders their application to systems spanning large time- and length-scales. Success in investigating the long-timescale behavior of small to intermediate biomolecules has recently been found by combining high-throughput MD with Markov state models (MSMs). However, scaling to larger length-scales or the interaction of many molecules remains challenging within this model. More suitable for these situations are particle-based reaction-diffusion (PBRD) simulations, that sacrifice molecular detail in order to reach larger spatial and temporal scales. Thus, a combination of both, i.e. a coupling between MSMs and PBRD simulations (MSM/RD) would be highly desirable, as it could efficiently produce simulations at large length- and time-scales that additionally conserve the characteristic kinetics observed at the atomistic detail. Such a scheme imitates the interaction of the molecules in proximity by an MSM derived from molecular simulations and treats particles independently if they are far away. In this chapter a first step towards

MSM/RD is introduced by laying out a general theory of coupling and proposing an implementation for association/dissociation of a protein with a small ligand $A + B \rightleftharpoons C$. The proposed model is validated on a toy system and an exemplary application is provided to the diffusion of carbon monoxide into the heme cavity of myoglobin.

2.1. Introduction

Understanding life processes on a fundamental level constitutes a big challenge as they are inherently multi scale and span multiple length and timescales. This challenge becomes even more difficult to tackle as the different scales are often tightly bounded and no single experiment or simulation can probe all time- and length-scales at a high enough resolution comprehensively. In computer simulations, this dilemma can be mitigated by multiscale techniques: different parts of the system are described by a high- and a low-resolution models, and these parts are coupled to give rise to a hybrid simulation. A famous example of such a multiscale model in biophysical chemistry is the coupling of quantum mechanics and molecular mechanics (QM/MM) [40].

This chapter lays the foundations for a hybrid simulation technique that couples two scales that are particularly useful to model intracellular dynamics: a Markov state model (MSM) of the molecular dynamics (MD) scale that describes structural changes of biomolecules and their complexes, and the reaction-diffusion scale that describes diffusion, association and dissociation on the lengthscale of a cell. This novel approach is called MSM/RD, as it combines the simulation models chosen at these scales:

(i) MSMs of the molecular scale: MD simulation allow for probing molecular processes at atomic detail, but its usefulness has long been limited by the sampling problem. Recently, the combination of hard- and software for high-throughput MD simulations [41–44] with MSMs [9, 45, 46] has enabled the extensive statistical description of protein folding and conformation changes [47–50], as well as the association of proteins with ligands [51–55] and even other proteins [56]. Using multi-ensemble Markov models (MEMMs) [57–60], MSMs can be derived that even capture the kinetics of ultra-rare events beyond the seconds timescale at atomistic resolution [2, 61]. MSM approaches can thus model the long-lived states and transition rates of molecular detail interactions, but the cost of atomistic MD sampling limits them to relatively small biomolecules and complexes.

(ii) Reaction-diffusion (RD) scale: While atomic detail is relevant for some processes that affect the cellular scale, it is neither efficient nor insightful to maintain atomic resolution at all times for cellular processes. A sensible choice for the cellular scale is particle-based reaction-diffusion (PBRD, Section 1.4.4) dynamics as a reference model. PBRD simulates particles, representing individual copies of proteins, ligands or other metabolites. Particles move in space via diffusion and reactive species will react with a probability according to their reaction rate when being close. Here, a reaction may represent molecular processes such as binding, dissociation, conformational change, or actual enzymatic reactions. PBRD acknowledges that

> chemical reactions are inherently discrete and stochastic in nature [62], and that diffusion in cells is often not fast enough to justify well-stirred reaction kinetics [63–65]. A large number of recent software packages and codes implement some form of PBRD [66–74], see also the reviews [38, 75]. The effect of crowders and complicated boundaries such as membranes on the particle diffusion can be represented by including interaction forces on the RD scale [72].

In the limit that the conformational transitions of all molecules are fast, the MSM dynamics of each molecule effectively averages, and the interaction between the molecules (e.g. association) occurs with suitably averaged rates, reducing the problem to PBRD. However, when the lifetimes of some conformations are long compared to the typical time between two molecular interactions, or even the time between successive rebinding events of two molecules, the conformation dynamics of molecules described by the MSM part couples with the PBRD dynamics. MSM/RD opens up the possibility to simulate and analyze such effects quantitatively. For example, bimolecular binding rates from MD-derived MSMs can be inaccurate due to periodic boundary effects and a short-lived dissociated state in comparison to the MSM lag-time [56]. MSM/RD can overcome these issues by extending the diffusion domain available lessening the periodic boundary effects and increasing the lifetime of the dissociated state.

The ultimate aim of MSM/RD is to produce an efficient multiscale simulation that reproduces the essential statistical behavior of a practically unaffordable large-scale MD simulation by employing only statistics obtained from simulations of the constituent biomolecules in small solvent boxes. Developing a full theory involving rotational diffusion, three- or more-body interactions, hydrodynamics would be highly complex. Here a first stab is taken towards this goal by coupling MSM and PBRD scales for bimolecular systems without large-scale hydrodynamic interactions. An extension of this work has recently been developed [76], which extends the current method to the interaction of multiple molecules and rotational motion which allows to study more complex systems, such as protein-protein association.

A theory of MSM/RD for bimolecular systems is derived as depicted in Fig. 2.1. When the two molecules are far from each other, they both undergo a diffusion process. When they come close to each other, molecular interactions, modeled with MD-derived MSMs, need to be taken into account. Furthermore, an algorithm to couple the MSM and PBRD scales is developed for the special case of a protein interacting with a ligand, which is one of the main advances in this chapter. This is not a trivial undertaking since one needs to solve two problems:

(i) couple the MSM and PBRD part in such a way that the correct macroscopic rates and equilibrium probabilities are recovered

(ii) develop a suitable MSM discretization such that this coupling can be made

The validity of the theory and algorithms is demonstrated on a toy model of protein-ligand interaction and on binding of carbon monoxide to myoglobin.

In related work, Refs. [77, 78] have coupled MD with a diffusion scheme. The work [79] further incorporates milestoning theory [80] to compute the local kinetic information

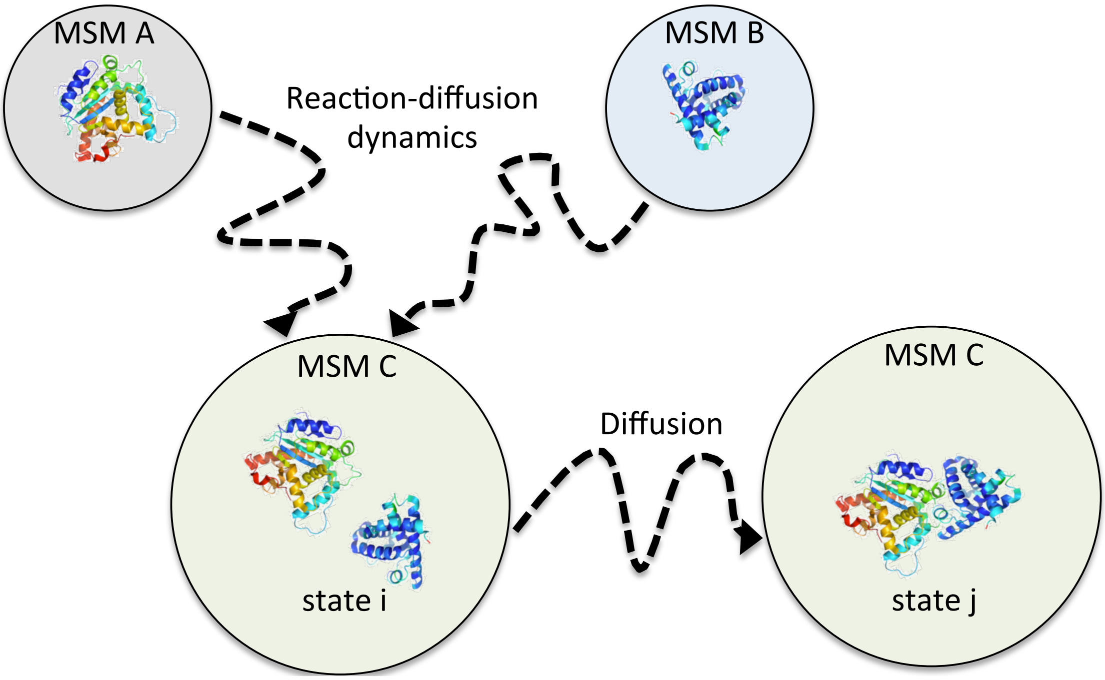

Figure 2.1.: Sketch of the MSM/RD scheme. When molecules A and B are not in close proximity, they diffuse freely. When A and B are close, they merge into a complex particle C which itself diffuses and whose internal dynamics are encoded by coupled MSM state transitions. When the molecules transition into a dissociated state, they are again separated into two separately diffusing particles A and B with initial positions depending on the last MSM state. Note that in the dissociated state, molecules A and B could also potentially undergo conformational changes encoded in independent MSM state transitions. Reprinted from *The Journal of Chemical Physics* “MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations”, Dibak et al., 2018, with the permission of AIP Publishing.

in terms of transitions between milestones via short MD runs. In contrast with this work, MSM/RD does not employ direct MD simulations at the "small" scale, but represents the small scale by an MSM as this allows for operating on roughly the same timesteps for the small and the large scales. Other works have proposed alternative schemes to couple random walks (MSMs) with Brownian diffusion schemes, some examples can be seen in [81–83]. However, these works focus on specific contexts that are not directly applicable for coupling MD-derived MSMs with reaction-diffusion schemes.

2.2. MSM/RD: coupling Markov state models and reaction-diffusion

In this section, a theoretical description for MSM/RD is developed. The relevant scenarios for MSM/RD can be classified by the number of interacting particles, or the related reaction order:

(i) First-order reactions: isolated diffusing particles can be modeled by an MSM obtained from MD simulations in a solvent box. The MSM directly translates into a set of unimolecular reactions that can be implemented in standard PBRD software. As long as the particles don't interact, the only effect of different states on the dynamics are changes between different diffusion constants/tensors.

(ii) Second-order reactions: interactions between two molecules that can be modeled as bimolecular reactions including protein-ligand or protein-protein association ($\mathrm{A} + \mathrm{B} \rightarrow \mathrm{C}$). As soon as the complex C has been formed, its dynamics may be described by state transitions of an MSM of the complex.

(iii) Higher-order reactions: simultaneous interactions between more than two molecules.

Here, the focus lays solely on the treatment of second-order reactions. First-order reactions are trivial state changes of a particle that are occurring as part of the MSM dynamics. Consistent with current conventions in PBRD frameworks, higher order reaction can be broken down to second-order reactions. An extension to treat higher-order reactions explicitly is suggested in Section 2.5.

In order to derive the theory for second-order reactions, the focus lies on the dynamics of two molecules, A and B. For the sake of simplicity, it is assumed that the two molecules do not have conformational changes of their own, so they can only diffuse and interact with each other. However, it is straightforward to extend MSM/RD to include conformational changes (first-order reactions) coupled with second-order reactions.

2.2.1. The *ground truth* model with full dynamics

Ground truth is a term often used in machine learning that refers to a reference model with respect to which modeling errors are measured. In the present context, the *ground truth* model contains the two (or more) solute molecules whose interactions will be later approximated by an MSM in a *large-scale* simulation, i.e. a simulation box that is not truncated after a small solvent boundary as customary for MD simulation. Importantly, there is no universally correct ground truth, but this model employs the MD simulation setup and dynamical model chosen by the user for the modeling task at hand. This choice includes the MD force field, solvation conditions and ion concentration, the protonation state at the pH of interest or even constant-pH simulations [84], the treatment of electrostatics, the thermostat, the integrator and time step, etc.

If such a large-scale model were simulated for a long time or with many trajectories, it would give rise to statistical properties of the solute molecules that MSM/RD aims to reproduce, such as their equilibrium constants and association rates. However, such a simulation is in general inefficient or infeasible, and the aim is to reproduce its statistical properties using an MSM/RD model that is parametrized only using small MD simulations of the constituent solute molecules and complexes.

For simplicity, the MSM/RD theory is derived using all-atom explicit solvent MD simulations with a Langevin thermostat as the *ground truth*, as this setup is frequently used for MD simulations. However, the MSM/RD results apply more generally, e.g. to different choices of thermostats or integrators, as the MSM limit for long-time description of the dynamics and the overdamped limit for long-time and large-scale description of the solute transport are achieved from a large family of ground truth models.

Langevin dynamics evolve as:

$$\partial_t \mathbf{x}_k(t) = \mathbf{v}_k(t), \quad m_k \partial_t \mathbf{v}_k(t) = -\nabla_k U\left(\mathbf{x}(t)\right) - \gamma_k \mathbf{v}_k + \boldsymbol{\xi}_k(t), \tag{2.1}$$

where $\mathbf{x}_k$ represents the three-dimensional position of the k^{th} atom in the system (including the solvent), $\mathbf{x}(t) = [\mathbf{x}_1(t), \ldots, \mathbf{x}_k(t), \ldots, \mathbf{x}_N(t)]$, N the total number of atoms, U is the potential energy and $-\nabla_k U$ is the force acting on the k^{th} particle, m_k is the k^{th} particle mass, γ_k is the k^{th} damping coefficient, and $\boldsymbol{\xi}_k(t)$ is a Gaussian random force such that the expectations of its components satisfy $\mathbb{E}[\xi_{k,i}(t)] = 0$ (zero mean) and $\mathbb{E}[\xi_{k,i}(t)\xi_{k,j}(s)] = 2k_{\mathrm{B}}T\gamma_k\delta_{ij}\delta(t-s)$ with $k_{\mathrm{B}}T$ being the thermal energy according to the fluctuation-dissipation theorem. Simulations utilize finite-time-step approximations of Eq. (2.1) and use it to generate stochastic trajectories.

A suitable description of the evolution of densities is found in terms of the transfer operator, which has been introduced in Section 1.3. Here it's discretized form is used in the framework of an MSM to model the close distance interaction of the molecules.

A suitable scale separation can be found by treating molecules A and B differently when they are close together (interacting) and far apart (non-interacting). More specifically these scales are defined by the distance between the centers of mass of A and B, r_{AB} and are introduces as:

(i) MSM domain: molecules are in the *interaction* region $I = \{\mathbf{x} \mid r_{AB}(\mathbf{x}) < R\}$.

(ii) PBRD domain: molecules are in the *outside* region $O = \{\mathbf{x} \mid r_{AB}(\mathbf{x}) \geq R\}$.

The definition of the interaction radius R will be investigated later. First, the dynamics in the respective domains is further investigated.

2.2.2. Markov state models for interacting molecules

Molecules within a distance closer than R are considered to be interacting, hence the corresponding subset of state space is called I. The kinetics in I are fully described by the transfer operator Eq. (1.8), which can be approximated by an MSM derived from a MD simulation that fully includes I (usually plus some extra space, because MD simulations typically employ periodic rather than spherical boundary conditions). This approximation implicitly assumes that the interaction forces between proteins or protein-ligand pairs have decayed to zero at distances R or greater. Note that this assumption requires that the MD simulation conducted to parametrize an MSM/RD model has a sufficiently large simulation box, suitable electrostatics treatment and solvation conditions (ions etc.) such that in the dissociated state the solutes can be in any orientation without significantly interacting with each other or with their periodic images.

The dynamics in the interaction region I will be approximated by an MSM (see Section 1.3.1). Estimating a high-quality MSM from MD simulation data can be quite complex. It typically involves:

(i) Mapping the MD coordinates to a set of features, such as residue distances, contact maps or torsion angles

(ii) reducing the dimension to slow collective variables (CVs), often based on the variational approach or conformation dynamics [85, 86] or its special case time-lagged independent component analysis (TICA) [87, 88] – see [89, 90] for an overview

(iii) (optionally) embedding the resulting coordinates in a metric space whose distances correspond to some form of dynamical distance [91, 92]

(iv) discretizing the result space using data-based clustering [45, 93, 94], typically resulting in 100-1000 discrete states

(v) estimating the transition matrix $\mathbf{T}(\tau)$ or a transition rate matrix $\mathbf{K}$ with $\mathbf{T}(\tau) = \exp(\tau\mathbf{K})$ at some lag time τ, and validating it [9, 20, 48, 95]

(vi) As a final step the MSM may be coarse-grained to few metastable states [96–99]

The MSM software packages PyEMMA [93] and MSMbuilder [94] can greatly help to simplify this process and make it reproducible.

In the case where there are well-defined meta-stable regions in phase space, one can greatly reduce the number of states in the MSM. One way to simplify the MSM construction process above and to directly end up with a few-state MSM is to employ VAMPnets, where the complex MSM construction pipeline is replaced by a neural network that is trained using the variational approach for Markov processes [100]. Alternatively, one can replace the discretization step (*iv*) above by employing a core set approach that was derived in [95] and further analyzed in [101]. The essential idea is to define the states as cores around the metastable regions. Due to the metastability, the probability of finding the system outside of the metastable regions is very small, so to a good approximation the kinetics can be described as a core-to-core jump process [101]. This approach will be employed throughout this chapter and explained in more detail in Section 2.3.

2.2.3. Reaction-diffusion dynamics for noninteracting molecules

When molecules are far apart, and thus in the PBRD domain defined by $r_{AB}(\mathbf{x}) \geq R$, they are not directly interacting. As the dynamics of the two molecules are independent, it is convenient to only track the net diffusion of the centers of mass, $\boldsymbol{r}_A$ and $\boldsymbol{r}_B$. Furthermore, it is assumed that the dynamics in the PBRD domain can be tracked by coarse timesteps of at least Δt which exceeds the typical velocity autocorrelation time (picoseconds). At such timescales, the fast dynamics corresponding to the solvent are averaged out. It is possible that even longer timesteps are made using an event-based integration scheme such as first-passage kinetic Monte Carlo (FPKMC) algorithm, Green's function reaction dynamics (GFRD)[77] or MD-GFRD [65, 68, 73, 74, 78]. At such timesteps, the Langevin equation (2.1) becomes an overdamped Langevin equation for the centers of mass of the two molecules, i.e. the motion is governed by free diffusion:

$$\partial_t \mathbf{r}_A(t) = \boldsymbol{\xi}_A(t), \quad \partial_t \mathbf{r}_B(t) = \boldsymbol{\xi}_B(t), \tag{2.2}$$

where $\boldsymbol{\xi}_A(t)$ and $\boldsymbol{\xi}_B(t)$ are independent white noise vectors with each of their components satisfying $\mathbb{E}[\xi_{K,i}(t)] = 0$ and $\mathbb{E}[\xi_{K,i}(t)\xi_{K,j}(s)] = 2D_K\delta_{ij}\delta(t-s)$, with $K \in \{\mathrm{A}, \mathrm{B}\}$. D_A and D_B are the net diffusion coefficients for the centers of mass, which can be obtained e.g. from MD simulations. Keeping track of only the centers of mass is not sufficient in cases where the orientation of particles is of significance to the binding process. In these cases the rotational diffusion has also taken into account in the PBRD domain. However, as rotational diffusion is not relevant for the examples discussed in this chapter, it is not further discussed here. For an extension to rotating interaction see [76] or more generally [102, 103].

In the present case the frame of reference is simply fixed in $\boldsymbol{r}_A(t)$. This assumes that the rotation of A is slower than the diffusion of B, which is true for protein-ligand systems. It is further assumed that B is a small molecule such that its orientation is not very relevant, as it will be the case in this implementation of the scheme. This simplifies Eq. (2.2) into a simple diffusion in $\boldsymbol{r}_B$ only

$$\partial_t \mathbf{r}_B(t) = \boldsymbol{\xi}(t), \tag{2.3}$$

with the components of $\boldsymbol{\xi}(t)$ again satisfying $\mathbb{E}[\xi_i(t)] = 0$ and $\mathbb{E}[\xi_i(t)\xi_{,j}(s)] = 2(D_A + D_B)\delta_{ij}\delta(t-s)$.

2.2.4. MSM/RD coupled dynamics

The present coupled model only considers interactions between up to two molecules. This is a frequent assumption in PBRD [67, 68, 72, 74] but may be restrictive from a molecular standpoint. It is assumed that simultaneous reactions between three or more molecules such as $\mathrm{A} + \mathrm{B} + \mathrm{C} \rightarrow \mathrm{D}$ can always be broken down into $\mathrm{A} + \mathrm{B} \rightarrow \mathrm{AB}$; $\mathrm{AB} + \mathrm{C} \rightarrow \mathrm{D}$ or other bimolecular pathways, and therefore the focus lies on MSM/RD involving two molecules. As the dynamics in the I and O region are given in terms of states and coordinates respectively, one needs to recognize that $\mathbf{x}$ and $\mathbf{y}$ in the transfer density $p(\mathbf{y} \rightarrow \mathbf{x}; \tau)$ can be either coordinates $\mathbf{c}$ (center-of-mass position and perhaps orientation of the molecule) or states s (metastable regions in the coordinate space). In order to implement the coupling, the following two quantities are defined:

- $p_{\text{entry}}\left(\mathbf{c}_t \rightarrow \mathbf{x}_{t+\Delta}; \Delta\right)$, transfer probability of starting in coordinates $\mathbf{c}_t$ just inside the MSM domain ($r_{AB}(\mathbf{c}_t) < R$) conditioned on hitting only one state $\mathbf{x}_{t+\Delta} = s_{t+\Delta}$ in the MSM domain (transition event) *or* on exiting once the MSM domain $\mathbf{x}_{t+\Delta} = \mathbf{c}_{t+\Delta}$ (return event)

- $p_{\text{exit}}\left(s_t \rightarrow \mathbf{x}_{\text{t}+\Delta}; \Delta\right)$, transfer probability of starting in state s_t conditioned on exiting once the MSM domain $\mathbf{x}_{t+\Delta} = \mathbf{c}_{t+\Delta}$ (exit event) *or* hitting once any other state ($\mathbf{x}_{t+\Delta} = s_{t+\Delta}$).

Once these transfer probabilities are known, the basic MSM/RD algorithm can be introduced where τ_{RD} and τ_{MSM} correspond to the diffusion and MSM time-step, respectively. Algorithm 1 shows the details of the implementation.

There are additional issues in specific scheme implementations, such as estimating the unknown conditional transfer probabilities, and choosing the MSM discretization and R such that the overall discretization error is small, among others. These issues are non-trivial and could potentially be tackled with different approaches. To asses the accuracy of the given approach, relevant macroscopic observable are compared between MSM/RD simulations and the ground truth. This comparison quantifies how well the ground truth can be approximated by the MSM/RD simulation. A possible implementation of the scheme is proposed in the following section.

2.3. An MSM/RD implementation for protein-ligand systems

This section introduces an implementation of the MSM/RD scheme for a special class of systems: the binding of a small ligand to a protein – a case that is e.g. relevant in the study of protein-drug binding kinetics [104]. While the theory described above is more general, implementations to more challenging systems such as protein-protein interaction are left open to be treated in future contributions. A step towards such systems has been taken recently in [76]. As a starting point, the macromolecule A is considered fixed at the origin with fixed orientation and the ligand B freely diffusing around it with an overall diffusion constant $D = D_A + D_B$. The macromolecule has several possible binding sites given by some interaction potential. In order to present the MSM/RD scheme in detail, three different types simulations of simulations are introduced:

(i) **Reference simulation** (ground truth, if available): MD simulation of B and its interaction with A in a large spherical domain with radius R_s. Unfortunately, reference simulations of realistic systems are in general not computationally feasible due

Algorithm 1: Basic MSM/RD algorithm

```
input  : Initial mode (PBRD or MSM)
         initial condition (coordinates $\mathbf{c}_0$ or state $s_0$, respectively)
         $t \leftarrow 0$
while $t \leqslant t_{\text{final}}$ :
    if in PBRD mode :
        Propagate $\mathbf{c}_t \to \mathbf{c}_{t+\tau_{\text{PBRD}}}$ by diffusion
        Update time $t \leftarrow t + \tau_{\text{PBRD}}$
        if $r_{AB}(\mathbf{c}_t) < R$ :                                  // enter MSM domain
            Sample next event $(\mathbf{x}_{t+\Delta},\Delta)$ from $p_{\text{entry}}(\mathbf{c}_t \to \mathbf{x}_{t+\Delta};\, \Delta)$
            if transition event :
                Map to state $s_{t+\Delta} \leftarrow \mathbf{x}_{t+\Delta}$
                Update time $t \leftarrow t + \Delta$
                Switch to MSM mode
            else:                                                        // return event
                Map to coordinates $\mathbf{c}_{t+\Delta} \leftarrow \mathbf{x}_{t+\Delta}$
                Update time $t \leftarrow t + \Delta$
    else:                                                                // in MSM mode
        if $s_t \neq s_{t-\tau_{\text{MSM}}}$ or previous mode $\neq$ MSM mode :
            Sample next event $(\mathbf{x}_{t+\Delta}, \Delta)$ from $p_{\text{exit}}\,(s_t \to \mathbf{x}_{\text{t}+\Delta};\, \Delta)$
            if exit event :
                Map to coordinates $\mathbf{c}_{t+\Delta} \leftarrow \mathbf{x}_{t+\Delta}$
                Update time $t \leftarrow t + \Delta$
                Switch to PBRD mode and break current loop iteration
            Propagate $s_t \to s_{t+\tau_{\text{MSM}}}$ using the MSM
            Update time $t \leftarrow t + \tau_{\text{MSM}}$
```

to the time and lengthscales of the simulation. Nonetheless, reference simulations of simple systems are used to verify the MSM/RD scheme and validate its use in more complex systems.

(ii) **Small-scale simulation** (MD simulation): analogous to the reference simulation with the difference that B is constrained to a small box with periodic boundary conditions, see Fig. 2.2 (a). As the potential is negligible outside this box, the main interaction dynamics are extracted from this simulation's data into an MSM. This simulation is used to parametrize the MSM/RD model.

(iii) **MSM/RD simulation** (hybrid model): couples the MSM for short-range interactions derived from the MD simulation *(ii)* with a diffusion scheme for the long-range, see Fig. 2.2 (c). The goal of the scheme is to approximate the ground truth dynamics given by the reference simulation *(i)*.

2.3.1. Estimation

Dynamic properties need to be extracted from the small-scale simulation *(ii)* in order to parametrize the MSM/RD scheme. These properties are the state-to-state dynamics and the coupling between the MSM and the PBRD domain. The state-to-state dynamics are estimated using an MSM, and the coupling is given in terms of entry and exit events from the MSM domain. As these events might happen on different timescales, it is favorable to be free from the fixed time-step that the MSM requires to be well equilibrated. Therefore, trajectory statistics are used for the entry and exit events.

MSM

As a first step for the construction of the Markov model and MSM/RD parametrization, a discrete representation of the underlying data needs to be found. This work uses the core MSM approach [101], which requires the definition of cores as metastable regions of phase space. Cores are given by spherical domains around the metastable regions in the MD simulation and can be found using a clustering algorithm. In the core MSM approach a discrete trajectory is constructed by assigning the last visited state-index to each point in the trajectory. Note the trajectory may leave the core of a given state and re-enter multiple times without transitioning to other states. Using this discretization technique, the discrete trajectories are truncated into three types of trajectories as shown in Fig. 2.2 (b): *i*) entry trajectories that start just inside the MSM domain and either leave the domain next or hit a core inside the domain. *ii*) transition trajectories that start in a state and hit another state as next event and *iii*) exit trajectories that start in a state and leave the MSM domain as next event. These trajectories are used to estimate the transfer densities and to parametrize the MSM/RD simulation.

The MSM for the short-range interactions is built using the full discrete trajectories and the exit trajectories, Fig. 2.2 (b). Following the methods from [9], a transition matrix $\mathbf{T}(\tau)$ is estimated, where the entries are the transition probabilities $\mathbf{T}_{ij}$ from state i to j.

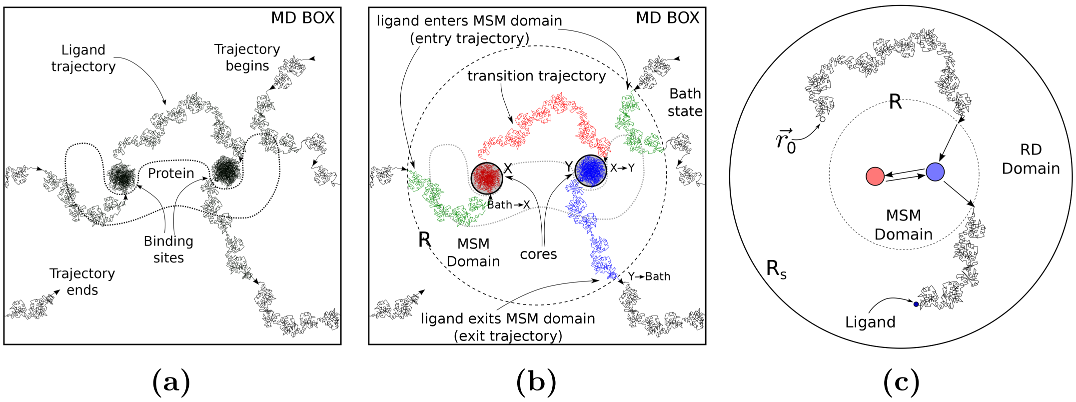

Figure 2.2.: Illustrations of a MD trajectory, its classification to extract the relevant dynamics and the MSM/RD scheme. **(a):** Illustration of a trajectory of a ligand in a MD simulation within a box with periodic boundaries (small-scale simulation). Note that there are two metastable regions, e.g. binding sites on a protein, where the ligand stays for a longer time. **(b):** Illustration of truncation and classification of the trajectory. The MSM domain is chosen so the interaction potential is effectively zero outside this region (bath state); the cores X and Y are chosen to represent the metastable regions in phase space. The truncated trajectories are classified into entry trajectories (green), transition trajectories (red) and exit trajectories (blue), which are used for the coupling in the MSM/RD scheme. In order to obtain the MSM for the MSM/RD scheme, the system is also classified into three states, the bath state and the two cores X and Y; it is also shown when the transition between these states occur along a trajectory. **(c):** Representation of the MSM/RD scheme. The full trajectories from the MD-simulation are used to derive an MSM to model the dynamics in the MSM domain. The entry and exit trajectories from the MD-simulation are used to couple the Brownian dynamics in the diffusion domain with the dynamics in the MSM domain. Reprinted from *The Journal of Chemical Physics* "MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations", Dibak et al., 2018, with the permission of AIP Publishing.

Using the discrete trajectories, count matrices $\mathbf{C}^{\text{full}}_{ij}(\tau)$ from the complete data set and $\mathbf{C}^{\text{exit}}_{ij}(\tau)$ from the exit trajectories are created, which count all the transitions from state i to j at a lag time τ observed in the respective datasets. As the coupling between the MSM and PBRD domain is handled separately, the MSM dynamics only accounts for transitions amongst the cores and therefore the counts arising from exit trajectories have to be subtracted

$$\mathbf{C}_{ij}(\tau) = \mathbf{C}^{\text{full}}_{ij}(\tau) - \delta_{ij}\mathbf{C}^{\text{exit}}_{ij}(\tau), \tag{2.4}$$

where δ_{ij} denotes the Kronecker delta

$$\delta_{ij} = \begin{cases} 1 & i = j \\ 0 & i \neq j. \end{cases} \tag{2.5}$$

The transition matrix is then obtained from the given counts $\mathbf{C}_{ij}$ using a maximum likelihood estimator. Note that the commonly used reversible estimator [20], that ensures detailed balance in the transition matrix, can not be used here, as detailed balance can not be assumed for these types of count matrices.

Entering the MSM domain

The protocol to enter the MSM domain from the PBRD domain is constructed with the entry trajectories as defined above. It consists of generating a list $L_{\text{entry}} = \{\mathbf{c}_{\text{entry}}, \mathbf{x}_{\text{end}}, \Delta\}$ of all start coordinates $\mathbf{c}_{\text{entry}}$(just inside the MSM domain) and endpoints $\mathbf{x}_{\text{end}}$ of entry trajectories and their corresponding times Δ. The endpoints may be either MSM states or coordinates in the PBRD domain, see Fig. 2.2 (b). The ensemble of trajectories in this list estimates the conditional transfer probability $p_{\text{entry}}\left(\mathbf{c}_t \to s_{t+\Delta};\, \Delta\right)$ (Section 2.2.4) for several times Δ. In the MSM/RD simulation samples are drawn from this list of entry points.

Exiting the MSM domain

For each state s of the MSM, all exit and transition trajectories are collected. Additionally their end-coordinate or -state along with their respective exit time is compiled in the lists $L_{\text{exit},s} = \{\mathbf{c}_{exit}, \Delta\}$ and $L_{\text{trans},s} = \{s_{\text{trans}}, \Delta\}$. The ensemble of trajectories in these list estimates the conditional transfer probability $p_{\text{exit}}\left(s_t \to \mathbf{c}_{\text{t}+\Delta};\, \Delta\right)$ (Section 2.2.4) for several times Δ. The probability of an exit event $P_{\text{exit},s}$ is simply estimated as the ratio of exiting trajectories over the total numbers of trajectories,

$$P_{\text{exit},s} = \frac{\#\text{ of trajectories in } L_{\text{exit},s}}{\#\text{ of trajectories in } L_{\text{exit},s} \text{ and } L_{\text{trans},s}}. \tag{2.6}$$

2.3.2. The MSM/RD algorithm for protein-ligand systems

Based on the estimated quantities defined in the Section 2.3, an implementation for protein-ligand systems is introduced based of the MSM/RD algorithm from Section 2.2.4. The algorithm is shown in Algorithm 2.

Algorithm 2: MSM/RD algorithm for protein-ligand systems

```
input : Initial mode (PBRD or MSM)
        initial condition (coordinates $\mathbf{c}_0$ or state $s_0$, respectively)
        t = 0
while $t \leqslant t_{final}$ :
    if in PBRD mode :
        Propagate $\mathbf{c}_t \to \mathbf{c}_{t+\tau_{\mathrm{PBRD}}}$ by diffusion
        $t \leftarrow t + \tau_{PBRD}$
        if $r_{AB}(\mathbf{c}_t) < R$ :                          // enter MSM domain
            Select trajectory from $L_{\mathrm{entry}} = \{\mathbf{c}_{\mathrm{entry}}, \mathbf{x}_{\mathrm{end}}, \Delta\}$
            with $\mathbf{c}_{\mathrm{entry}}$ closest to $\mathbf{c}_t$
            if $\mathbf{x}_{\mathrm{end}}$ is a state :                  // enter MSM mode
                Map to state $s_{t+\Delta} \leftarrow \mathbf{x}_{\mathrm{end}}$
                Update time $t \leftarrow t + \Delta$
                Switch to MSM mode
            if $\mathbf{x}_{\mathrm{end}}$ are coordinates :             // remain in PBRD mode
                Map to coordinates $\mathbf{c}_{t+\Delta} \leftarrow \mathbf{x}_{\mathrm{end}}$
                Update time $t \leftarrow t + \Delta$
    else:                                                    // in MSM mode
        if $s_t \neq s_{t-\tau_{\mathrm{MSM}}}$ or previous mode $\neq$ MSM mode :
            Sample exit event with $P_{\mathrm{exit},s}$
            if exit event :
                Uniformly select trajectory from $L_{\mathrm{exit},s} = \{\mathbf{c}_{\mathrm{exit}}, \Delta\}$
                Map to coordinates $\mathbf{c}_{t+\Delta} \leftarrow \mathbf{c}_{\mathrm{exit}}$
                Update time $t \leftarrow t + \Delta$
                Switch to PBRD mode and break current loop iteration
        Propagate $s_t \to s_{t+\tau_{\mathrm{MSM}}}$ using $\mathbf{T}(\tau_{\mathrm{MSM}})$
        Update time $t \leftarrow t + \tau_{\mathrm{MSM}}$
```

The diffusion in the PBRD domain is done using a Euler-Maruyama discretization of Eq. (2.3) [105]. A higher efficiency when simulating the diffusion step can be achieved with event-based algorithms, like FPKMC or eGFRD [65, 68, 73, 78] for systems with low particle concentrations. In order to optimize the efficiency of the algorithm, the entry points of entry trajectories are classified into equal area bins on the sphere. This allows the algorithm to find the closest trajectory to a given entry point more efficiently. The partition of the sphere was done following Ref. [106].

2.3.3. Verification of the MSM/RD scheme

In order to verify the MSM/RD scheme, systems are used for which it is feasible to perform a reference simulation. The internal dynamics is verified by comparing the first passage times (FPTs) distributions and mean first passage times (MFPTs) for each pair of metastable states within region I between the MSM/RD and reference simulations. The ground truth MFPTs are estimated by computing the FPTs $\mathbf{t}_{ij}^{\text{ref}}$, where the initial conditions are chosen as the minima $\mathbf{M}_i$ and the system is propagated following the reference simulation until hitting state j (conditioned on not leaving the MSM domain). For the MSM/RD scheme, the FPTs $\mathbf{t}_{i,j}^{\text{MSM}}$ are computed by placing the particle in state i and propagating the system following the MSM/RD scheme until state j is hit. If the particle exits the MSM domain before reaching state j, the trajectory is not taken into account. With a sufficiently large number of generated sampled, the distributions of FPTs can be estimated by histograms. The MFPTs are estimated as $\boldsymbol{\tau}_{ij}^{\text{ref}}=\overline{\mathbf{t}_{ij}^{\text{ref}}}$ and $\boldsymbol{\tau}_{ij}^{\text{MSM}}=\overline{\mathbf{t}_{ij}^{\text{MSM}}}$, respectively. The MFPT relative error between the MSM/RD and the reference simulations is estimated as

$$(\mathbf{E}_{\text{rel}})_{\text{ij}} = \frac{\boldsymbol{\tau}_{ij}^{\text{ref}} - \boldsymbol{\tau}_{ij}^{\text{MSM}}}{\boldsymbol{\tau}_{ij}^{\text{ref}}}. \tag{2.7}$$

Furthermore, in order to verify the coupling between the PBRD and MSM domain, the unbinding rate k_{off}, the binding rate k_{on} and equilibrium constant $K_{\text{eq.}}$ are compared between the ground truth and MSM/RD simulations. The binding rate and equilibrium constant are calculated for different particle concentrations c by fixing the radius R_s of the simulation domain such that $c = 1/V_{\text{PBRD}}$,with V_{PBRD} the volume of the PBRD domain.

2.4. Experiments

In this section, the MSM/RD scheme from Section 2.3 is applied to two systems. The first is a simple model of a ligand diffusing in a potential landscape, which is used to verify that the MSM/RD scheme reproduces the correct dynamics. The second corresponds to a more realistic MD system, where the binding of carbon monoxide to myoglobin is studied.

2.4.1. Ligand diffusion in potential landscape

The MSM/RD scheme is implemented in a simple model, where the reference simulation is available. The model consists of a ligand B under over-damped Langevin dynamics in a three-dimensional potential landscape

$$\partial_t \mathbf{x}(t) = -\frac{1}{\gamma}\nabla U(\mathbf{x}) + \boldsymbol{\xi}(t), \tag{2.8}$$

with U the interaction potential with some macromolecule A fixed at the origin, γ the damping coefficient, and each component of the noise satisfies $\mathbb{E}[\xi_i(t)] = 0$ and $\mathbb{E}[\xi_i(t)\xi_j(s)] = 2D\delta_{ij}\delta(t-s)$ with $D = k_\mathrm{B}T/\gamma$ the diffusion coefficient. A trajectory density plot of the potential landscape chosen is shown in Fig. 2.3 (a), and it consists of nine Gaussians with different depths and widths

$$U(\mathbf{r}) = -\sum_{i=1}^{9} s_i \mathcal{N}(\mathbf{r}; \boldsymbol{M}_i, \boldsymbol{\Sigma}_i), \tag{2.9}$$

where $\mathcal{N}(\mathbf{r}; \boldsymbol{M}_i, \boldsymbol{\Sigma}_i)$ denotes a Gaussian centered at minimum $\boldsymbol{M}_i$ with covariance matrix $\boldsymbol{\Sigma}_i$, s_i denotes a scale factor. The small-scale simulation consists of Euler-Maruyama numerical realizations of Eq. (2.8) under this potential constrained to a box with an edge length of 6 nm with periodic boundary conditions. The reference simulation is analogous to the small-scale simulation with the difference that it uses a larger spherical domain with reflective boundary conditions at a range of radii corresponding to simulations at different ligand concentrations.

Parametrization of the MSM/RD scheme

A radius of $R = 2.5\,\mathrm{nm}$ is selected for the MSM domain (I region) since outside this domain the potential [Eq. (2.9)] is essentially zero. For the estimation of the MSM, 120 small-scale simulation trajectories are generated. Each of these trajectories is run for 10^7 steps with a time-step of $\Delta t = 10^{-4}\,\mathrm{ns}$ and ever tenth step is recorded. This results in a total simulation time of $t = 1.2 \times 10^5\,\mathrm{ns}$.

The cores are defined as spheres with radius 0.2 nm around the minima $\boldsymbol{M}_i$, and the count matrix of transition between cores is generated from the trajectories following Eq. (2.4). A maximum likelihood estimator (implemented in PyEMMA [93]) is then applied to the count matrix to yield the MSM. From these trajectories the lists L_entry, $L_{\mathrm{exit},s}$, $L_{\mathrm{trans},s}$ and $P_{\mathrm{exit},s}$ are also generated which were introduced in Section 2.3.1.

The timescales of the eigenmodes are estimated for different MSM lag times to test how well the underlying process is estimated by the MSM. The timescales have small variations for different lag times [Fig. 2.3 (d)], which means the system can be considered Markovian for all lag times. However, the lag time has to be chosen with care, as a too large lagtime would mean that relevant fast timescales are not properly resolved which would result in significant errors. For all further analyses, a lag time of $\tau_\mathrm{MSM} = 500\Delta t = 0.05\,\mathrm{ns}$ is considered to be an optimal compromise.

Comparison of dynamic properties

The binding rate is computed by first estimating the first passage time from a uniformly sampled location closed to the simulation boundary $r = R_S - \delta$ to any MSM state. The distance to the boundary is chosen to be $\delta = 0.05\,\mathrm{nm}$ and 10^4 simulations are used to obtain and estimate of the average of the $\mathrm{MFPT_{on}}$, from which the binding rate is calculated as $k^*_\mathrm{on} = 1/\mathrm{MFPT_{on}}$. This procedure is performed for both the MSM/RD

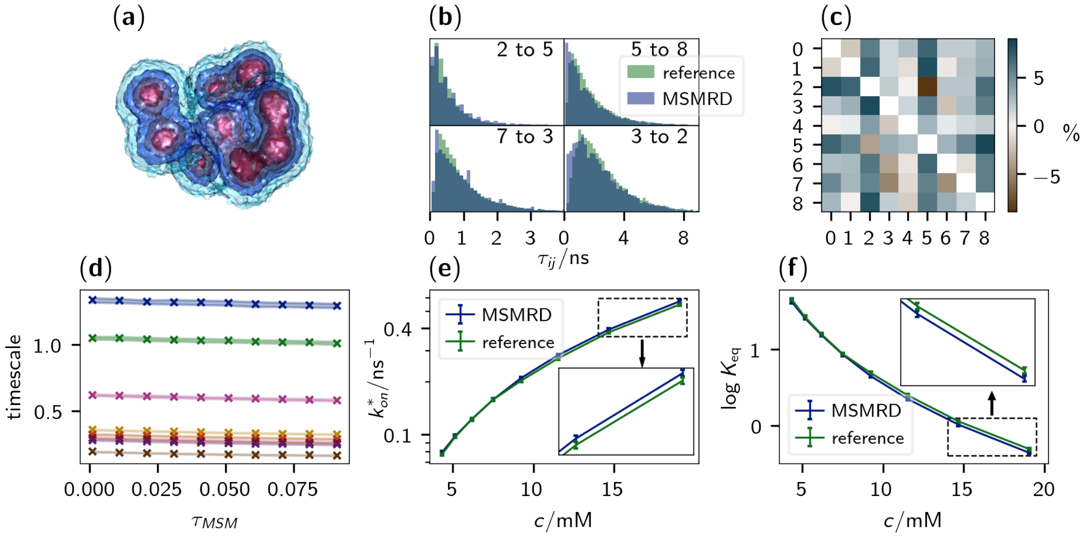

Figure 2.3.: Visualization and verification of results for the model of ligand diffusing in a potential landscape. **(a):** Density plot of the position of the ligand in the three dimensional potential. Red indicates regions of higher density while blue indicates regions of lower density. **(b):** Comparison of first passage times distribution histograms for the transitions with the highest error in **(c)**. The left pane corresponds to transitions with negative relative error, and the right pane to transitions with positive relative error. **(c):** Relative error of MFPTs conditioned on not leaving the MSM domain between the MSM/RD and the reference simulation. **(d):** Implied timescales of the MSM. The shaded area represents the standard deviation of the bootstrapping sample. Well converged timescales are observed for all considered lag times. **(e):** The rate k_{on} as function of the concentration of the system for the MSM/RD and reference simulations. **(f):** Same plot as **(e)** but for the logarithm of the equilibrium constant $\log(K_{eq})$. The error bars in **(e)** and **(f)** represent the 95 % confidence interval using a bootstrapping approach. Reprinted from *The Journal of Chemical Physics* “MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations”, Dibak et al., 2018, with the permission of AIP Publishing.

and the reference simulation, and excellent agreement is observed between the two, Fig. 2.3 (e).

For the unbinding rate, the inverse process is investigated by starting in an MSM state and propagating the dynamics until crossing a boundary defined by a sphere with radius $2.7\,\mathrm{nm} > R$. The reference simulation yields a value of $0.402^{0.404}_{0.400}\,\mathrm{ns}^{-1}$ (sub- and superscript indicate lower and upper bound of the 95% percentile) and the MSM/RD simulation a value of $k_{\mathrm{off}} = 0.400^{0.402}_{0.398}\,\mathrm{ns}^{-1}$. Furthermore, the logarithm of the equilibrium constant $\log(K_{\mathrm{eq}}) = \log(k_{\mathrm{off}}/k^{*}_{\mathrm{on}})$ is computed for both models and for the chosen values of concentrations, resulting in accurate reproduction of the reference values by the MSM/RD scheme, Fig. 2.3 (f). These experiments verify that the coupling between the MSM domain and the RD domain works consistently in the MSM/RD simulation scheme.

As a next step it is to be demonstrated, that also the dynamics between the states inside the MSM are reproduced to a high accuracy. This can be shown by comparing MFPTs between all pairs of states conditioned on not leaving the MSM domain. In the reference simulation this is done by placing the particle at position $\boldsymbol{M}_i$ and propagating the system until state j is reached. If the particle leaves the MSM domain before reaching state j, this trajectory is discarded. For the MSM/RD simulation, the simulation is simply started in state i and propagated until state j is hit, while discarding trajectories that leave the MSM domain. This procedure is repeated until 10^4 successful trajectories are found for both simulations, which are averaged to obtain the MFPTs. The relative errors are calculated with Eq. (2.7); all relative errors are below 9 % [Fig. 2.3 (c)]. It is further observed that negative errors arise for state pairs that are close together and thus have short passage times. For these transitions, the MFPT in the MSM/RD simulation is typically overestimated, as short processes are truncated in the MSM estimation. Moreover, it is observed that the highest positive errors arise for transitions which are far apart. These are the hardest to sample since for these transitions there is a very high number of possible long and non-direct transition trajectories, which are less likely to be observed. Fig. 2.3 (b) gives a closer look at the four transitions with the highest relative error by directly comparing the histograms of their FPTs. Even though these transitions have the highest errors, it is observed that the distributions match well. Therefore, it is verified that the MSM/RD scheme also describes the internal dynamics accurately.

2.4.2. Binding of CO to myoglobin

Myoglobin is a globular protein which is responsible for the transport of oxygen in muscle tissue. As a competitor, carbon monoxide (CO) can bind to it, thus disabling its function to transport oxygen. This system is an interesting application of the MSM/RD scheme. The binding process of CO to myoglobin has recently been studied in MD simulations by de Sancho et al. [54], whose data is used to parametrize the MSM/RD scheme. The dataset consist of MD trajectories of 20 CO molecules and one myoglobin protein for a total simulation time of 500 ns. The MD simulation is confined to a periodic box with edge length of 5 nm. Despite the fact that only one CO molecule can reside in the binding pocket, the error of treating 20 CO molecules as being statistically independent is small within statistical uncertainty as described in [54]. Therefore, the trajectories of the 20

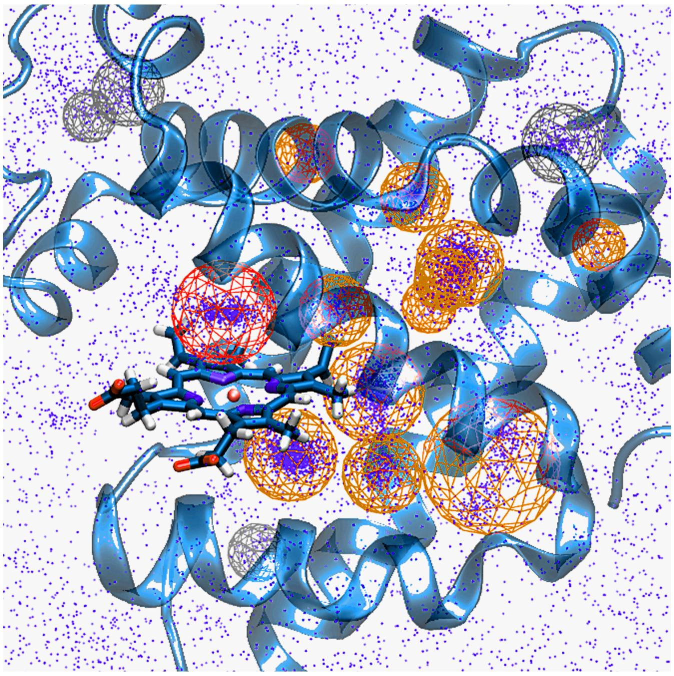

Figure 2.4.: Definition of the cores (wire frame spheres) within the myoglobin molecule. The red sphere indicates the bound state and orange spheres the metastable states. These are resolved in the MSM. The gray spheres correspond to the states that were not in the connected set and therefore discarded. The blue dots are positions of the CO molecules for every 50th frame in the vicinity of the protein. Reprinted from *The Journal of Chemical Physics* “MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations”, Dibak et al., 2018, with the permission of AIP Publishing.

CO molecules are treated independently, effectively increasing the total simulation time to 10 µs.

Parametrization of MSM/RD scheme

In order to parametrize the scheme, all frames are first aligned using the C_α atoms of the myoglobin as reference. The cores are identified using the *density-based spatial clustering of applications with noise* algorithm (DBSCAN) [107] on the aligned trajectories, which finds a total of 16 metastable regions/cores. The positions and size of the cores are shown in Fig. 2.4, where it can be observed that the algorithm correctly identifies regions of high ligand density, including the myoglobin bound state which is indicated in red. The radius of the spherical cores is chosen such that 80 % of the datapoints that were assigned to the respective state are inside the core.

Four states are discarded as they are not part of the largest connected set of the MSM.

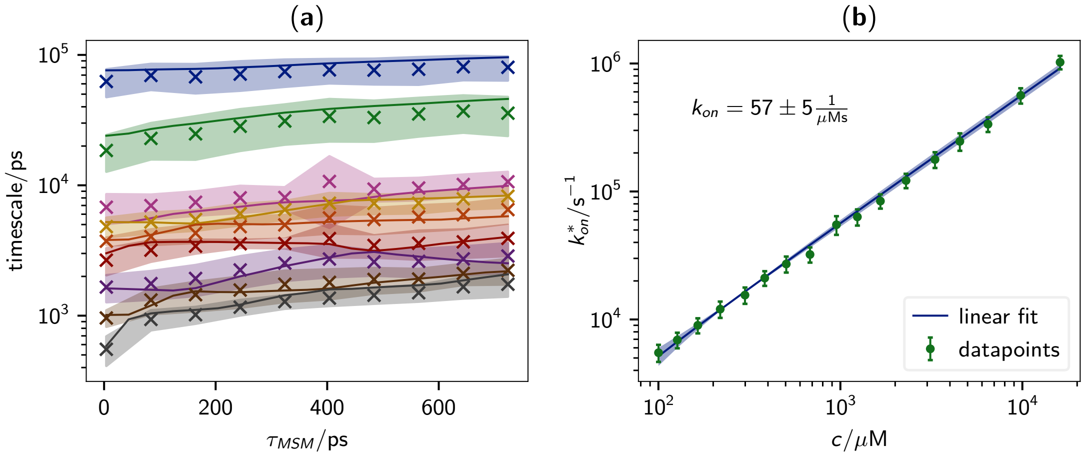

Figure 2.5.: Results of the CO-myoglobin system. **(a):** Implied timescales of the dynamics of the CO myoglobin system. The datapoints and shaded area denote the sample mean and standard deviation of the bootstrapping samples over the trajectories: from the 20 given trajectories 20 are sampled with replacement. Over this sample the discretization process is run which returns a sample of timescales. The trajectory-samples which are not ergodic or do not lead to a connected count matrix are considered invalid and discarded. Solid lines are found using the full dataset. **(b):** Reaction rate as estimated from multiple simulations at different concentrations. Reprinted from *The Journal of Chemical Physics* "MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations", Dibak et al., 2018, with the permission of AIP Publishing.

As the simulation box had been set up to just contain the protein and a 1 nm solvent layer, the largest MSM domain that still fits inside the box, $R = 2.5\,\mathrm{nm}$, is chosen for the border between MSM and PBRD region. Analogous to the previous example, the methods from Section 2.3 are used to estimate an MSM for the close-range dynamics and generate L_{entry}, $L_{\mathrm{exit},s}$, $L_{\mathrm{trans},s}$ and $P_{\mathrm{exit},s}$ to couple the dynamics in the two domains.

The implied timescales of the MSM are computed and a lag time of 150 ps is selected where timescales are sufficiently converged, Fig. 2.5 (a). The diffusion constant is computed using the mean squared displacement (MSD) of the parts of the CO trajectories that are far from the protein, with $D = \Delta\mathrm{MSD}(t)/6\Delta t$. Using a linear regression, a diffusion constant of $D_{\mathrm{CO}} = 2.5\,\mathrm{nm}^2\,\mathrm{ns}^{-1}$ is found, which is comparable to the experimental value, which is in the range of $D_{\mathrm{CO}} = 2.03\,\mathrm{nm}^2\,\mathrm{ns}^{-1}$ (at 20 °C) to $D_{\mathrm{CO}} = 2.43\,\mathrm{nm}^2\,\mathrm{ns}^{-1}$ (at 30 °C) [108].

Comparison of dynamic properties

As in the previous example, binding rates are computed by sampling positions uniformly in the PBRD domain and simulating the MSM/RD model until it reaches the bound state. For each concentration, 200 trajectories are run to estimate the binding rate k^*_{on}. These rates are plotted against the concentration and shown in Fig. 2.5 (b). The reaction rate $k_{\mathrm{on}} = 57^{62}_{52}\,\mu\mathrm{M}^{-1}\,\mathrm{s}^{-1}$ is obtained as the slope of the linear fit. For the unbinding rate, simulations are started in the bound state and MFPTs are collected for leaving the MSM domain, resulting in a rate of $k_{\mathrm{off}} = 19.0^{19.2}_{18.8}\,\mu\mathrm{s}^{-1}$. The resulting equilibrium constant is found as $K_{\mathrm{eq}} = k_{\mathrm{on}}/k_{\mathrm{off}} = 3.0^{3.3}_{2.7}\,\mathrm{M}^{-1}$ and is similar to $3.6\,\mathrm{M}^{-1}$ found by de Sancho et al. [54], both of which are close to the experimental value of $2.2\,\mathrm{M}^{-1}$[109]. Table 2.1 shows a comparison of the resulting rates of the different methods. The binding and unbinding rate found by de Sancho et al. [54], although yielding a similar equilibrium constant, are both nearly an order of magnitude faster than the ones obtained with MSM/RD (see 2.1). The first indication that the present rates are an improved estimate is the fact that the kinetics (both the MSM relaxation timescales and k_{on}) are independent of the lag time, Fig. 2.5 (a).

To validate that the MSM/RD estimates of k_{off} and k_{on} have been estimated without significant bias, it must be shown that they are statistically consistent with the ground truth, which in this case is a sufficiently large and sufficiently long MD simulation. Here, k_{off} can be estimated directly by counting the frequency of ligand dissociation events from the binding pocket in the underlying MD simulations. Since there are not sufficient full dissociation pathways from the bound to the dissociated states in the MD data in order to make a statistically relevant comparison, a more precise estimate is obtained by computing the MFPT using an MSM directly constructed from the original MD data with the same discretization as used in the MSM/RD model. This resulted in a reference estimate of $23.4^{46.6}_{11.6}\,\mu\mathrm{s}^{-1}$ (95 % percentile computed with 1000 bootstrap samples), which is consistent with the MSM/RD estimate (Table 2.1).

Unfortunately, this method is not as accurate for the binding rate k_{on}, which is notoriously difficult to estimate from small MD simulation boxes, where the length of trajectory segments in which the ligand stays in the dissociated state without touching the protein or crossing the periodic boundary are short compared to lagtimes τ used in an MSM approach, resulting in biased estimates [56]. Therefore, another Myoglobin MD simulation is performed in an eightfold larger periodic box (edge length 10 nm) with the same CO concentration as in the small MD simulation (resulting in 160 CO molecules) for a total simulation time of 405 ns. For this data, a direct MSM estimate of the binding rate yields $74.7^{130.9}_{29.9}\,\mu\mathrm{M}^{-1}\,\mathrm{s}^{-1}$ (95 % percentile computed with 1000 bootstrap samples). As a result, the MSM/RD binding and dissociation rates are consistent with standard estimates computed directly from MD simulation, and the MSM/RD modeling error can be concluded to be statistically insignificant.

Given the consistency of the models, a further comparison is taken with results of experimental measurements, which is essentially a test of the MD model (e.g. force field, thermostat, integrator). These are yet a factor 4-5 slower than the estimates obtained from MSM/RD ($k_{\mathrm{on}} = 12\,\mu\mathrm{M}^{-1}\,\mathrm{s}^{-1}$ and $k_{\mathrm{off}} = 5.3\,\mu\mathrm{s}^{-1}$ found in [109]), confirming that

	MSM/RD	Reference (approx. ground truth)	MSM in [54]	Experiment [109]	Unit
k_{on}	$57.0_{52.0}^{62.0}$	$74.7_{27.9}^{130.9}$	647	12	$\mathrm{M}^{-1}\,\mu\mathrm{s}^{-1}$
k_{off}	$19.0_{18.8}^{19.2}$	$23.4_{11.6}^{46.6}$	179	5.3	$\mu\mathrm{s}^{-1}$
K_{eq}	$3.0_{2.7}^{3.3}$	$3.19_{2.6}^{3.8}$	3.6	2.2	M^{-1}

Table 2.1.: Rates and equilibrium constants for Myglobin-CO estimated from different methods. The reference values approximate the ground truth by conducting a standard MSM-based MFPT estimate from the MD simulation (for k_{on} a larger simulation box was used to allow for a generous definition of the dissociated state). Lower- and uppercase numbers show 95 % confidence intervals of respective numbers

the major part of the difference between the estimates in [54] and the experimental values could be removed by the fact that MSM/RD is a significantly more accurate model of the binding kinetics.

2.5. Discussion and outlook

Combining MSMs derived from MD simulations with PBRD simulations is a promising simulation scheme for reaching longer time- and length-scales. In this chapter, MSM/RD was introduced and developed, which proposes a coupling of the two methods based on the distance of the interacting molecules: The molecule's interaction in close vicinity is resolved by an MSM by stochastically traversing between its states; in the far distance regime, where particles no longer interact, the molecules are diffusing independently. The MSM extracts the characteristic features of the dynamics from several short MD simulations. This allows for the generation of new data with great accuracy at a much faster rate than the original MD simulation. This is a clear advantage to comparable approaches that rely on running explicit MD simulations when molecules are in vicinity [78, 103]. Specifically, a simulation algorithm for protein-ligand systems is proposed which allows to study the binding kinetics of such systems. It is observed that the proposed schemes yields more accurate binding rates than conventional MSM approaches as it decreases boundary effects and increases the lifetime of the dissociated state. The scheme can be, in principle, coupled to any PBRD scheme, like over-damped Langevin dynamics, Langevin dynamics, GFRD [73, 74], and FPKMC algorithm [68], which could yield additional efficiency and accuracy.

For verification, the MSM/RD scheme was implemented for a simple ligand diffusion model consisting of a particle diffusing in a potential landscape, which mimics the interaction between a ligand and a protein. It is observed that the combined MSM/RD scheme reproduces the correct dynamics of the binding/unbinding processes as compared to the reference simulation. Especially, it was able to produce an MSM that accurately captures the interaction dynamics with a relatively small amount of data. This hints towards the model being scalable to larger systems with computationally feasible amounts of MD simulation data.

As a biologically relevant application, the MSM/RD scheme was applied to the binding of CO to myoglobin. After successfully extracting a self-consistent MSM and the coupling dynamics from the MD data, it is observed that the equilibrium constant is consistent with previous experiments and computational results [54, 109]. It is further demonstrated that the MSM/RD estimates are consistent with the underlying MD simulations and, in particular, the estimated association rate is consistent with the association rate estimated from a reference MD simulation conducted in a large simulation box that was not used to parametrize the MSM/RD model. This is a significant improvement over Ref. [54], where tenfold higher rates were estimated.

In a recent development [76], the MSM/RD scheme has been extended to enable rototranslational diffusion of the particles, independent conformational changes, as well as interaction between multiple molecules. To this end, the theory has been extended to treat the multiscale scheme in the framework of hybrid switching diffusion. From this theory, new simulation algorithms were derived which show promising results for toy systems with anisotropic binding patches and self-assembly of multiple particles. With the addition of these features, biologically relevant scenarios can be simulated. If, for example, different conformations of the protein have different ligand association/dissociation rates and the conformational changes of the protein are rare events, then the conformational dynamics and the ligand-binding dynamics are non-trivially coupled at high ligand concentrations. Such coupled binding kinetics is e.g. observed in the trypsin-benzamidine system [53].

Another biological relevant example on a larger scale is the activation of the calcium sensor synaptotagmin in neuronal synapses [110]. Here, a locally very high calcium concentration is created by the opening of voltage-gated calcium channels as a response to an electric signal. Synaptotagmin then binds up to five calcium ions while going through different conformations. At the same time, the local calcium concentration is reduced as ions leave the vicinity due to diffusion. If synaptotagmin successfully binds enough calcium ions and transitions into an active conformation, it can catalyze the fission of neuronal vesicles, which transduces the signal to the postsynaptic side. Such scenarios could be simulated with MSM/RD, by resolving the channels, the synaptotagmin proteins, and the ions as individual particles. In this scenario, the binding/dissociation kinetics and conformational changes of synaptotagmin would be encoded in several MSMs, which describe the kinetics of the interacting molecules. This would allow studying the whole biological process in one single MSM/RD simulation.

This example illustrates that MSM/RD could be used to study a wide range of systems and future iterations of the method, including further details of the underlying system could allow for computational studies of biological processes that are beyond reach with other methods currently available.

3. Diffusion-influenced reaction rates in the presence of pair interactions

The results of this chapter have been published in:

> Manuel Dibak, Christoph Fröhner, Frank Noé, and Felix Höfling. "Diffusion-influenced reaction rates in the presence of pair interactions". In: *The Journal of Chemical Physics* 151.16 (2019), p. 164105. DOI: `10.1063/1.5124728`

Parts of text and figures have been adopted unchanged in this chapter. Reproduced from *The Journal of Chemical Physics* "Diffusion-influenced reaction rates in the presence of pair interactions", Dibak et al., 2019, with the permission AIP Publishing.
Manuel Dibak (MD) and Christoph Fröhner (CF) contributed equally to this work. In particular, the contributions of the authors were as follows: Felix Höfling, MD and CF conceived the project and laid out the theory. MD implemented the semi-analytical solution. CF ran and analyzed the iPRD simulations. MD and CF visualized the data. All contributors wrote the manuscript.

Summary

In a coarse-grained picture, the interaction between two molecules in solution can be described in terms of bimolecular reactions. Microscopically, bimolecular reactions can be seen as a two step process: the reactants firstly have to encounter by diffusion and in a second step, a reaction can take place. The kinetics of these processes depends, among other factors, on intermolecular forces, such as steric repulsion or electrostatic interactions. In this chapter, an extension to the Doi reaction model by interaction potentials is established. This is a key ingredient for interacting particle reaction-dynamics (iPRD) simulations. Specifically, for a bimolecular reaction, the model parameters of the microscopic model are related to the macroscopic reaction rate constants, bridging the gap between the micro- and macroscopic description. To this end, the corresponding reaction-diffusion equation is solved in the steady-state, yielding semi-analytical expressions for the macroscopic rate constant and the density profiles. Analytical expressions are provided for the limiting cases and a perturbative solution is found in the limit slow reactions. A computationally inexpensive numerical scheme is provided for the general case, allowing for probing the whole range from slow to fast reactions. It is observed that the resulting rate uniquely decomposes into encounter and formation rate, with the interaction potential influencing both subprocesses. As examples, the influence of a harmonic repulsion and the Lennard-Jones potential on these rates are discussed in

detail. The theoretical analysis is corroborated by extensive iPRD simulations of the systems under investigation and excellent agreement is observed.

3.1. Microscopic reaction model

Following the discussion in Section 1.4, this chapter considers a solution of substances A and B, that undergo the reaction

$$\mathrm{A} + \mathrm{B} \longrightarrow \mathrm{A}^* + \mathrm{B}, \tag{3.1}$$

for which the product A* of the reaction falls out of scope, such that it does not need to be considered further. The concentrations c_A and c_B of A and B molecules, respectively, are assumed to be both so dilute that interactions between molecules of the same type can safely be ignored. (Otherwise, the reaction kinetics would non-trivially depend on c_A and c_B and the reaction rate would not be a well-defined constant.) Further, the concentration of B molecules is assumed to be much smaller than that of A, $c_B \ll c_A$, i.e., A molecules are abundant relative to Bs and there is no competition for reactants between the B molecules. Equivalently, substance B is highly diluted, and the problem can be rephrased as that of a single B molecule surrounded by A molecules in a large, yet finite volume V. It is convenient to switch to the reference frame of the B molecule, and a spherical volume V of radius L is chosen; see Fig. 3.1 for an illustration. In a finite amount of time and for sufficiently large V, the B molecule absorbs only a negligible fraction of As so that a quasi-steady state can be assumed with the concentration c_A being constant at the boundary ∂V of the volume. This setup can entirely be studied by only considering the concentration field of species A, $c_A(\mathbf{r}, t)$, in the frame of reference of a single B particle. As this is the only concentration field of interest, the subscript A will be omitted throughout the rest of this chapter $c(\mathbf{r}, t) = c_A(\mathbf{r}, t)$.

The Doi model [30–33] is used as microscopic reaction model, in which A and B molecules diffuse in space with diffusion constants D_A and D_B, respectively, forming a reactive complex whenever an A is separated from a B by less than the reaction distance R. This reactive complex undergoes reaction (3.1) with a microscopic rate constant or *propensity* λ, thus effectively removing A molecules from the system with a frequency K. More precisely, given a reactive complex, reaction events are triggered by a Poisson clock with parameter λ. The throughput or velocity of reaction (3.1) is then given by

$$\frac{\mathrm{d}c_{A^*}}{\mathrm{d}t} = K c_B\,, \tag{3.2}$$

where c_{A^*} is the overall concentration of the reaction product A*.

Similarly to Debye's work [28], and as commonly done in iPRD simulations [112], the focus here is on situations where A and B molecules interact physically with each other according to an isotropic pair potential $U(\mathbf{r}) = U(|\mathbf{r}|)$; the vector $\mathbf{r}$ denotes the separation of an AB pair. In this setup, the reaction-diffusion equation [Eq. (1.21)] can be formulated in terms of the average concentration field $c(\mathbf{r}, t)$ of A molecules and the

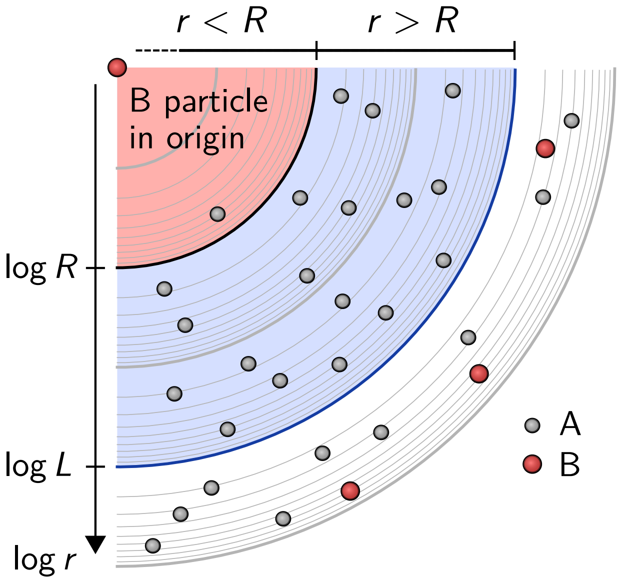

Figure 3.1.: System of reactive molecules. Molecules of species A diffuse in space and can react with B molecules if their distance r is smaller than the reaction radius R. If B particles are scarce, a reasonable assumption is that there is no competition between them and one can treat only one of them within a spherical domain of radius $L \gg R$. For the analytical treatment, $L \to \infty$, whereas for numerical methods and simulations L is finite. Reprinted from *The Journal of Chemical Physics* "Diffusion-influenced reaction rates in the presence of pair interactions", Dibak et al., 2019, with the permission of AIP Publishing.

corresponding flux (density) $\mathbf{j}(\mathbf{r},t)$. The reaction operator then takes the shape of a reaction propensity $a(\mathbf{r})$ yielding

$$\partial_t c(\mathbf{r},t) = -\nabla\cdot\mathbf{j}(\mathbf{r},t) - a(\mathbf{r})\,c(\mathbf{r},t)\,, \tag{3.3a}$$

$$\mathbf{j}(\mathbf{r},t) := -D\,\mathrm{e}^{-\beta U(\mathbf{r})}\nabla\Big[\mathrm{e}^{\beta U(\mathbf{r})}c(\mathbf{r},t)\Big]\,, \tag{3.3b}$$

with $a(\mathbf{r}) \geqslant 0$ and $D = D_A + D_B$ the relative diffusion constant of the particles; $\beta = 1/k_\mathrm{B}T$ denotes the inverse of the thermal energy scale. Within the Doi model, the propensity $a(\mathbf{r})$ is implemented in terms of the Heaviside step function,

$$a(\mathbf{r}) = \lambda\,\theta(R - |\mathbf{r}|), \tag{3.4}$$

such that the B molecule appears as a spherical reactive sink of radius R.

By isotropy of the setup, the steady flux $\mathbf{j}(\mathbf{r})$ of A molecules has only a radial component $j(r)$ that is a function only of the distance $r = |\mathbf{r}|$ to the B molecule. It determines the reaction frequency K through the surface integral

$$K = -\int_{|\mathbf{r}|=R} \mathbf{j}(\mathbf{r})\cdot\mathbf{n}\,\mathrm{d}\sigma = -4\pi R^2 j(R)\,, \tag{3.5}$$

with the surface normal $\mathbf{n}$ pointing outwards; the minus sign arises due to the fact that particles flow from the boundary to the sink at the origin, $j(r) < 0$. On the other hand, the law of mass action [Eq. (1.23)] yields the reaction rate equation

$$\frac{\mathrm{d}c_{A^*}}{\mathrm{d}t} = k c_A c_B\,, \tag{3.6}$$

in terms of the macroscopic association rate constant k. Comparing to Eq. (3.2), the latter is related to the microscopic frequency K by $k = K/c_A$, and the reaction rate constant follows as

$$k = \frac{4\pi R^2 |j(R)|}{c_A}\,. \tag{3.7}$$

The goal of the following sections is to calculate the flux profile $j(r)$ of the quasi-steady state and thus the macroscopic rate k, focussing on their dependences on the microscopic reaction parameters, λ and R, and on the pair potential $U(r)$ between A and B molecules. Note once again that there is no interaction amongst A molecules.

3.2. Solution strategy and classical limiting cases

In this section, the general solution strategy is worked out for the reaction–diffusion equations, Eq. (3.3), and analytical solution to important subproblems are obtained, which resemble a number of classical results. The stationary solutions $c(\mathbf{r})$ obeys $\partial_t c(\mathbf{r}) = 0$, and thus Eq. (3.3a) reduces to

$$\nabla \cdot \mathbf{j}(\mathbf{r}) = -a(\mathbf{r})\, c(\mathbf{r}). \tag{3.8}$$

According to the quasi-steady state assumption, $c(\mathbf{r})$ further satisfies the Dirichlet boundary condition

$$c(\mathbf{r}) = c_A\,, \quad \mathbf{r} \in \partial V\,. \tag{3.9}$$

Restricting to isotropic potentials, the symmetry can be utilized by switching to a single radial coordinate, $r = |\mathbf{r}|$, with the convention that the flux $j(r) = \mathbf{j}(\mathbf{r}) \cdot \mathbf{r}/r$ points outwards:

$$\frac{1}{r^2}\, \partial_r r^2 j(r) = -\lambda\, \theta(R - r)\, c(r) \tag{3.10}$$

with

$$j(r) = -D \mathrm{e}^{-\beta U(r)} \partial_r \Big[\mathrm{e}^{\beta U(r)} c(r)\Big]. \tag{3.11}$$

In this case and for an infinitely large volume V, Eq. (3.9) simplifies to $c(r \to \infty) = c_A$.

To complete the boundary value problem for $c(r)$, the behavior at the coordinate origin needs to be specified, which is not obvious due to the interaction potential. The total flux through a ball B_ε of radius ε centered at $\mathbf{r} = 0$ obeys:

$$\int_{\partial B_\varepsilon} \mathbf{j}(\mathbf{r}) \cdot \mathbf{n}\, \mathrm{d}\sigma = -\int_{B_\varepsilon} a(\mathbf{r})\, c(\mathbf{r})\, \mathrm{d}^3 r\,, \tag{3.12}$$

invoking Gauss' theorem and inserting Eq. (3.8). Continuity of the solution $c(\mathbf{r})$ together with the choice for $a(\mathbf{r})$ according to Eq. (3.4) yields $4\pi\varepsilon^2\, j(\varepsilon) \simeq -\lambda c(0) \cdot 4\pi\varepsilon^3/3$, and thus

$$j(0) = 0\,. \tag{3.13}$$

It implies a Robin boundary condition for the concentration profile,

$$\lim_{r \to 0} \left[\beta U'(r)\, c(r) + \partial_r c(r)\right] = 0, \tag{3.14}$$

which is satisfied by a Boltzmann distribution scaled by a constant factor:

$$c(r) \sim \exp\big(-\beta U(r)\big)\,, \quad r \to 0, \tag{3.15}$$

which captures the r-dependence asymptotically.

Note that the preceding derivation does not apply for potentials $U(r)$ that diverge as $r \to 0$. In this case, the current $\mathbf{j}(\mathbf{r})$ is not defined at the origin, $\mathbf{r} = 0$, and, strictly speaking, this point must be excluded from the integration domain B_ε, which forbids the application of Gauss' theorem. Yet, the extension of Eq. (3.15) to diverging potentials,

$U(r \to 0) = +\infty$, is motivated physically as it is improbable that any A molecule reaches the centre of the reaction volume: an upper bound on $c(r)$ is given by the equilibrium distribution, describing the non-reacting case. In particular, $c(r)$ is continuous in $r = 0$ and so is $\nabla \cdot \mathbf{j}(r)$ by Eq. (3.8), justifying the use of Gauss' theorem *a posteriori*.

The step-like reaction propensity in Eq. (3.10) suggests splitting the domain at the reaction boundary, $r = R$, and to find separate solutions $c_{\gtrless}$ and $j_{\gtrless}$ in both subdomains, $r \gtrless R$. By inspection of the r.h.s. of Eqs. (3.10) and (3.11), the flux $j(r)$ is finite and continuous at this interface, which implies that $c(r)$ is continuously differentiable at $r = R$. This provides the interface conditions

$$c_>(R) = c_<(R)\,, \tag{3.16}$$

$$j_>(R) = j_<(R) = -K/4\pi R^2, \tag{3.17}$$

making use of Eq. (3.5) in the last step. Matching the solutions of both subdomains will thus yield the sought-after reaction frequency K.

3.2.1. Outer solution

In the outer domain ($>$), where $R \leqslant r < \infty$, Eq. (3.10) reduces to an equation for the flux alone, $\partial_r r^2 j_>(r) = 0$. Integration from the lower boundary, Eq. (3.17), to some $r > R$ yields:

$$j_>(r) = -\frac{K}{4\pi r^2}\,, \tag{3.18}$$

with unknown rate K. The functional dependence on r is readily understood by the fact that, in the absence of reactions, the integral flux through spheres of radius r is constant (Gauss' theorem). In particular, the solution is compatible with the no-flux condition, $j_>(r \to \infty) = 0$, which is implied by the upper boundary, $c_>(r \to \infty) = c_A$, together with the vanishing force, $-\nabla U(r \to \infty) = 0$, and using Eq. (3.11).

A next step is to calculate the concentration profile $c_>(r)$ from Eqs. (3.9) and (3.11). Introducing

$$g(r) := \mathrm{e}^{\beta U(r)} r^{-2} \tag{3.19}$$

for brevity, one finds $(K/4\pi D)\, g(r) = \partial_r \left[\mathrm{e}^{\beta U(r)} c_>(r)\right]$, and after integration over $[r, \infty)$:

$$c_>(r) = \mathrm{e}^{-\beta U(r)} \left[c_A - \frac{K}{4\pi D} \int_r^\infty g(s)\mathrm{d}s\right]\,, \tag{3.20}$$

which is Debye's classical result [28]. If the interaction potential is not present ($U = 0$), this reduces to the familiar solution of the Dirichlet–Laplace problem:

$$c_>(r) = c_A - \frac{K}{4\pi D}\frac{1}{r}\,. \tag{3.21}$$

For diffusion-limited reactions, that is when product formation is fast and $k_\mathrm{f} \gg k_\mathrm{e}$ in Eq. (1.20), particles almost surely react on the surface of the reaction volume and the concentration inside vanishes: $c_<(r) = 0$ for $r \leqslant R$. Then by continuity of $c(r)$ at the interface of the subdomains, Eq. (3.20) is amended by $c_>(R) = 0$ and can be solved for K. This yields the Debye reaction rate constant $k = K/c_A$, which it identified as the encounter rate k_e in the presence of a pair potential:

$$k_\mathrm{e} = 4\pi D \Big/ \int_R^\infty g(s)\mathrm{d}s\,. \tag{3.22}$$

The corresponding concentration profile is given by Eq. (3.20) and reads

$$c_>(r) = c_A \mathrm{e}^{-\beta U(r)} \int_R^r g(s)\,\mathrm{d}s \Big/ \int_R^\infty g(s)\,\mathrm{d}s. \tag{3.23}$$

In particular, $c_>(r)$ is independent of the diffusion constant D. For $U(r) = 0$, these results recover Smoluchowski's rate constant [14] $k = 4\pi DR$ and the profile $c_>(r) = c_A(1 - R/r)$.

3.2.2. Inner solution without potential

In the absence of an interaction potential, Eqs. (3.10) and (3.11) simplify drastically and the concentration inside $c_<(r)$ the reaction volume, $0 \leqslant r \leqslant R$, obeys the Helmholtz equation

$$\left(\partial_r^2 + \frac{2}{r}\partial_r - \kappa^2\right) c_<(r) = 0 \tag{3.24}$$

with the inverse length $\kappa := \sqrt{\lambda/D}$, describing the penetration depth into the reactive domain. The flux takes the form $j_<(r) = -D\partial_r c_<(r)$, which turns the boundary conditions for the flux, Eqs. (3.13) and (3.17), into von Neumann conditions for the concentration, $c_<'(0) = 0$ and $c_<'(R) = K/4\pi DR^2$. Equation (3.24) is equivalent to $(\partial_r^2 - \kappa^2)[rc_<(r)] = 0$, and the boundary value problem is solved by [63]

$$c_<(r) = \eta\,\frac{\sinh(\kappa r)}{\kappa r} \tag{3.25}$$

with the constant η fixed by the upper boundary; in particular, η is proportional to the reaction frequency K. Matching inner and outer solutions for $c(r)$, Eqs. (3.21) and (3.25), at the interface, $r = R$, leads to $\eta = c_A/\cosh(\kappa R)$, and Doi's result for the reaction rate constant [31, 63] follows:

$$k = 4\pi DR\left[1 - \frac{\tanh(\kappa R)}{\kappa R}\right]. \tag{3.26}$$

The solution naturally decomposes as in Eq. (1.20) into Smoluchowski's encounter rate $k_\mathrm{e} = 4\pi DR$, see Eq. (3.22), and a formation rate

$$k_\mathrm{f} = 4\pi DR[\kappa R \coth(\kappa R) - 1], \tag{3.27}$$

with $\coth(x) = 1/\tanh(x)$. In the fast-diffusion limit, $\kappa R \ll 1$, i.e., when the reaction propensity λ is low, the formation rate $k_\mathrm{f} \simeq (4\pi/3)R^3\lambda$ is simply the product of the reaction volume $V_R = (4\pi/3)R^3$ and the propensity, reflecting well-mixed conditions inside the reaction volume ($c_<(r) = \mathit{const}$). For fast reactions, $\kappa R \gg 1$, one obtains $k_\mathrm{f} \simeq 4\pi R^2\kappa^{-1}\lambda$, which is interpreted as reactions being restricted to a volume $4\pi R^2\kappa^{-1}$, that is a thin shell of radius R and width κ^{-1}.

3.3. Reaction rates and spatial distributions in the presence of an interaction potential

For the general solution to the reaction–diffusion problem, Eqs. (3.10) and (3.11), in the presence of an interaction potential, it remains to find a solution inside the reaction radius (inner domain) and to match it with Eq. (3.20). As boundary condition $j_<(0) = 0$ is used, Eq. (3.13), and the equations for the current $j_<(r)$ are solved first.

3.3.1. Constant potential inside the reaction volume

For the special case that the interaction potential is constant within the reaction volume, i.e., $U(r) = U(R)$ for $r \leqslant R$, the solution is analytically accessible. This may be useful in modelling reactions in electrolytes while neglecting excluded volume effects and is discussed as a preliminary to the general case. The inner solution is then equals the non-interacting case, Eq. (3.25), and can be matched with Eq. (3.20) to find the reaction rate constant

$$k = 4\pi D \left(\frac{R\, g(R)}{\kappa R \coth(\kappa R) - 1} + \int_R^\infty g(r)\, \mathrm{d}r \right)^{-1} . \tag{3.28}$$

In particular, the encounter rate k_e is equal to Debye's result, Eq. (3.22), whereas the formation rate is suppressed by a factor $R^2 g(R) = \mathrm{e}^{\beta U(R)}$ relative to the non-interacting value, Eq. (3.27), and the total rate is the harmonic mean of both, Eq. (1.20).

3.3.2. Solution for arbitrary potentials

Proceeding along the lines of the potential-free case, Section 3.2.2, a solution to Eqs. (3.10) and (3.11) is found inside the reaction volume, $0 \leqslant r \leqslant R$, subject to the boundary conditions Eqs. (3.13) and (3.17). Applying the differential operator $\mathrm{e}^{-\beta U(r)}\partial_r \mathrm{e}^{\beta U(r)}$ on both sides of Eq. (3.10) and identifying the flux on the right hand side, one finds the following Dirichlet problem for the dimensionless function $\psi(r) := -4\pi r^2 j_<(r)/K$:

$$\psi''(r) + \left(\beta U'(r) - \frac{2}{r}\right)\psi'(r) - \kappa^2\psi(r) = 0\,, \tag{3.29a}$$

$$\psi(0) = 0\,, \quad \text{and} \quad \psi(R) = 1\,. \tag{3.29b}$$

In the absence of an explicit solution, the method of finite differences [113] is used to compute, in particular, the derivative on the reaction boundary, $\psi'(R)$. The latter determines the concentration on the boundary via Eq. (3.10):

$$c_<(R) = \psi'(R)K/4\pi R^2\lambda\,. \tag{3.30}$$

The reaction frequency K is obtained by matching inner and outer solutions for the concentration, Eq. (3.16). Employing the numerical value for $\psi'(R)$ and the previous result, Eq. (3.20), to find

$$\frac{K}{4\pi R^2\lambda}\,\psi'(R) = \mathrm{e}^{-\beta U(R)}\left[c_A - \frac{K}{4\pi D}\int_R^\infty g(s)\,\mathrm{d}s\right]. \tag{3.31}$$

Solving for $K = k/c_A$, yields an exact, closed expression for the macroscopic rate constant k, which is one of the main results:

$$k = 4\pi D\left[\int_R^\infty g(s)\mathrm{d}s + \frac{g(R)\,\psi'(R)}{\kappa^2}\right]^{-1}; \tag{3.32}$$

the pair potential enters through the function $g(r) := \mathrm{e}^{\beta U(r)}r^{-2}$. The result naturally displays the decomposition of Eq. (1.20), and the formation rate is identified as

$$k_\mathrm{f} = \frac{4\pi\lambda}{g(R)\,\psi'(R)}\,, \tag{3.33}$$

which appears to be proportional to the reaction propensity λ; in fact, the value of $\psi'(R)$, as given by Eqs. (3.29), indirectly depends on λ as well. Note that the diffusion-limited encounter rate k_e is the same as for the Debye problem, see Eq. (3.22), and the classical result, $k = k_\mathrm{e}$, is recovered in the limit of instantaneous reactions, $\lambda \to \infty$, i.e., for vanishing k_f^{-1}.

An alternative expression for the formation rate k_f in terms of the concentration $c(R)$ is obtained by substituting $\psi'(R)$ using Eq. (3.30) and $K = kc_A$, which yields $k_\mathrm{f} = kc_A\mathrm{e}^{-\beta U(R)}/c(R)$. Employing the decomposition of the total rate k [Eq. (1.20)] and solving for k_f, one finds

$$k_\mathrm{f} = k_\mathrm{e}\left[\frac{c_A\mathrm{e}^{-\beta U(R)}}{c(R)} - 1\right]. \tag{3.34}$$

Interestingly, the formation rate is fully specified by the encounter rate k_e and the concentration at the reaction boundary relative to its equilibrium value. However, the computation of $c(R)$ requires the full solution of the reaction–diffusion problem.

The concentration profile $c(r)$ follows from integration of Eq. (3.11) in terms of $\psi(r)$ and using continuity, Eq. (3.16), to eliminate $c_<(R)$ to find

$$c(r) = c_A\mathrm{e}^{-\beta U(r)}\left[1 - \frac{k}{4\pi D}\int_r^\infty g(s)\psi(s)\mathrm{d}s\right], \tag{3.35}$$

with the convention $\psi(r) = 1$ for $r > R$. Alternatively, the density profile can also be found by Eq. (3.10), from the solution $\psi(r)$ as $c_<(r) = \psi'(r)K/4\pi r^2\lambda$. However, it is observed that the numerical integration in Eq. (3.35) yields smaller errors.

3.3.3. Perturbative solution for slow reactions

Slow reactions, $\lambda \ll DR^2$, corresponding to a well-mixed reaction volume, are described by a large penetration depth $\kappa^{-1} \gg R$. This suggests to expand the concentration profile $c_<(r)$ in the small parameter $\kappa R \ll 1$, introducing functions $c_0, c_1, \dots$:

$$c_<(r) = c_0(r) + (\kappa R)^2 c_1(r) + \mathcal{O}((\kappa R)^4)\,; \tag{3.36}$$

here, terms of order $(\kappa R)^4$ and higher are neglected. Corresponding fluxes $j_0(r), j_1(r), \dots$ are defined by virtue of Eq. (3.11). Inserting the expansion into Eq. (3.10) for $r \leqslant R$ and sorting by powers of $\kappa^2 = \lambda/D$, one finds that the 0^{th} order is satisfied by the equilibrium distribution in the absence of reactions:

$$c_0(r) = c_A \mathrm{e}^{-\beta U(r)}\,, \tag{3.37}$$

which is accompanied by a vanishing flux, $j_0(r) \equiv 0$, due to detailed balance. The flux $j_1(r)$ at order $(\kappa R)^2$ obeys

$$\frac{1}{r^2}\,\partial_r r^2 j_1(r) = -\kappa^2 D\, c_0(r)\,, \tag{3.38}$$

which can be integrated to yield

$$j_1(r) = -\frac{\kappa^2 D c_A}{r^2} \int_0^r \mathrm{e}^{-\beta U(s)} s^2 ds \tag{3.39}$$

for $0 \leqslant r \leqslant R$, where the boundary condition $j(0) = 0$ [Eq. (3.13)] was used. With this, the reaction rate constant k follows from Eq. (3.7) straightforwardly:

$$k = \kappa^2 D \int_0^R \mathrm{e}^{-\beta U(r)}\, 4\pi r^2 \mathrm{d}r + O\big((\kappa R)^4\big)\,. \tag{3.40}$$

It allows for a simple interpretation valid for slow reactions: the macroscopic rate $k \simeq \lambda V_{\text{eff}}$ is the product of the reaction propensity λ and an effectively accessible reaction volume [37],

$$V_{\text{eff}} = \int_{|\mathbf{r}| \leqslant R} \mathrm{e}^{-\beta U(\mathbf{r})}\, \mathrm{d}^3 r\,. \tag{3.41}$$

3.3.4. Numerical details

The computation of the reaction rate [Eq. (3.32)] for arbitrary potentials and reaction parameters requires the numerical solution of the boundary-value problem, Eq. (3.29), and of the integral, Eq. (3.22). To check the numerical implementation, it is compared

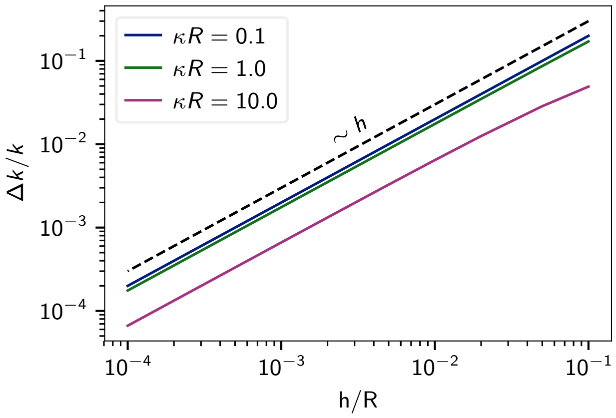

Figure 3.2.: Relative error $\Delta k/k$ of the reaction rate constant k of the numerical solution [Eq. (3.32)] with respect to the analytical solution [Eq. (3.43)] for a diverging potential [Eq. (3.42)]. The numerical result is obtained for different discretisation widths h given in units of the reaction radius R and for different reactivities κR. The dashed line depicts a linear scaling, $\Delta k/k \sim h$. Reprinted from *The Journal of Chemical Physics* "Diffusion-influenced reaction rates in the presence of pair interactions", Dibak et al., 2019, with the permission of AIP Publishing.

to the analytically exactly tractable, albeit peculiar case of a logarithmic potential,

$$U(r) = \begin{cases} -2k_B T \log(r/R), & r < R \\ 0, & \text{otherwise.} \end{cases} \tag{3.42}$$

With this, $g(r) = R^{-2}\,\theta(R-r)$ is a step function, and the coefficient $\beta U'(r) - 2/r$ in Eq. (3.29a) reduces to $-4/r$. The differential equation can be solved using computer algebra, yielding $\psi'(R)$ and the reaction rate according to Eq. (3.32) as

$$k = 2\pi D R \left\{ 3 - \frac{(\kappa R)^2}{(\kappa R)^2 - 2[\kappa R \coth(\kappa R) - 1]} \right\}. \tag{3.43}$$

The Debye rate was computed via the adaptive quadrature routines from QUADPACK. For numerical solutions to Eq. (3.29), the method of finite differences [113] was used by discretising the domain $[0, R]$ into N sub-intervals of equal size $h := R/N$. Note that at the outer most grid points, $r = 0$ and $r = R$, Eq. (3.29a) does not require evaluation if central differences are used to compute $\psi'(r)$ and $\psi''(r)$ from $\psi(r)$. For a range of values of κR, the error Δk between the numerical and the analytical results for the rate was computed, see Fig. 3.2. The relative error $\Delta k/k$ scales approximately linearly with h and decreases with increasing κR. For the worst case studied, $\kappa R = 0.1$, it was concluded that an accuracy better than 10^{-3} is reached by choosing a grid spacing of $h = 10^{-4}R$, which is still well feasible in terms of computational costs. This value of h is used for all subsequent calculations.

Finally, it was checked that all terms in Eq. (3.29a) are bounded. In particular, that the term $[\beta U'(r) - 2/r]\psi'(r)$ vanishes in the limit $r \to 0$. The expression is proportional to $[\beta U'(r) - 2/r]r^2 c(r)$ after re-substituting $\psi(r)$ and using Eq. (3.10). Further, it is anticipated that the concentration profile is bounded from above by the equilibrium distribution, $c(r) \leqslant c_A \mathrm{e}^{-\beta U(r)}$, as reactions can only lower the concentration in the reaction volume, see Fig. 3.8. With this, $(2/r)\, r^2 c(r) \to 0$ and $|\beta U'(r)\, c(r)| \leqslant c_A \left|\partial_r \mathrm{e}^{-\beta U(r)}\right|$, and it remains to show that $\left|\partial_r \mathrm{e}^{-\beta U(r)}\right| \xrightarrow{r \to 0} 0$. This is fulfilled by certain logarithmic potentials, such as in Eq. (3.42), and by algebraically diverging potentials, $\beta U(r \to 0) \simeq a r^{-m}$ with $a, m > 0$. In the latter case, putting $y := r^{-m}$ one finds $\left|\partial_r \mathrm{e}^{-\beta U(r)}\right| \simeq am\, y^{(m+1)/m} \mathrm{e}^{-ay} \to 0$ as $y \to \infty$.

3.4. iPRD simulations

Complementary to the preceding theoretical analysis, extensive simulations of the microscopic reaction-diffusion dynamics in the steady state were performed in order to measure the absolute reaction rate k of the reaction (3.1) and the radial distribution function $c(r)$ of A molecules relative to a B molecule.

3.4.1. Simulation setup and protocol

Stochastic simulations of the interacting particle-based reaction–diffusion dynamics (iPRD) are performed with the software ReaDDy 2 [69, 112], which integrates the motion of particles and reactions between them explicitly in three-dimensional space. In ReaDDy, time is discretised into steps of fixed size Δt. A single step consists of first integrating the Brownian motion of molecules via the Euler–Maruyama scheme and then handling reaction events according to the Doi model (Section 3.1). After each step, one can evaluate observables, such as the positions of particles or the number of reactions that occurred.

The simulation setup is constructed spherically symmetric around a single B molecule in the coordinate origin, as depicted in Fig. 3.1. In particular, a spherical domain of finite radius L was used, which will be filled with A molecules such that at the boundary, $r = L$, the concentration $c(L)$ of A molecules matches a given constant. Within the whole domain, A particles diffuse subject to the interaction potential $U(r)$, whereas the B molecule is fixed in space; here, only potentials that are cut off at a distance $r_c < L$ are studied. The conversion reaction (3.1) takes place with reaction propensity λ inside the sphere with $r \leqslant R$. A large number of simulations is run for varying propensity λ and different potentials $U(r)$, see below. Simulation units were chosen such that distances are measured in terms of the reaction radius R, energies in terms of the thermal energy $k_B T$, and times in terms of the combination $\tau_\mathrm{d} := R^2/D$, which is proportional to the time to explore the reaction volume by diffusion. The parameters used are listed in Table 3.1, in particular, a time step $\Delta t = 10^{-4}\tau_\mathrm{d}$ was used throughout production runs.

Quantity	Symbol	Value	Unit
Propensity of reaction (3.1)	λ	varies	τ_d^{-1}
Soft repulsion strength	b	40	$k_\mathrm{B}T/R^2$
Soft repulsion range	r_0	1	R
LJ interaction strength	ε	1	$k_\mathrm{B}T$
LJ interaction range	σ	$(26/7)^{-1/6}$	R
LJ cutoff radius	r_c	2.5	R
Integration time step	Δt	10^{-4}	τ_d
Radius of simulation domain	L	10	R
Width of factory shell	ΔL	5	R
Number of factory particles	N_f	1.5×10^4	1
Propensity to create A	f_+	0.01	τ_d^{-1}
Propensity to absorb A	f_-	0.01	τ_d^{-1}

Table 3.1.: Parameters used in the particle simulations. Basic units of length, time, and energy are R, $\tau_\mathrm{d} := R^2/D$, and $k_\mathrm{B}T$, respectively.

This chosen time step is sufficiently small to be suitable for the Lennard-Jones potential, which generally calls for much smaller integration steps than the harmonic repulsion due to an increased stiffness.

Aiming at the simulation of a stationary reaction kinetics, the domain is coated by a *factory shell*, with radial coordinates in $r \in [L, L + \Delta L]$, that yields a constant supply of A molecules. Adjacent to the shell, for $r \geqslant L + \Delta L$, an external harmonic potential is added that prevents A molecules from escaping and thereby closing the simulation domain. The factory shell contains N_f factory (F) particles, which are fixed in space at random positions according to a uniform distribution. F particles create and absorb A molecules through the reversible reaction

$$\mathrm{F} \underset{f_-}{\overset{f_+}{\rightleftharpoons}} \mathrm{F} + \mathrm{A}\,. \tag{3.44}$$

The forward reaction has propensity f_+ and is of fission type: a new A molecule is placed at a random distance $d \in [0, R_f]$ from the active F particle. The backward reaction is of fusion type, by which an A molecule is absorbed with propensity f_- if it is closer than R_f to an F particle. Due to the fact that the number of F particles is conserved, the factory reactions (3.44) are pseudo-unimolecular, i.e. they can be reduced to

$$\mathrm{A} \rightleftharpoons \varnothing, \tag{3.45}$$

which leads to a steady-state concentration $c(L)$ of As. The latter depends also on the outflux $K = 4\pi L^2 |j(L)|$ of A molecules, which can diffuse freely into and out of this shell and migrate towards the origin due to the reaction of interest, Eq. (3.1). Lacking an *a priori* knowledge of the concentration $c(L)$ and the concentration c_A in the far field ($r \to \infty$), simulations with a certain set of parameters N_f, f_+, f_-, and R_f are run and the resulting value of c_A is estimated accurately from the observed steady-state profile $c(r)$. Specifically, the solution $c(r) = c_A - K/4\pi Dr$ [Eq. (3.21)], is fit to the data for $c(r)$ in the range $\max(R, r_c) \leqslant r \leqslant L$, where both interactions and reactions are absent and A molecules diffuse freely. This yields the extrapolated concentration at far distances, $c(r \to \infty) = c_A$. Note that the reaction frequency K is directly available from the simulation by counting reaction events.

The above procedure relies on the fact that shifting the upper boundary from infinity to $r = L$ merely shifts the concentration $c(r)$ by an additive constant, leaving the integral flux through spheres of radius r unchanged, provided that r is outside of the interaction range. This is a consequence of Gauss's theorem, see also Eq. (3.11). Therefore, simulation results with a finite volume can be mapped exactly to the infinite case upon using the effective far-field concentration c_A as determined above.

A data production cycle starts with uniformly distributing A molecules in the factory shell with a concentration that roughly anticipates the expected c_A. This initial state is relaxed by evolving the reaction–diffusion dynamics for a time span of $t_{\mathrm{eq}} = 300\tau_{\mathrm{d}}$, by executing 3×10^5 integration steps with a coarser time step size of $\Delta t = 10^{-3}\tau_{\mathrm{d}}$. Equilibration is verified by observing that the number of A particles does not vary significantly. The time step is then decreased to $\Delta t = 10^{-4}\tau_{\mathrm{d}}$ and the system equilibrated for another time span of $30\tau_{\mathrm{d}}$. During the subsequent production run of length similar to t_{eq}, the two main observables are recorded:

(i) the concentration profile $c(r)$ as the radial distribution function (RDF) of A molecules relative to the B molecule in the center, and

(ii) the number of reactions (3.1) that were performed in each integration step, yielding the reaction frequency K and thus the macroscopic reaction rate constant $k = K/c_A$.

Observing the RDF in the case without a reaction and comparing it against the Boltzmann distribution is used to verify the time step.

One such simulation procedure took roughly 512 hours on a single CPU. Simulations were run for 3 different potentials and 5 different propensities. Statistical averages were taken over 13 independent realisations, altogether yielding 195 simulations that were run in parallel. The cumulative CPU time amounts to 100,000 hours.

3.4.2. Pair potentials

In the following, two different isotropic pair potentials are considered for the interaction between A and B molecules, and are compared to the non-interacting case ($U = 0$). The employed potentials are visualized in Fig. 3.3, and all relevant parameters are given in Table 3.1. The first potential describes an ultra-soft steric repulsion, which is common for macromolecules such as polymer rings [114]. For simplicity, it is assumed that A and B molecules repel each other only when their centres are within a cutoff radius r_0, and a harmonic form of interaction is used:

$$U(r) = \frac{1}{2}b(r - r_0)^2 \,, \quad r \leqslant r_0 \,, \tag{3.46}$$

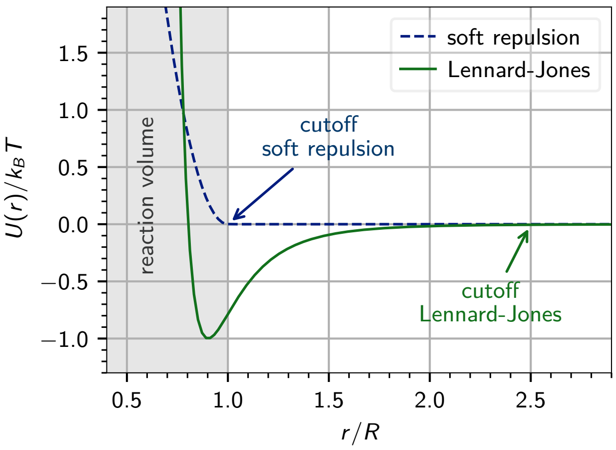

Figure 3.3.: Pair potentials $U(r)$ used in the study of the steady-state reaction kinetics [Eqs. (3.46) and (3.47)] for the parameters given in Table 3.1. The separation r of molecule centres is given in units of the reaction radius R, and the potential energy U is given in terms of the thermal energy k_BT; the shaded region marks the reaction sphere in which reactions (3.1) can occur. Arrows indicate the location of the interaction cutoffs. Reprinted from *The Journal of Chemical Physics* "Diffusion-influenced reaction rates in the presence of pair interactions", Dibak et al., 2019, with the permission of AIP Publishing.

and $U(r) = 0$ otherwise; here, $b > 0$ is a harmonic spring constant chosen to be stiff, $br_0 \gg k_B T$, and the cutoff is set to be equal to the reaction radius, $r_0 = R$.

The second potential is a commonly truncated form of the Lennard-Jones (LJ) potential, which combines a strong steric repulsion of nearly overlapping molecules with a short-range attraction due to van der Waals forces:

$$U(r) = 4\varepsilon \left[(\sigma/r)^{12} - (\sigma/r)^6\right] \theta(r_c - r), \tag{3.47}$$

with σ and $\varepsilon > 0$ being a length and an energy, respectively, that set the range and the strength of the interaction. The value of ε is also the depth of the potential well at $r = \sigma$. Here σ is chosen such that the potential minimum lies *within* the reaction volume, specifically, the inflection point of $U(r)$ is set at the boundary, $R = (26/7)^{1/6}\sigma \approx 1.24\sigma$. The attractive part of the interaction is truncated at $r_c = 2.5R$.

3.5. Results and discussion

3.5.1. Macroscopic rates

Simulation results for the reaction rate constant k as a function of the propensity $\lambda = \kappa^2 D$ are shown in Fig. 3.4 for the above potentials. They are compared to the theoretical predictions from the reaction–diffusion problem, Eq. (3.3), as follows: For the non-interacting case ($U = 0$), the exact solution is available in closed form, Eq. (3.26). For the soft repulsion and the LJ potential, the solution is available only in quasi-analytic form, Eq. (3.32), i.e., the final expressions for k are explicit in terms of a numerical quadrature as in the Debye problem and the numerical solution to a one-dimensional boundary value problem in the interior of the reaction sphere, see Section 3.3.4. As dimensionless control parameter the combination $\kappa R = R\sqrt{\lambda/D}$ is chosen, which distinguishes the reaction- and diffusion-limited regimes, $\kappa R \ll 1$ and $\kappa R \gg 1$, respectively. Equivalently, $(\kappa R)^2 = \lambda\tau_\mathrm{d}$ controls the reaction propensity relative to the diffusion time $\tau_\mathrm{d} = R^2/D$.

For all choices of the potential, the agreement between theory and simulations is excellent, see Fig. 3.4 (a). In all three cases, the reaction rate k increases monotonically with the reaction propensity λ and saturates at Debye's result, Eq. (3.22), for a diffusion-limited reaction ($\kappa R \to \infty$). In this limit, the reaction occurs almost surely upon first contact and details inside of the reaction volume become irrelevant, the formation rate diverges, $k_\mathrm{f} \to \infty$. Note that for the truncated soft repulsion, Eq. (3.46), the limiting value equals the Smoluchowski rate as the potential is zero in the outer domain. For slow reactions, $\kappa R \ll 1$, the initial increase of k depends quadratically on κR and it coincides with the prediction $k \simeq \lambda V_\mathrm{eff}$ of perturbation theory, Eq. (3.40). This regime is better visualised by normalising k with the perturbation result for the non-interacting case, $k^{(0)} = \lambda V_R$, where $V_R = (4\pi/3)R^3$, see Fig. 3.4 (b). From the limit $\kappa R \to 0$ it is evident that also the constant of proportionality V_eff as calculated from Eq. (3.41) matches very well with the numerical results. For $\kappa R = 0.2$, noticeable relative deviations are seen in the simulation data, indicating that the slow-reaction regime is challenging to explore by the particle-based approaches such as iPRD. The figure shows further that the

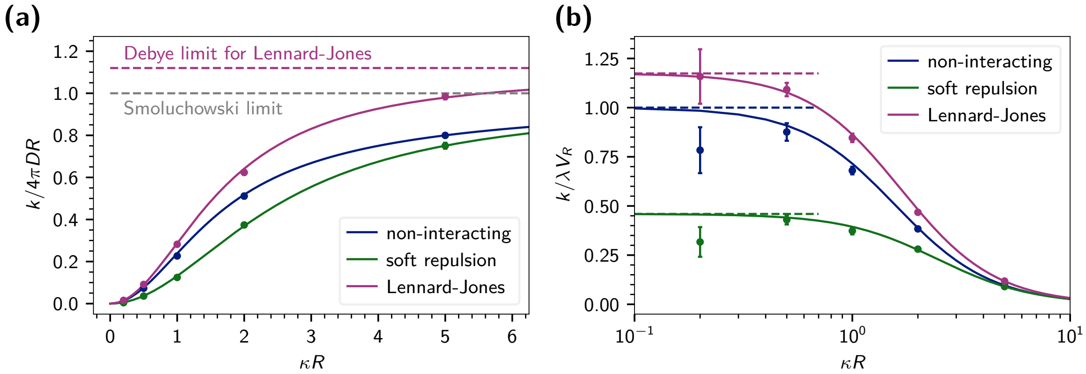

Figure 3.4.: **(a):** Macroscopic rate constant k as a function of the reactivity κR with the inverse penetration depth $\kappa = \sqrt{\lambda/D}$ and the reaction radius R for different pair potentials $U(r)$. Data are given relative to the Smoluchowski rate constant $4\pi DR$ (grey dashed line) in terms of the relative diffusion constant $D = D_\mathrm{A} + D_\mathrm{B}$ and the reaction radius R. Symbols are results of interacting particle-based stochastic simulations of the reaction–diffusion process (iPRD simulations). Solid lines show theoretical predictions obtained from exact expressions [non-interacting case, Eq. (3.26)] or quasi-analytic solutions [soft harmonic repulsion and LJ potential, Eq. (3.32)] of the reaction–diffusion problem, Eqs. (3.3). The pink dashed line indicates the Debye limit, Eq. (3.22), for the LJ potential. **(b):** Macroscopic rate constant k as a function of the reactivity κR normalized by the perturbative solution $k^{(0)} \simeq \lambda V_R$ of the non-interacting case for slow reactions [Eq. (3.40)]. Dashed lines indicate the ratios of the accessible to the total reaction volume V_eff/V_R for each potential [Eq. (3.41)], which is the prediction of the perturbation theory. Reprinted from *The Journal of Chemical Physics* "Diffusion-influenced reaction rates in the presence of pair interactions", Dibak et al., 2019, with the permission of AIP Publishing.

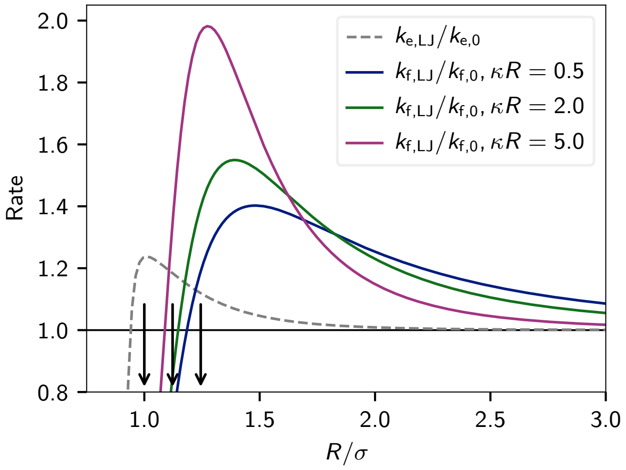

Figure 3.5.: Dependence of the partial reaction rates k_e and k_f on the attractive part of the LJ potential with depth $\varepsilon/k_BT = 1$, which is tested by varying the interaction range σ for fixed reaction radius R. The rates are normalised by their values for the non-interacting case, Eqs. (3.22) and (3.27). Black arrows indicate the zero crossing, the minimum, and the inflection point of the Lennard-Jones potential. Reprinted from *The Journal of Chemical Physics* "Diffusion-influenced reaction rates in the presence of pair interactions", Dibak et al., 2019, with the permission of AIP Publishing.

perturbation solution deviates by no less than 10% from the full solution for $\kappa R \lesssim 0.5$.

How is the reaction rate constant k changed due to the presence of the investigated potentials? A repulsion within the reaction volume slows down the reaction relative to the non-interacting case, which can be attributed to the greatly diminished accessible reaction volume (Fig. 3.4, soft repulsion). The effect is most pronounced for slow reactions, which are most sensitive to a reduction of the actual penetration depth relative to its value κ^{-1} of the free case. Evaluating Eq. (3.41) for the specific harmonic repulsion used here, V_{eff} and thus k are reduced by a factor of ≈ 2.2 relative to the non-interacting case.

An attractive interaction between A and B molecules, on the other hand, is expected to enhance the encounter rate k_{e} and thus to speed up the overall reaction. Already the short-ranged well of the truncated LJ potential, Eq. (3.47), suffices to increase k_{e} by 12 % with respect to the free case, Eq. (3.22). Noting that only the part of the potential outside of the reaction volume, $r > R$, contributes to k_{e}, the dependence on the attraction can be tested by varying the interaction range σ at fixed R, see Fig. 3.5. The encounter rate becomes maximal at $\sigma = R$, i.e., when the integral in Eq. (3.22) is taken over the full domain where the potential is negative, $U(r) < U(r \to \infty)$.

The ramifications of the potential on the formation rate k_{f} are more subtle: the strongly repulsive part of the LJ potential should lead to a decrease as the accessible reaction volume is diminished. At the same time, the potential well induces an enrichment of

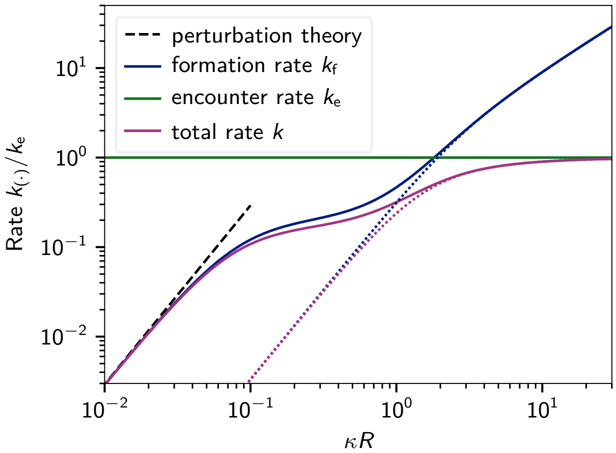

Figure 3.6.: Encounter, formation and total rate constants as a function of the reactivity κR by changing the propensity $\lambda = \kappa^2 D$ for a Lennard-Jones potential with energy $\varepsilon/k_B T = 13$ and reaction radius $\sigma/R = 0.1$. The black dashed line shows the perturbative solution where $k \propto \kappa^2$, and the dotted lines the respective rates in the non-interacting case. Reprinted from *The Journal of Chemical Physics* “Diffusion-influenced reaction rates in the presence of pair interactions”, Dibak et al., 2019, with the permission of AIP Publishing.

A molecules at the boundary of the reaction volume, which would increase k_f. The combination of both can lead to a non-monotonic dependence of the formation rate on the position of the reaction boundary relative to the potential well, which indeed is observed in the numerical solutions to Eq. (3.33), see Fig. 3.5. The position of the maximum in k_f depends on κR and shifts towards larger σ/R for higher reaction propensity. For the parameters given in Table 3.1, the effectively accessible reaction volume is *increased* by $\approx 17\%$ over the free volume V_R [Fig. 3.4 (b)], and for all κR the overall rate constant k is larger than for non-interacting molecules.

By the Markov property of the microscopic reaction–diffusion process, the total reaction rate constant k is the harmonic mean of the partial rates for encounter and formation, Eq. (1.20), and thus, k is bounded from above by the smaller rate: $k \leqslant \min(k_\mathrm{e}, k_\mathrm{f})$. The relative importance of both processes depends on the rescaled reaction propensity κR, which is nicely seen from Fig. 3.6 for the Lennard-Jones potential with $\sigma/R = 0.1$ and $\varepsilon/k_B T = 13$. One reads off that the formation and diffusion-limited regimes, where the other contribution can safely be neglected, are delimited by $\kappa R \lesssim 10^{-1}$ and $\kappa R \gtrsim 10^1$, respectively. Inbetween, there is a wide window of propensities, where both processes enter the overall rate constant. Here, an enhanced availability of reactants due to the deep potential well compensates a slower reaction propensity so that the formation rate displays an approximately plateau-like behaviour for $0.1 \lesssim \kappa R \lesssim 0.5$. For sufficiently fast reactions, the accumulation disappears and k_f starts increasing again towards its large κR behaviour, $k_\mathrm{f} \sim \kappa R$, which resembles the potential-free case as reactions are confined

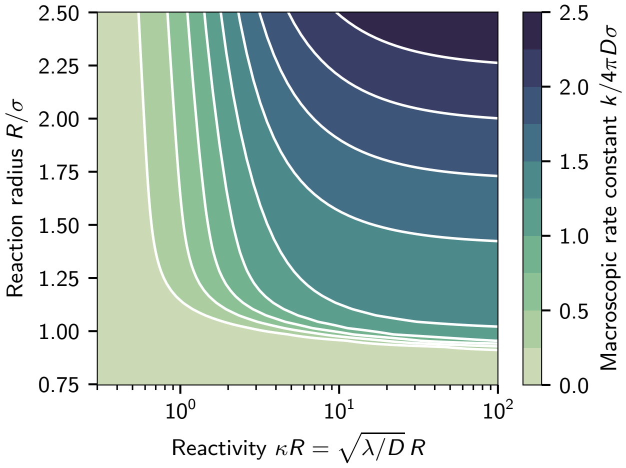

Figure 3.7.: The macroscopic rate constant k in the presence of a Lennard–Jones potential with particle diameter σ and energy depth that is equal to the thermal energy $\varepsilon = k_BT$. Here k is a function of the unit–less reactivity $\kappa R = \sqrt{\lambda/D}\,R$ and a function of the reaction radius R, with the microscopic rate constant λ, relative diffusion constant D. k is given in units of $4\pi D\sigma$, which is the encounter rate up to particle diameter if no reaction and potential would be present. Reprinted from *The Journal of Chemical Physics* "Diffusion-influenced reaction rates in the presence of pair interactions", Dibak et al., 2019, with the permission of AIP Publishing.

to a thin shell near $r = R$. Note that k_f is a monotonic function of κR, which follows from Eq. (3.34) and anticipating the monotonic decrease of $c(R)$ as κR increases, see Fig. 3.8.

Motivated by the practical question how to choose the model parameters λ and R for given reaction rate k and diffusivity D and given interaction potential, the dependence of k on both the propensity κR and the reaction radius R/σ was further scrutinized, exemplified for the Lennard-Jones potential (Fig. 3.7). For slow reactions, $\kappa R \lesssim 1$, the rate constant k is insensitive to the reaction radius. In the diffusion-limited regime, $\kappa R \gtrsim 10$, the rate constant k mainly depends on the reaction radius R/σ and is insensitive to the value of κR. Inbetween, $1 \lesssim \kappa R \lesssim 10$, both parameters must be adjusted carefully. From physical considerations, the reaction radius R should be comparable to the molecular radius σ, which delimits the freedom in the choice of λ.

3.5.2. Concentration profiles

Simulation results for the concentration profile $c(r)$, more precisely, the radial distribution of A molecules relative to Bs, are shown in Fig. 3.8 for three different propensities λ, expressed in terms of $\kappa = \sqrt{\lambda/D}$, and for the different interactions considered above. The data are compared to the theoretical predictions developed in Sections 3.2 and 3.3, and

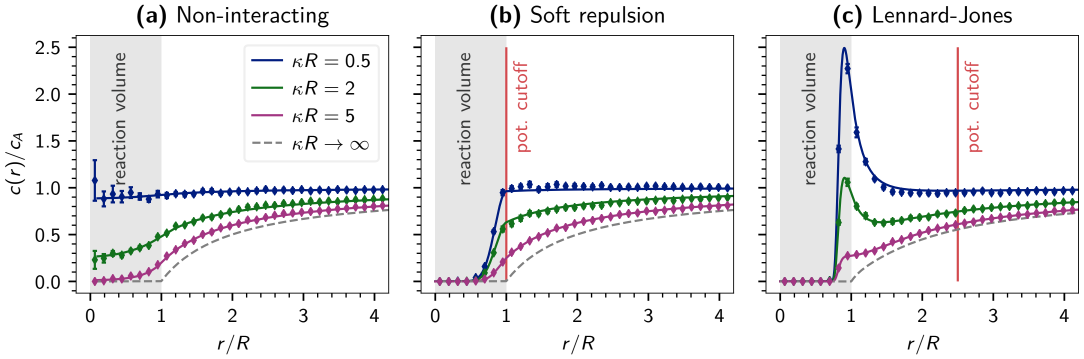

Figure 3.8.: Radial distribution $c(r)$ of A molecules around a B molecule for different reaction propensities λ, here expressed by $\kappa = \sqrt{\lambda/D}$. The panels show results for **(a)** the non-interacting case, **(b)** the soft harmonic repulsion [Eq. (3.46)], and **(c)** a truncated LJ potential [Eq. (3.47)]. Data points are results from iPRD simulations, and solid lines theoretical predictions from Eqs. (3.21) and (3.25) for the non-interacting case and from numerical solutions to Eqs. (3.20) and (3.29) otherwise. Grey dashed lines represent the limit $\kappa R \to \infty$ of almost sure reactions upon contact [Eq. (3.23)]. Grey shaded areas mark the interior of the reaction volume ($r \leqslant R$), and vertical lines indicate the respective positions r_c of the potential cutoffs. Reprinted from *The Journal of Chemical Physics* "Diffusion-influenced reaction rates in the presence of pair interactions", Dibak et al., 2019, with the permission of AIP Publishing.

the quantitative agreement is very good for all cases studied. Thus, the iPRD simulations corroborate the theoretical analysis and the numerical results, which in turn are used to validate the implementation of the simulation algorithm.

For the non-interacting case [Fig. 3.8 (a)], closed analytic expressions for $c(r)$ inside and outside of the reaction volume are available, Eqs. (3.25) and (3.21), respectively. For the soft repulsive and the LJ potentials [Eqs. (3.46) and (3.47)], profiles in the outer domain are obtained from Eq. (3.20) by a quadrature, and in the inner domain from the numerical solution for $\psi'(r)$ of the boundary value problem, Eq. (3.29). At distances $r > r_c$, where neither a reaction can occur nor a potential is present, the constant flux implies for the profile, $c(r) = c_A(1 - k/4\pi Dr)$, see Eq. (3.21).

For slow reactions, $\kappa R \ll 1$, the concentration profile at leading order in κR is expected to equal the equilibrium distribution, $c_0(r) = c_A \mathrm{e}^{-\beta U(r)}$, subject to the specific boundary condition $c(r \to \infty) = c_A$ [Eq. (3.37)]. Indeed, for $\kappa R = 0.5$ both the numerical and simulation results for $c(r)$ are hardly distinguishable from $c_0(r)$ in all three cases studied, see Fig. 3.8; for $U = 0$ it holds $c_0(r) = c_A$ everywhere. Upon increasing κR, the concentration is decreasing uniformly and, in the limit of an instantaneous product formation, $\kappa R \to \infty$, the profile $c(r)$ vanishes inside the reaction volume and approaches Debye's solution, Eq. (3.23), outside as expected. For the non-interacting case and the soft repulsive potential, the latter simplifies to Smoluchowski's result, $c(r) = c_A(1 - R/r)$ for $r \geqslant R$; for the truncated LJ potential used here, the differences are small and hardly seen in the graph (Fig. 3.8 (c)). Summarising, the equilibrium distribution and Debye's solution constitute upper and lower bounds on $c(r)$.

After having understood these limits, the consequences of the interaction potential on the profiles will be discussed in more detail: Adding a soft repulsion within the reaction volume to mimic an excluded volume largely reduces the probability of finding a particle inside the reaction volume [Fig. 3.8 (b)] and thus suppresses the product formation rate k_f [see also Fig. 3.4 (b)]. Yet, the effect is more pronounced for slow reactions as the interior of the reaction volume becomes less and less accessible upon increasing κR, and it is concluded that the repulsion is particularly relevant for slow reactions. The attractive well of the LJ potential on the other hand induces an enrichment of A molecules near the reaction boundary, which is more developed for smaller κR [Fig. 3.8 (c)].

3.6. Conclusion

Bimolecular association processes $\mathrm{A} + \mathrm{B} \to \mathrm{C}$ are fundamental in reactive systems and can describe many different types of interactions. In this chapter reactions of this type were studied for molecules that diffuse in space and exhibit interactions between A and B molecules through a radial potential $U(r)$. With Doi's volume reaction model as a microscopic descriptor, the macroscopic reaction constant k, as well as the concentration profiles $c(r)$ of A particles around B particle were calculated as functions of microscopic reaction propensity λ and the reaction radius R. In this picture the B molecules are assumed to be so dilute that interaction between B and reactive competition can be neglected. The explicit dependence of the model on λ allows systematically probing the kinetics from the well-mixed to the diffusion-limited regime. These regimes can be distinguished by the dimensionless quantity $\kappa R := \sqrt{\lambda/D}R$, which is denoted as the *reactivity* of an AB pair. The length κ^{-1} describes how far molecule centers can penetrate the reaction volume of radius R before they react, with $D := D_\mathrm{A} + D_\mathrm{B}$ being the relative diffusion constant. This approach bridges between the two well-studied cases $\kappa R \ll 1$ (reaction-limited or well-mixed) and $\kappa R \gg 1$ (diffusion-limited or fast-reaction limit), where particles immediately react upon contact.

Over the entire spectrum of κR values and for arbitrary pair potentials, the analytical result for the reaction rate constant exhibits the Markovian decomposition $k^{-1} = k_\mathrm{e}^{-1} + k_\mathrm{f}^{-1}$ into encounter k_e and formation k_f rates [Eq. (1.20)], where k_e is always given by Debye's result Eq. (3.22).

In the well-mixed limit, the rate constant is dominated by the formation rate k_f. Here, perturbation theory (see Section 3.3.3) in small values of κR can be used to find the reaction rate constant $k = \lambda V_\mathrm{eff}$, with the effective reaction volume V_eff given by Eq. (3.41). In the absence of a potential, V_eff simplifies to the volume of the reactive sphere $V_R = (4\pi/3)R^3$. The other side of the spectrum constitutes the diffusion-limited regime, which is dominated by the encounter rate k_e: reactions occur almost surely upon entering the reaction volume. The expression for k derived here reproduces the Smoluchowski encounter rate $k = 4\pi DR$ in the absence of potentials and Debye's result [28], when particles diffuse subject to an interaction potential $U(r)$.

In the application-relevant diffusion-influenced regime (see Section 3.3), where k_e is of comparable magnitude as k_f, semi-analytical expressions were obtained for the rate k and the local concentration $c(r)$ that require numerical evaluation [Eqs. (3.32) and (3.35)]. Practically, one has to solve a one-dimensional boundary value problem for the reaction-diffusion equation inside the reaction volume and compute an integral over the domain outside the reaction volume; the computational costs of both tasks are negligible. The numerical scheme was tested against explicit analytic solutions for a logarithmically repulsive potential, for which the system could be solved using a computer algebra system. A closed expression for the rate k is given for general potentials outside the reaction volume in the case that molecules do not interact if their centers are within the reaction volume [Eq. (3.28)]. This may be useful to model, e.g., reactions in electrolytes while neglecting excluded volume.

The detailed dependence of the rate k on the reactivity parameter κR was studied for two different potentials: a soft harmonic repulsion inside the reaction volume, and a truncated Lennard-Jones potential combining excluded volume and attraction. The numerical results for the rate k and the concentration $c(r)$ show excellent agreement with extensive iPRD simulations.

The presented model is at the core of iPRD simulations, which permit studying reaction processes in cells with spatial resolution at different levels of coarse-graining. In these simulations, the interaction potential $U(r)$ could either be chosen *ad hoc*, based on physical insight, or determined as a potential of mean force from atomistic simulations [59, 115, 116]. The found relations between the microscopic parameters λ, R and the macroscopic rate k enable the development of quantitative iPRD simulations, i.e. the parameters can be tuned, such that experimental values of k are reproduced. In the present study, the dilute limit was explored, where interaction between particles of the same species can be ignored. This serves as a well-defined starting point to study effects arising from higher concentration and crowding on the reaction rate and distribution of particles.

Part II.

Enhanced sampling techniques for molecular systems in equilibrium

4. Introduction

The investigation of many problems in statistical physics includes the computation of equilibrium quantities, common examples of quantities are the magnetization in spin systems [117, 118] or the binding affinity in biological systems [119, 120]. These quantities are usually defined as observables of the system under investigation. More specifically, an observable $\langle O\rangle$ is defined as the ensemble average

$$\langle O\rangle = \int_{\Gamma} O\left(\mathbf{x}\right)\mu\left(\mathbf{x}\right)\, \mathrm{d}\mathbf{x}, \tag{4.1}$$

with configuration space Γ, the Boltzmann distribution $\mu\left(\mathbf{x}\right) = \mathcal{Z}^{-1}\exp\left[-\beta U\left(\mathbf{x}\right)\right]$, where $U\left(\mathbf{x}\right)$ denotes the energy of the system, $\mathcal{Z} = \int_{\Gamma}\exp\left[-\beta U(\mathbf{x})\right]\, \mathrm{d}\mathbf{x}$ the canonical partition function and $\beta = \left(k_{\mathrm{B}}T\right)^{-1}$ the inverse thermal energy with temperature T and Boltzmann constant k_{B}. Solving this integral is generally very difficult, as systems under investigation commonly have many degrees of freedom, making it very high dimensional. A common set of techniques designed to solve such integrals numerically are Monte Carlo (MC) methods. In these, the ensemble average Eq. (4.1) is approximated using samples $\left\{\mathbf{x}_i\right\}_{i\geqslant 0}$ from the stationary distribution $\mu(\mathbf{x})$. This can be denoted using the Dirac measure as $\mu(\mathbf{x}) \approx \sum_{i\geqslant 0}\delta\left(\mathbf{x}-\mathbf{x}_i\right)$. Eq. (4.1) thus reduces to

$$\langle O\rangle \approx \overline{O}_N = \frac{1}{N}\sum_{i\geqslant 0} O\left(\mathbf{x}^{(i)}\right). \tag{4.2}$$

According to the central limit theorem, it can be shown that this method converges to the correct value, $\lim_{N\to\infty}\overline{O}_N \to \langle O\rangle$, if the observable has a finite variance $\sigma^2 = \mathrm{var}\left[O\right] < \infty$ [121]. Furthermore, assuming that the samples are independent and identically distributed (i.i.d), it can be shown that this limit is reached independently of the dimensionality of the underlying system as

$$\overline{O}_N \approx \mathcal{N}\left(\langle O\rangle, \frac{\sigma^2}{N}\right), \tag{4.3}$$

Thus, the error of the MC method scales as $\mathcal{O}\left(N^{-1/2}\right)$. While in low dimensions other numerical integration methods converge considerably faster, their complexity rapidly increases with the dimensionality of the system whereas the rather slow convergence of the MC methods is independent of the dimensionality of the system. Therefore, they are a common and often the only feasible choice to compute observables in high-dimensional systems. It should be further noted, that the error in all MC methods scales proportionally to $N^{-1/2}$, and the common problem when designing new methods is usually to improve the factor of proportionality.

MC samples are either generated dynamically or statically: dynamic Monte Carlo methods use a Markovian stochastic process to generate a new sample $\mathbf{x}'$ given the current state $\mathbf{x}$ of the system. Evolving this process results in a trajectory of configurations that are used as samples. Note that this dynamics does not need to correspond to the *physical* dynamics of the system. Even though the Markov property holds for this process, a sequence of samples will still inhibit a certain amount of correlation which reduces the effective sample size. This is further discussed in Section 4.2.

Static MC generates a set of samples in one shot, e.g. by transforming samples from a different distribution to the target distribution. Doing this in an exact fashion is generally not feasible for high-dimensional distribution. Generally, samples will be drawn from a biased distribution $q(\mathbf{x})$. If $\mu(\mathbf{x})$ and $q(\mathbf{x})$ have a sufficient amount of energetic overlap, a reweighting technique can be applied in order to remove any bias from the samples. This again reduces the number of effective samples and faces limitations in high dimensions. This will be further discussed in the following section.

4.1. Importance sampling and reweighting

Directly generating samples that follow the Boltzmann distribution $\mu(\mathbf{x})$ is a challenging task. Suppose however it would be possible to generate i.i.d. samples from an approximation distribution $q(\mathbf{x})$ with sufficient energetic overlap with $\mu(\mathbf{x})$. Then, a method called importance sampling [122] can be applied. The latter gives a weight w_i to all samples $\mathbf{x}_i$ drawn from $q(\mathbf{x})$ and is able to compute unbiased expectation values with respect to $\mu(\mathbf{x})$. This can be shown directly from the definition of the expectation by inserting a factor of unity

$$\langle O\rangle = \int_\Gamma O(\mathbf{x})\mu(\mathbf{x})\mathrm{d}\mathbf{x} = \int_\Gamma O(\mathbf{x})\frac{\mu(\mathbf{x})}{q(\mathbf{x})}q(\mathbf{x})\mathrm{d}\mathbf{x}. \tag{4.4}$$

With samples $\mathbf{x}_1, \ldots, \mathbf{x}_N \sim q(\mathbf{x})$, an estimator of the expectation above can be found as

$$\widetilde{O}_N = \frac{\sum_i O(\mathbf{x}_i)w_i}{\sum_i w_i}, \tag{4.5}$$

with the unnormalized weights $w_i \propto \mu(\mathbf{x}_i)/q(\mathbf{x}_i)$. Defining the normalized weights as $\widehat{w}_i = w_i/\sum_i w_i$, this further simplifies to $\widetilde{O}_N = \sum_i \widehat{w}_i O(\mathbf{x}_i)$. Note that finding a good trial distribution q is a highly nontrivial task, especially in high-dimensional systems, where the fraction of accessible volume is vanishingly small. In these cases, a poor approximate q may lead to only a few of the samples having substantial weight. To asses the quality of the generated samples, the efficiency of the sampling has to be investigated.

4.1.1. Sampling efficiency of a static Monte Carlo method

As the weights $\widehat{w}_i$ in an importance sampling scheme can become vanishingly small, the question arises, how a certain number of weighted samples from the approximated distribution q compares to a set number of samples from the correct target distribution μ. One way to assess this quantity is to compare the variance between an estimate of a reweighted observable and direct samples from the target for a fixed number N of i.i.d. samples. This is expressed in the estimated sample size

$$\mathrm{ESS} = N \frac{\mathrm{var}_\mu \left[\overline{O}_N\right]}{\mathrm{var}_q \left[\widetilde{O}_N\right]}. \tag{4.6}$$

Following Liu [123], the ESS can be approximated in terms of the coefficient of variation

$$\mathrm{cv}^2(w) = \frac{\sum_i (w_i - \bar{w})^2}{(N-1)\bar{w}^2}, \tag{4.7}$$

with the sample mean $\bar{w}$, as

$$\mathrm{ESS} = \frac{N}{1 + \mathrm{cv}^2(w)}. \tag{4.8}$$

Consider the simple illustrative example of a multivariate Gaussian $\mu(\mathbf{x}) = \mathcal{N}(\mathbf{x}; 0, \mathbb{1})$, with $\mathbb{1} \in \mathbb{R}^{d\times d}$ that is to be approximated with samples from another multivariate Gaussian with a slightly higher variance $q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; 0, \mathbb{1}(1+\varepsilon))$, with $\varepsilon > 0$. For this simple example, the effective sample size can be computed analytically

$$\mathrm{ESS} \approx \frac{N}{(1 + \varepsilon^2/(2\varepsilon + 1))^{d/2}}. \tag{4.9}$$

This example demonstrates that the effective sample size exponentially decreases with dimension d. It is therefore expected that this procedure exhibits limitations when it comes to scaling to systems consisting of many degrees of freedom. In this example, even a small "error" ε in the proposal distribution leads to an effective sample size that vanishes exponentially with the dimensionality [124]. Moreover, Bengtsson et al. [125] have shown that the maximum weight $\max(\hat{w}_i)$ approaches unity if the number of samples N grows slower than $d^{1/3}$, and thus, the computation of expectations as in Eq. (4.4) fails.

4.2. Markov chain Monte Carlo

Generating independent samples as required in the static MC method described above is in general a very difficult task. The key idea of Markov chain Monte Carlo (MCMC) methods is that, with the correct update scheme, a Markov chain can be guaranteed to converge to the stationary distribution μ. This update rule is given by some stochastic process that generates a transition from state $\mathbf{x}$ to state $\mathbf{y}$ with probability $p(\mathbf{x} \to \mathbf{y})$. In order to guarantee convergence to the stationary distribution, this update rule has to fulfill two conditions [126]

(i) irreducibility: for each pair $\mathbf{x}, \mathbf{y} \in \Gamma$, where Γ denotes the configuration space, there

exists an $n \geqslant 0$, s.t. $p(\mathbf{x} \to \mathbf{y}) = \int \mathrm{d}\mathbf{x}_1 \cdots \int \mathrm{d}\mathbf{x}_N p(\mathbf{x} \to \mathbf{x}_1) \ldots p(\mathbf{x}_N \to \mathbf{y}) > 0$, or in other words each state $\mathbf{x}$ can be reached from any other state $\mathbf{y}$ in a finite number of steps

(ii) stationarity: for each $\mathbf{y} \in \Gamma$ $\int \mathrm{d}\mathbf{x} \mu(\mathbf{x})\, p(\mathbf{x} \to \mathbf{y}) = \mu(\mathbf{y})$

The first condition is hard to prove for a general sampler and update rules are commonly constructed s.t. this property is fulfilled by construction. In the context of physical problems it is also commonly denoted as *ergodicity.* The second condition is usually fulfilled by picking an update rule that fulfills detailed balance

$$\mu\left(\mathbf{x}\right) p\left(\mathbf{x} \to \mathbf{y}\right) = \mu\left(\mathbf{y}\right) p(\mathbf{y} \to \mathbf{x}), \tag{4.10}$$

where $p\left(\mathbf{x} \to \mathbf{y}\right)$ denotes the transition probability from $\mathbf{x}$ to $\mathbf{y}$ of some stochastic process $K\left(\mathbf{x} \to \mathbf{y}\right)$, which is denoted as *transition kernel.* One way to enforce detailed balance for a stochastic process $K\left(\mathbf{x} \to \mathbf{y}\right)$ that proposes transitions $\mathbf{x} \to \mathbf{y}$ with probability $p_{\text{prop}}\left(\mathbf{x} \to \mathbf{y}\right)$ is using an acceptance/rejection step, which is commonly denoted as Metropolis-Hastings algorithm [127]. To this end, the transition probability is split into a proposal and an acceptance probability

$$p\left(\mathbf{x} \to \mathbf{y}\right) = p_{\text{prop}}\left(\mathbf{x} \to \mathbf{y}\right) p_{\text{acc}}\left(\mathbf{x} \to \mathbf{y}\right). \tag{4.11}$$

using the detailed balance condition Eq. (4.10), a criterion for the acceptance probability can be found as

$$p_{\text{acc}}\left(\mathbf{x} \to \mathbf{y}\right) = p_{\text{acc}}\left(\mathbf{y} \to \mathbf{x}\right) \frac{\mu\left(\mathbf{y}\right) p_{\text{prop}}\left(\mathbf{y} \to \mathbf{x}\right)}{\mu\left(\mathbf{x}\right) p_{\text{prop}}\left(\mathbf{x} \to \mathbf{y}\right)}, \tag{4.12}$$

which serves as a condition on the acceptance probability.
With $a := \left[\mu\left(\mathbf{y}\right) p_{\text{prop}}\left(\mathbf{y} \to \mathbf{x}\right)\right] / \left[\mu\left(\mathbf{x}\right) p_{\text{prop}}\left(\mathbf{x} \to \mathbf{y}\right)\right]$, the detailed balance equation is fulfilled for any function fulfilling $p_{\text{acc}}(a)/p_{\text{acc}}(1/a) = a$. The function that maximizes the acceptance probability is the Metropolis-Hastings criterion [122]

$$p_{\text{acc}}\left(\mathbf{x} \to \mathbf{y}\right) = \min\left\{1, \frac{\mu\left(\mathbf{y}\right) p_{\text{prop}}\left(\mathbf{y} \to \mathbf{x}\right)}{\mu\left(\mathbf{x}\right) p_{\text{prop}}\left(\mathbf{x} \to \mathbf{y}\right)}\right\}, \tag{4.13}$$

and most MCMC algorithms use this criterion in order to enforce detailed balance. An example of the standard MCMC algorithm is shown in Algorithm 3. MCMC algorithms differ mainly in the used proposal kernels. A simple implementation draws the proposed new configuration $\mathbf{y}$ from a Gaussian centered at the current position $p_{\text{prop}}\left(\mathbf{x} \to \mathbf{y}\right) = \mathcal{N}\left(\mathbf{y}; \mathbf{x}, \sigma^2\right)$, with fixed variance σ^2. In this case the proposal probability is symmetric, meaning $p_{\text{prop}}(\mathbf{x} \to \mathbf{y}) = p_{\text{prop}}(\mathbf{y} \to \mathbf{x})$ and the acceptance probability Eq. (4.13) simplifies to

$$p_{\text{acc}}\left(\mathbf{x} \to \mathbf{y}\right) = \min\left\{1, \mu\left(\mathbf{y}\right)/\mu\left(\mathbf{x}\right)\right\} = \min\left\{1, \exp\left[-\beta\left(U(\mathbf{y}) - U(\mathbf{x})\right)\right]\right\}. \tag{4.14}$$

While this approach is straightforward, it also has several drawbacks: Convergence of

Algorithm 3: Standard MCMC algorithm

```
input  : $l_s = [\,]$ : empty list for samples
         $p_{\text{prop}}$ : proposal density
         $N_{\text{iterations}}$ : number of generated samples
         $\mathbf{x}_0$ : initial configuration
$\mathbf{x} \leftarrow \mathbf{x}_0$
while $i \leqslant N_{\text{iterations}}$ :
    $\mathbf{y} \leftarrow$ sample from $p_{\text{prop}}(\mathbf{x} \to \mathbf{y})$ ;
    $p_{\text{acc}} \leftarrow p_{\text{acc}}(\mathbf{x} \to \mathbf{y})$ (Eq. 4.13)
    if $r \sim \mathcal{U}(0,1) < p_{\text{acc}}$ :
        $\mathbf{x} \leftarrow \mathbf{y}$
    $l_{\text{s}}$.append ($\mathbf{x}$)
    $i \leftarrow i + 1$
output : list of samples $l_{\text{s}}$
```

this method is rather slow, as samples are highly correlated due to them being only small displacements in configurations space. The relation between efficiency and correlation is discussed in the following section. Furthermore crossing energy barriers is difficult, as steps climbing free energy hills are more likely to be rejected, thus further increasing correlations in the sampled configurations. This led to the development of many methods that aim towards overcoming these problems and lead to a more rapid convergence of the estimator $\overline{O}$. These will be discussed in Section 4.4.

4.3. Efficiency of an MCMC sampler

As noted above, the samples generated from a sequential Monte Carlo sampler are correlated, as a proposed state $\mathbf{y}$ is conditioned on the previous state $\mathbf{x}$. While the error in the sampled observable is of the order $\mathcal{O}(N^{-1/2})$, the pre-factor of that scaling depends on the underlying sampling algorithm. Following [123], the efficiency of an MCMC sampler generating a sequence of samples $\mathbf{x}_1, \ldots, \mathbf{x}_N$, all distributed according to the stationary distribution $\mu(\mathbf{x})$, can be assessed by observing the variance of the estimator of the observable under investigation $\overline{O} = \frac{1}{N}\sum_i O(\mathbf{x}_i)$. The variance is given by

$$\operatorname{var}\left[\overline{O}\right] = \operatorname{var}\left[\frac{\sum_{i=1}^{N} O(\mathbf{x}_i)}{N}\right] = \frac{\sigma^2}{N}\left[1 + 2\sum_{i=1}^{N-1}\left(1 - \frac{i}{N}\right)\varrho_i\right] \approx \frac{\sigma^2}{N}\left[1 + 2\sum_{i=1}^{\infty}\varrho_i\right], \quad (4.15)$$

where $\sigma^2 = \operatorname{var}[O]$ and $\varrho_i = \operatorname{corr}\left[O(\mathbf{x}_1), O(\mathbf{x}_{1+i})\right]$ denotes the normalized time lagged autocorrelation of O. Here, the stationarity of the Markov chain was used in the first equality, stating that for $j > i$ the autocorrelation only depends on the lag time $\operatorname{corr}\left[O(\mathbf{x}_i), O(\mathbf{x}_j)\right] = \operatorname{corr}\left[O(\mathbf{x}_1), O(\mathbf{x}_{1+j-i})\right]$. Furthermore, terms of $\mathcal{O}(N^{-2})$ were omitted in the approximation.

Defining the integrated autocorrelation time as

$$\tau_{\text{int}}(O) = \frac{1}{2} \sum_{i=-\inf}^{\inf} \varrho_i = \frac{1}{2} + \sum_{i=1}^{\infty} \varrho_i, \tag{4.16}$$

the variance of the sampler of the mean can be expressed as

$$N \text{var}\,(O) = 2\tau_{\text{int}}(O)\sigma^2. \tag{4.17}$$

Comparing the variance Eq. (4.17) with that of an estimator using i.i.d. samples Eq. (4.3), one observes that the error is equivalent to a sampler using $N_{\text{eff}} = N/2\tau_{\text{int}}(O)$ i.i.d. samples. Thus, this quantity can be considered the *effective sample size* of a sampler and it can be used as a basis for comparison between different types of samplers. Furthermore, this gives rise to the sampling efficiency, which can be defined as effective sample size per sample

$$\eta = \frac{N_{\text{eff}}}{N} = (2\tau_{\text{int}})^{-1} \tag{4.18}$$

Generally, the sampler with larger efficiency η is preferable.

For large lag times i is is often observed that the autocorrelation time exhibits an exponential decay $|\varrho_i| = \exp\left[-i/\tau_{\text{exp}}\,(O)\right]$ on a timescale $\tau_{\text{exp}}(O)$. With the definition of the integrated autocorrelation time Eq. (4.16) it can be shown that

$$\tau_{\text{int}}(O) = \frac{1}{2} + \sum_{i=1}^{\infty} \varrho_i = \sum_{i=0}^{\infty} \mathrm{e}^{-i/\tau_{\text{exp}}(O)} - \frac{1}{2} = \frac{1}{2} \coth\left(\frac{1}{2\tau_{\text{exp}}(O)}\right) \approx \tau_{\text{exp}}(O), \tag{4.19}$$

where the geometric series was used in the third equality and the approximation holds for large $\tau_{\text{exp}}(O)$ [123].

In the case of a discrete phase space S the transition probabilities are given by the transition matrix $\mathbf{T}$. The slowest *relaxation time*, which is defined by the timescale of the slowest process of the system $\tau_{\text{exp}} = \sup_{O \in L^2(S)} \tau_{\text{exp}}(O)$ is then connected to the second largest eigenvalue λ_2 of the transition matrix $\mathbf{T}$ through $\tau_{\text{exp}} = -\ln \lambda_2$ [128].

This allows for estimating the efficiency of a sampler using Gaussian displacements in a system exhibiting a free energy barrier in configuration space. Returning to the example of transitioning an energy barrier in configuration space (Section 1.2), where the transition rate is proportional to $k_{\text{K}} = \exp\left(-E_b/k_{\text{B}}T\right)$. Standard MCMC with Gaussian displacements approximates overdamped Brownian dynamics [129] and therefore the transition rate of such an MCMC sampler is expected to be proportional to k_{K}. Thus, the integrated autocorrelation time is also approximated to be $\tau_{\text{exp}} = k_{\text{K}}^{-1}$, meaning that the efficiency $\eta = 1/2\tau_{\text{int}} \propto k_{\text{K}}$ of a Gaussian MCMC sampler exhibits Arrhenius-like scaling $\log \eta \propto (k_{\text{B}}T)^{-1}$. This results in extremely slow convergence when high barriers are present in the energy landscape. This problem led to a vast development of enhanced MCMC methods. The following section introduces some of the most common techniques in this field.

4.4. Enhanced Markov chain Monte Carlo techniques

In this section, several improvements of the MCMC algorithm utilizing symmetric Gaussian displacements are discussed.

4.4.1. Multi-temperature methods

A set of common techniques that aims at increasing the sampling efficiency utilizes the fact that the transition rate over energy barriers k_{K} exhibits Arrhenius-like scaling, $\log k_{\mathrm{K}} \propto (k_{\mathrm{B}}T)^{-1}$, and thus samplers running at higher temperatures are producing uncorrelated samples at a higher frequency.

The two most widely recognized methods in this class are simulated [130, 131] and parallel tempering [132–134] which operate on a family of distributions $\left\{\mu_i\left(\mathbf{x}\right) = \mathcal{Z}_i^{-1} \exp -U(\mathbf{x})/k_{\mathrm{B}}T_i\right\}_{i\leqslant N_T}$ parameterized by a set of N_T temperatures $\{T_i\}_{i\leqslant N_T}$. In these types of sampling algorithms, the lowest temperature is the one of interest and higher temperatures are chosen with the sole purpose of enhancing the transition rates over the free energy barriers. The method then utilizes temperature exchange moves to mix the configurations between the different temperatures. In order to fulfill detailed balance, these methods rely on a significant overlap of the energetic distributions at different temperatures, therefore these have to be chosen carefully. Specifically, for a system of size N, the number of required temperature approximately scales as $\sqrt{N}$ [134]. One way of reducing the required number of replicas is by utilizing nonequilibrium simulations [135]. These are performed by swapping the temperatures between the proposed configurations and allowing for relaxation of the energies in the process. The heat that is dissipated in the process has to be taken into account when accepting the resulting move.

Simulated tempering (ST) defines the target distribution on an augmented space $\mathcal{A} = \Gamma \times \mathcal{I}$, where $\mathcal{I} = \{1, \ldots, N_T\}$ is a set of discrete temperature indices. The density of $(\mathbf{x}, i) \in \mathcal{A}$ is given by

$$\mu_{\mathrm{ST}}\left(\mathbf{x}, i\right) \propto c_i \exp\left[-U(\mathbf{x})/k_{\mathrm{B}}T_i\right], \tag{4.20}$$

where the c_i are pre-defined constants that need to be set such that each temperature level roughly has the same probability of being visited. The ST proposal kernel consists of a convex combination of two steps that either operate on the state space variable $\mathbf{x}$ and propose a step $(\mathbf{x}, i) \to (\mathbf{y}, i)$, or operate on the temperature index i, and propose a step $(\mathbf{x}, i) \to (\mathbf{x}, j)$. Both steps enforce detailed balance in the augmented space $(\mathbf{x}, i)$. To obtain the samples at the correct temperature, a final marginalization step has to be taken. Intuitively this algorithm implements Simulated Annealing [136] into an MCMC sampler which can be understood as heating up the system to overcome energy barriers with cooling down to relax the energies, similar to annealing procedures in metallurgy. Temperature swaps are usually proposed between neighboring temperatures only, in order to maximize the acceptance probability of these moves. Under the assumption of tuning the parameters c_i to achieve equal population across all temperatures, the diffusion in the temperature space takes the form of a symmetric random walk in one dimension. This results in a mixing time of $\mathcal{O}(N_T^2)$, which limits the number of temperatures that can be used efficiently [123].

Furthermore, under the assumptions above, only a number of $1/N_T$ of the samples are drawn from the distribution at the temperature of interest, which leads to an efficiency of $\eta \approx 1/\left(N_T \tau_{\text{int}}\right)$. Assuming that the sampler in the augmented space transitions energy barriers as quickly as the highest temperature, the integrated correlation time is reduced to $\tau_{\text{int}} \approx \exp(E_B/k_{\text{B}}T_{\max})$. With these assumptions, the highest temperature must be at least $T_{\max} = T_0 \ln \tau_{\text{int}}/\left(\ln \tau_{\text{int}} - \ln N_T\right)$ for the sampler to be as efficient as just running sampling at the lowest temperature N_T times longer.

Parallel Tempering (PT) also operates at multiple temperatures. The major difference is, that N_T copies, each at its own temperature T_i, are propagated simultaneously. Thus, instead of handling an augmented space as in ST, the sampler operates in the product space $S_{T_1} \times \cdots \times S_{T_N}$. The joint probability of samples in this space is given by

$$\mu_{\text{PT}}(\mathbf{x}_1, \ldots, \mathbf{x}_{N_T}) = \prod_{i=1}^{N_T} \mu^{T_i}(\mathbf{x}_i). \tag{4.21}$$

Again, the system can be propagated in two different fashions. Either all samples are propagated at their respective temperature $(\mathbf{x}_1, \ldots, \mathbf{x}_{N_T}) \rightarrow \left(\mathbf{y}_1, \ldots, \mathbf{y}_{N_T}\right)$, or samples can be swapped between two different temperatures $(\mathbf{x}_1, \ldots, \mathbf{x}_i, \ldots \mathbf{x}_j, \ldots, \mathbf{x}_{N_T}) \rightarrow (\mathbf{x}_1, \ldots, \mathbf{x}_j, \ldots \mathbf{x}_i, \ldots, \mathbf{x}_{N_T})$. Accepting swapping moves with probability

$$p_{\text{acc}}\left((T_i, T_j) \rightarrow (T_j, T_i)\right) = \min\left\{1, \exp\left[U(\mathbf{x}_i) - U(\mathbf{x}_j)\right]\left[1/T_i - 1/T_j\right]\right\} \tag{4.22}$$

enforces detailed balance in the temperature product space and ensures convergence to the Boltzmann distribution at each temperature.

The efficiency of the scheme still depends on the chosen dynamics in the configuration space. In this scheme, several copies of the system are propagated at the same time. This also has to be reflected when computing the efficiency. Specifically, as only the lowest temperature is of importance, N_T update steps have to be performed in order to generate one new sample. Following Eq. (4.18), the efficiency for a parallel tempering sampler is therefore

$$\eta_{\text{PT}} = (2 N_T \tau_{\text{int}}(O))^{-1}. \tag{4.23}$$

An augmented scheme for improved dynamics at several temperatures including a deep learning approach will be discussed in Chapter 6.

Population Annealing Another multi-temperature method is population annealing [137]. In this algorithm, many replicas of the system are initialized at a high temperature. The temperature is then sequentially reduced $(T_{N-1}, \ldots, T_0)$, where $T_i > T_{i+1}$, and T_0 is the target temperature. At each step, the samples are redrawn with the (unnormalized) probabilities $w_j \propto \exp\left[-\left(\beta_i - \beta_{i-1}\right) E_j\right]$, where E_j is the energy of replica j. This ensures that the samples are close to equilibrium at the given temperature. However, the resampled configuration is going to be correlated, as it might contain multiple copies of the same configurations. Additionally, it might also be biased, as the low energy tails are expected to be missing in the ensemble. Both of these problems are partially addressed by running a fixed number of MCMC moves after re-sampling which ensures relaxation of the energies at the current temperature and additionally aids to decorrelate samples [138, 139].

Different versions of these algorithms exist when the ensemble under consideration is not that of constant temperature. A more general approach to the sampling with several ensembles is the replica-exchange or Hamiltonian exchange method [140].

4.4.2. Nonequilibrium path sampling

One potential strength of the MCMC method is that a proposed transition $\mathbf{x} \to \mathbf{y}$ does not need to follow a physical path in phase space. This is utilized in a class of methods [141, 142] which uses short out-of-equilibrium simulations in order to generate far-ranging MCMC moves. The recently developed non-equilibrium chain Monte Carlo [143] gives a general framework for using non-equilibrium processes as a proposal density. A proposal step is generated by applying a sequence of perturbation kernels $\alpha_i(\mathbf{x}, \mathbf{y})$ and propagation kernels $K_i(\mathbf{x}, \mathbf{y})$. The full propagation protocol is given by alternating the perturbation and propagation kernels $\Lambda = (\alpha_1, K_1, \ldots, \alpha_N, K_N)$. The perturbation kernels are meant to drive the system out of equilibrium, while the propagation kernels are meant to relax the system toward equilibrium. Applying these kernels consecutively results in a non-equilibrium path

$$\mathbf{x} = \mathbf{x}_0 \xrightarrow{\alpha_1} \mathbf{x}_1^* \xrightarrow{K_1} \mathbf{x}_1 \xrightarrow{\alpha_2} \cdots \xrightarrow{\alpha_N} \mathbf{x}_N^* \xrightarrow{K_N} \mathbf{x}_N = \mathbf{y}. \tag{4.24}$$

The reverse protocol, denoted as $\tilde{\Lambda} = \left(\tilde{\alpha}_N, \tilde{K}_N, \ldots, \tilde{\alpha}_1, \tilde{K}_1\right)$, where the $\tilde{\zeta}$ denotes the inverse of the protocol ζ. Computing the exact proposal probability of such a move is in general not feasible, as the dynamics inherently are stochastic and depend on many intermediate states which would have to be integrated over. Thus fulfilling detailed balance with such a move is not tractable. However, the notion of superdetailed balance [144] can be fulfilled with such a move, which constitutes a detailed balance of the conditional path probabilities:

$$p_{\text{acc}} = \min\left\{1, \exp(-\beta \Delta U + \Delta S)\right\}, \tag{4.25}$$

where $\Delta S = \log p(\mathbf{x}_N \to \cdots \to \mathbf{x}_0; \tilde{\Lambda}) - \log p(\mathbf{x}_0 \to \cdots \to \mathbf{x}_N; \Lambda)$ denotes the difference in logarithmic path probability, between the forward and backward path under the propagation kernels Λ. A connection between the path probabilities and the dissipated heat is given by the Crooks fluctuation theorem [145]. The perturbation kernels in this method usually drive the system along a pre-defined reaction coordinate. An exact perturbation protocol has therein to be crafted by hand and the final acceptance ratio crucially depends on it.

This method has recently been applied in the machine learning community as a way to implement a more complex family of transformations called stochastic normalizing flows [146], which combine deterministic proposals with stochastic dynamics in order to generate a sample.

4.4.3. Biased potential methods

Several methods make use of a small set of coordinates, called collective coordinates, which are considered relevant for the dynamics of interest of the molecules. The system of interest is then biased along these coordinates, to either drive the system systematically or allow for faster transitions along them. Samples from the unbiased Boltzmann distribution can then be generated by a reweighting step Section 4.1. Metadynamics [147] biases the potential on the collective variables in an iterative fashion, where the bias potential is increased in areas where the system resides a long time, thus, pushing the system out of metastable states. The bias potentials are chosen to be Gaussian-shaped with a user-defined width. Recent developments suggest the usage of deep learning to find an optimal bias potential [148]. Umbrella sampling [149] runs several sampling iterations with bias potentials, so-called *umbrellas*, placed along a single collective coordinate, the reaction coordinate, and thus, pulls the system from one end to the other. All of these methods require prior knowledge of coordinates which are relevant for the dynamics of the molecule of interest.

4.4.4. Hamiltonian Monte Carlo

The autocorrelation of a Markov chain can further be reduced by designing better proposals. To this end, care has to be taken such that the benefit of generating better steps does not outweigh their computational cost. One important method for this is Hamiltonian Monte Carlo (HMC) [150]. The HMC method augments the configuration space with auxiliary variables $\mathbf{x} \to (\mathbf{x}, \mathbf{p})$, which are distributed as $p_A(\mathbf{p})$, propagates the sample in augmented space by some dynamics $(\mathbf{x}', \mathbf{p}') = \mathcal{D}(\mathbf{x}, \mathbf{p})$ and computes a proposal in the augmented space $\mathbf{x} \to (\mathbf{x}, \mathbf{p}) \to (\mathbf{x}', \mathbf{p}') \to \mathbf{x}'$. Defining the *Hamiltonian* in augmented space $\mathcal{H}(\mathbf{x}, \mathbf{p}) = U(\mathbf{x}) - k_\mathrm{B} T \log p_A(\mathbf{p})$, it can be shown that accepting a step generated in this fashion with probability

$$p_\mathrm{acc}\left((\mathbf{x}, \mathbf{p}) \to (\mathbf{x}', \mathbf{p}')\right) = \min\left\{1, \exp\left[-\beta \Delta \mathcal{H}\right] \left|\det J_\mathcal{D}(\mathbf{x}, \mathbf{p})\right|\right\}, \tag{4.26}$$

with $\Delta \mathcal{H} = \mathcal{H}(\mathbf{x}, \mathbf{p}) - \mathcal{H}(\mathbf{x}', \mathbf{p}')$, and $J_\mathcal{D}$ the Jacobian of the dynamics $\mathcal{D}$, fulfills detailed balance in configuration space. The augmented distribution p_A is commonly chosen to be a Maxwell Boltzmann distribution. With this choice, $\mathcal{H}$ takes the form of a classical Hamiltonian, which lead to naming this method *Hamiltonian Monte Carlo.* Furthermore, choosing a dynamics that preserves $\mathcal{H}$ resembles the classical Hamiltonian dynamics of the system. Integrating Hamilton's equations

$$\partial_t \mathbf{x} = \partial_{\mathbf{p}} \mathcal{H}, \quad \partial_t \mathbf{p} = -\partial_{\mathbf{q}} \mathcal{H}, \tag{4.27}$$

$$(\mathbf{x}', \mathbf{p}') = \mathcal{D}(\mathbf{x}, \mathbf{p}) = \int_0^{\Delta t} (\dot{\mathbf{x}}, \dot{\mathbf{p}}) \mathrm{d}t, \tag{4.28}$$

for some time Δt preserves the Hamiltonian $\Delta \mathcal{H} = 0$ and furthermore, according to Liouville's theorem, also preserves the volume $|\det J_{\mathcal{D}}| = 1$. Thus, such a proposed move has acceptance probability $p_{\text{acc}}(\mathbf{x} \to \mathbf{x}') = 1$. In practice, due to errors arising from numerical integration, the acceptance probability will eventually be lower. As numerical errors accumulate with an increase of the integration time Δt, a decrease of the acceptance probability is observed with an increase of the integration time. The correct choice of the integrator and integration time is important, and a trade-off between precision and computational cost has to be considered when selecting them. A common choice for the numerical integration method in this context is the leapfrog integrator [151].

However, the dynamics $\mathcal{D}$ that generates a new proposal does not need to have any physical meaning and in many situations, it is favorable to choose a function that generates far-reaching moves that have a lower acceptance rather than short-ranged moves with high acceptance, as these moves will drastically reduce the autocorrelation time. Methods such as *smart darting Monte Carlo* make use of this, which is introduced in the following section.

4.4.5. Smart darting Monte Carlo

The convergence rate of a MCMC sampler is limited by the transition rate over the largest free energy barriers in the system which is a function of the barrier height as discussed above. This limitation of the convergence rate is especially evident in situations where states are disconnected and finding e.g. a reaction coordinate that could be used to drive the system is challenging. Different approaches have been developed for these scenarios: Smart Darting Monte Carlo [152, 153] introduces kernels that propose far-ranging moves that allow for direct transitions between regions in phase space of high population. This approach completely ignores energy barriers between the metastable states. To this end, free energy minima $\mathbf{m}_i$ are identified in a first exploration step. Then small spheres are defined around these minima. In the sampling run, the sampler proposes direct transitions between these spheres by translation. Say the system is currently in sphere i, then a proposal to a randomly selected sphere j is generated by $\mathbf{y} = \mathbf{x} + \mathbf{d}_{ij}$, where $\mathbf{d}_{ij} = \mathbf{m}_j - \mathbf{m}_i$ denotes the displacement from center i to j. In order for detailed balance to hold, the resulting $\mathbf{y}$ then needs to be in sphere j. These moves are alternated with local exploration moves, such as standard MCMC. In high dimensions, however, the fraction of the spheres in the total volume becomes vanishingly small, and therefore finding a sphere by random exploration becomes unlikely.

This problem is partially circumvented in ConfJump [154] by mapping the current state to the closest energy minimum and attempting a long-range move by translation to another energy minimum. The method focuses on the sampling of molecules and these long-range moves are proposed by transforming the current configuration to internal coordinates and performing a translation there. Again is has to be checked that the probability of this proposal is balanced to ensure convergence.

In Chapter 5 a new sampling scheme along the lines of these methods is introduced. It uses general invertible functions in order to propose the new, far-reaching moves. These are parametrized by neural networks and trained on samples of the individual metastable states. This has significant advantages over using a simple translation between the metastable states, as the invertible function can learn the local free energy landscape around the minima and propose moves with higher acceptance probability accordingly.

4.5. Deep learning approaches for MC methods

With the current rise of attention that the field of deep learning has seen, new methods have been developed which use deep neural networks (DNNs) to construct MCMC moves. Recently proposed methods use DNNs for approximating the target Boltzmann distribution [155] or projecting onto high probability manifolds [156]. Two recent approaches use reversible network architectures to improve HMC: A-NICE-MC [157] proposes steps by applying volume-preserving transformations that operate on the augmented configuration space. Another approach [158] augments the leapfrog algorithm with DNNs and thereby altering the classical path of the system. Both are trained for sampling efficiency in an unsupervised fashion and therefore rely on random exploration of configuration space in order to find metastable states.

In Boltzmann generators [159, 160] reversible DNNs are used to draw statistically independent equilibrium samples of condensed matter systems and protein molecules in an importance sampling framework. This method uses deep learning for sampling a target distribution by combining an exact probability generator such as a normalizing Flow [161, 162] with reweighting (Section 4.1).

These new methods are facilitated by the recent development of new neural network architectures that enable for inexpensive inversion and computation of Jacobians, named *flows*

4.5.1. Flows

Both, MCMC and reweighting rely on the computation of a likelihood associated with a proposed move or sample $\mathbf{y}$, in order to fulfill detailed balance or compute weights respectively. If such a move or sample is generated by a function $\mathbf{y} = f(\mathbf{x})$, this likelihood can be computed by the change of variable formula, which states that for an invertible function $f : \mathbb{R}^d \to \mathbb{R}^d$, and a random variable X with distribution $p_X(\mathbf{x})$, the density of the transformed variable $\mathbf{y} = f(\mathbf{x})$ is given by

$$p_Y(\mathbf{y}) = p_X(\mathbf{x}) \left|\det J_f(\mathbf{x})\right|^{-1} . \tag{4.29}$$

Utilizing a neural network to parametrize the function f therefore requires two properties:

- invertibility: the inverse of a function f^{-1} needs to exist and its computation needs to be feasible
- the Jacobian determinant $\det J_f(\mathbf{x}) = \left|\frac{\partial f_i(\mathbf{x})}{\partial x_j}\right|$ needs to be easily computable.

While it is usually possible to compute the Jacobian of a deep neural network (DNN), only a few of the DNN-architectures are invertible. Furthermore, the computation of the Jacobian of a general function can become extremely computationally expensive. A class of neural network architectures that fulfill these conditions are called *flows*. Recently many new types of flows have been proposed [146, 157, 162–164]. A flow usually consists of a sequence of transformations f_i, with the total transformation given by the composition $\mathbf{y} = f_N \circ \cdots \circ f_1(\mathbf{x})$. Furthermore, let $\mathbf{x}^{(i)} = f_i \circ \cdots \circ f_1(\mathbf{x})$ denote the output of the i-th transformation. A wide variety uses so-called *coupling layers* to ensure invertibility and easy computation of the Jacobian. In coupling layers, the coordinates $\mathbf{x}^{(i)} \in \mathbb{R}^d$ are split into two channels $\mathbf{x}_1^{(i)} \in \mathbb{R}^{d_1}, \mathbf{x}_2^{(i)} \in \mathbb{R}^{d_2}$, with $d_1 + d_2 = d$. The update of one set of coordinates is then given by an invertible parametrized scalar function $h_\Theta(\cdot)$ whose parameters $\Theta = \Theta(\mathbf{x}_2^{(i)})$ are a function of the second set of coordinates. This function is applied element-wise on the first set of coordinated, while other set of coordinates is simply kept constant

$$\begin{aligned} \mathbf{x}_{1,j}^{(i+1)} &= h_{\Theta\left(\mathbf{x}_2^{(i)}\right)}\left(\mathbf{x}_{1,j}^{(i)}\right) \\ \mathbf{x}_2^{(i+1)} &= \mathbf{x}_2^{(i)}, \end{aligned} \tag{4.30}$$

which can be inverted by applying the inverse h_Θ^{-1} in the first line. Furthermore, it results in a triangular Jacobian, whose determinant is easily computed as the product over its diagonal entries

$$\det J_{f_i} = \prod_j \partial h_{\Theta(\mathbf{x}_2^{(i)})}(\mathbf{x}_{1,j}^{(i)}) / \partial \mathbf{x}_{1,j}^{(i)}. \tag{4.31}$$

In the following step, the coordinate sets are swapped, and $\mathbf{x}_{2,j}^{(i+2)} = h_{\Theta\left(\mathbf{x}_1^{(i+1)}\right)}(\mathbf{x}_{2,j}^{(i+1)})$, while $\mathbf{x}_1^{(i+2)} = \mathbf{x}_1^{(i+1)}$. Stacking multiple of these building blocks enables for complicated transformations, as information can flow between the different sets of coordinates. The Jacobian of the full transformation f is then given by the sum of the Jacobians of the partial transformations $\det J_f(\mathbf{x}) = \prod_i \det J_{f_i}$

The most simple transformation defined by this procedure is the *nonlinear independent component estimation* (NICE) [157], where $h_{\Theta(y)}(x) = x + \Theta(y)$. Here, Θ can be an arbitrary function. Thus, a *multilayer perceptron* (MLP), as a general function estimator, is a common choice for Θ. Further generalizations use an affine transformation, which is denoted as *real-valued non-volume preserving transformation* (RNVP) [162], or general invertible functions parametrized by splines [165]. The work presented in this thesis mostly makes use of the RNVP structure [162]. For this flow structure, the updates take

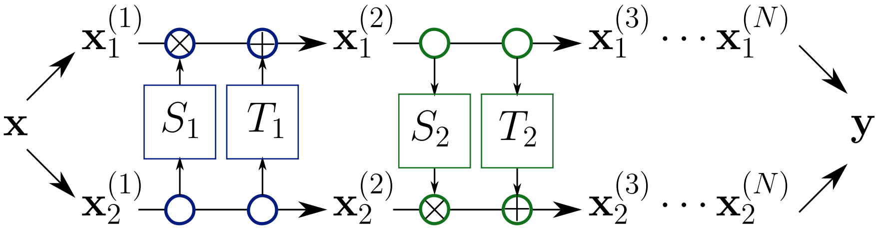

Figure 4.1.: Architecture of the RNVP networks. The input configuration $\mathbf{x}$ is separated into two channels $\mathbf{x}_1^{(1)}, \mathbf{x}_2^{(1)}$. At the first iteration only subsets $\mathbf{x}_2^{(1)}$ is fed into the neural networks S_1, T_1 allowing for nonlinear transformations. The outputs of the neural networks S_1, T_1 are respectively multiplied and added to the other subset $\mathbf{x}_1^{(1)}$ to arrive at $\mathbf{x}_1^{(1)}$, while $\mathbf{x}_2^{(1)}$ is kept constant. In the following iteration, the subsets are swapped and $\mathbf{x}_2^{(2)}$ gets transformed by functions depending only on $\mathbf{x}_1^{(1)}$. Stacking many of these iterations allows for a complex transformation that can easily be inverted.

the form

$$\begin{bmatrix} \mathbf{x}_1^{(i+1)} \\ \mathbf{x}_2^{(i+1)} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1^{(i)} \odot \exp\left[s^{(i)}\left(\mathbf{x}_2^{(i)}\right)\right] + T^{(i)}\left(\mathbf{x}_2^{(i)}\right) \\ \mathbf{x}_2^{(i)} \end{bmatrix}, \tag{4.32}$$

$$\begin{bmatrix} \mathbf{x}_1^{(i+2)} \\ \mathbf{x}_2^{(i+2)} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1^{(i+1)} \\ \mathbf{x}_2^{(i+1)} \odot \exp\left[s^{(i+1)}\left(\mathbf{x}_1^{(i+1)}\right)\right] + T^{(i+1)}\left(\mathbf{x}_1^{(i+1)}\right) \end{bmatrix}, \tag{4.33}$$

where $s_i, T_i : \mathbb{R}^{d_{1/2}} \to \mathbb{R}^{d_{2/1}}$ are DNNs and $\odot$ denotes the element-wise multiplication. Due to the simple form of the transformation, the log Jacobian can easily be computed as $\log|\det J_{f_i}| = \sum_j s_j^{(i)}$. A diagrammatic description of this transformation is shown in Fig. 4.1. As denoted later on in Chapter 6, it is in some cases favorable to keep transformations volume preserving, i.e. with $\log|\det J| = 0$. This can be achieved straightforwardly by subtracting the mean from the scaling parameters $s^{(i)} = \tilde{s}^{(i)} - \sum_j \tilde{s}_j^{(i)}/d_s$, where $\tilde{s}^{(i)}$ is the output of the MLP and d_s its dimensionality. This structure allows the flow to exchange volume between the coordinates while preserving the total volume under the transformation.

These flows can also be used in a static Monte Carlo framework in which they serve a proposal distribution for reweighting.

4.5.2. Normalizing flows and Boltzmann Generators

A class of methods [161, 166–168] utilizes flows in combination with the change of variable formula for densities [Eq. (4.29)] to transform samples $\mathbf{z}$ from a distribution $p_Z(\mathbf{z})$ to the target distribution $p_X(\mathbf{x})$ by a transformation $\mathbf{x} = g(\mathbf{z})$. The distribution p_Z is usually denoted as latent distribution and is commonly chosen such that i.i.d. samples can easily be generated from it. A common choice is a standard Gaussian distribution $p_Z = \mathcal{N}(0, 1)$, thus *normalizing* the configuration space distribution p_X by the transformation g^{-1}. Generating i.i.d. samples from $\mathbf{x} \sim p_X$ then becomes as easy as sampling $\mathbf{z} \sim p_Z(\mathbf{z})$ and computing $\mathbf{x} = g(\mathbf{z})$. However, finding the correct transformation g is a difficult task.

This is exactly what flows were designed for: approximate the function g by a flow f, which gives rise to a distribution $q_X(\mathbf{x}) = p_Z(f^{-1}(\mathbf{x}))|\det J_{g^{-1}}(\mathbf{x})|$. Further train f, such that $q_X \approx p_X$ by minimizing some (statistical) divergence $D(p_X||q_X)$. This method is called *normalizing flows* as it utilizes a flow to *normalize* a distribution. A recent development of normalizing flows in the field of statistical physics is Boltzmann Generators (BG) [146, 159], which aims at generating samples from the Boltzmann distribution $p_X = \mu$. To this end, some samples from the Boltzmann distribution, which follow some empirical sampling distribution $\varrho(\mathbf{x})$, are usually required for training. BG minimize a combination of the forward and backward Kulback-Leibler divergence $D_{KL}(q||p) = \int dx p(\mathbf{x}) \log p(\mathbf{x})/q(\mathbf{x})$ between $q(\mathbf{x})$ and $\mu(\mathbf{x})$ [159, 169]:

$$D_{KL}(\varrho_X||q_X) = \mathbb{E}_{\mathbf{x}\sim\varrho_X}\left[\frac{1}{2}\left\|f^{-1}(\mathbf{x})\right\|^2 - \log|\det J_{f^{-1}}(\mathbf{x})|\right] + \text{const.} \tag{4.34}$$

$$D_{KL}(q_X||\mu) = \mathbb{E}_{\mathbf{z}\sim p_Z}\left[\beta U(f(\mathbf{z})) - \log|\det J_f(\mathbf{z})|\right] + \text{const.}, \tag{4.35}$$

where the term Eq. (4.34) is commonly denoted as *negative log likelihood* $\mathrm{nll}(\varrho_X||q_X)$. Training a BG usually consists of minimizing a convex combination of Eqs. (4.34) and (4.35).

As recovering the exact Boltzmann distribution is extremely hard in high dimensions, a final reweighting step can be taken in order to obtain unbiased ensemble averages. As the flow provides an exact likelihood method, unnormalized weights can easily be computed in terms of importance sampling and using Eq. (4.4), these are found as

$$w = \frac{\exp\{-\beta U(f(\mathbf{z}))\}}{p_Z(\mathbf{z})}|\det J_f(\mathbf{z})|. \tag{4.36}$$

With these weights, the bias introduced by the deviation between the output distribution q_Z and the target distribution μ can be compensated, ensuring the convergence of observables.

Twofold contribution in the field of enhanced sampling methods has been made in the course of this thesis:

With Neural Mode Jump Monte Carlo, a new method has been developed that allows for MCMC proposals that directly transition between metastable states. To this end flows are trained to propose these transitions directly. This allows for circumventing the problem of slow convergence across free energy barriers. The method is introduced in the following chapter.

Furthermore, an extension in the BG framework has been developed that allows for generating output distributions that exhibit the same Arrhenius-like scaling behavior as the Boltzmann distribution. This extension allows for training at higher temperatures where high-quality training samples are more easily accessible. Following this training the BG can be used to sample at a lower temperature for more efficient sampling or at multiple temperatures to investigate temperature-dependent observables. The method can further be utilized in a PT framework to generate a highly efficient sampler. It is introduced in Chapter 6.

5. Neural Mode Jump Monte Carlo

The results of this chapter have been published in:

> Luigi Sbailò, Manuel Dibak, and Frank Noé. "Neural mode jump Monte Carlo". In: *The Journal of Chemical Physics* 154.7 (2021), p. 074101. DOI: `10.1063/5.0032346`

Parts of text and figures have been adopted unchanged in this chapter. Reproduced from *The Journal of Chemical Physics* "Neural mode jump Monte Carlo", Sbailò et al., 2021, with the permission AIP Publishing.
Luigi Sbailò (LS) and Manuel Dibak (MD) contributed equally to this work. In particular the contribution was as follows: all authors conceived the project, LS and MD worked out the theory, performed experiments and visualized the data. MD implemented the training and sampling algorithms. All contributors wrote the manuscript.

Summary

Markov chain Monte Carlo methods are a powerful tool for computing equilibrium observables in complex high-dimensional systems. In molecular systems these methods often exhibit a slow convergence, as transitions over large energy barriers drastically slow down the generation of uncorrelated samples. Methods such as smart darting Monte Carlo partially address this issue, by directly proposing moves between energy wells, thus completely ignoring the barriers in-between. However, these often face problems, as they can not account for the energy surface around these minima, which can be very rigid, leading to this method having a very poor acceptance probability. In this chapter, a novel method is proposed, which increases convergence in systems composed of many metastable states. This method aims to connect metastable regions directly using invertible functions that generate long-ranging proposal moves in the MCMC framework. These functions are parametrized by *flows*, a type of neural network architecture suitable for this application. During training, the flow optimizes the acceptance probability of large jumps between minima in the free energy landscape. A comprehensive theory, as well as a training scheme for the network, are provided and the method is demonstrated on example systems.

5.1. Introduction

The generation of long-range moves as proposed in the schemes introduced in Section 4.4.5 is challenging when the energy landscape is rough, since the potential energy surface in the region surrounding local minima can drastically change among the different minima. In this case, using trivial translation as long-range moves would most likely cause large energy differences, causing proposed moves likely to be rejected. In order to generate moves with high acceptance probability, trial moves should pair points in either metastable state in a way that minimizes the energy difference, thus constructing a bijection between configurations in either well.

Constructing such bijection manually would require detailed knowledge of the system and is practically impossible in multi-dimensional systems. Recent advances in the field of machine learning have permitted to deal with problems that were not solvable with a sole human understanding, and, more specifically, *flows* (Section 4.5.1) are an ideal tool to provide such a bijection.

In this chapter, neural mode jump Monte Carlo (Neural MJMC) is presented, a novel method to efficiently sample the equilibrium distribution of complex many-body systems with unbiased Markov chains. In this scheme, reversible neural networks are trained to propose long-range *neural* moves that directly connect different metastable states, as neural moves do not depend on the physical path connecting start and endpoint. The method requires prior knowledge of the position of the metastable states in configuration space, which could be obtained from e.g. x-ray scattering or NMR experiments from which starting points in different conformations can be generated [171]. Possible applications range from proteins with multiple metastable states (e.g. Kinases) to solid-state systems with multiple phases, where configurations in either phase can easily be generated, but observing the transition is rare. Local displacements, e.g. by Gaussian displacements, and neural moves are randomly alternated in a combined scheme to accelerate the convergence rate of Markov chains. Configurations from different metastable states are used to train the networks, which are optimized to produce high acceptance probability moves. The local moves ensure the ergodicity of the scheme, while neural moves accelerate convergence to equilibrium, realizing an accurate and deep exploration of the configuration space.

In Neural MJMC, the proposal probability is split into two steps: firstly, a proposal density is selected from a pre-defined list of proposal densities on the current state $\mathbf{x}$, then a new state $\mathbf{y}$ is drawn from the selected proposal density. Proposal densities are distinguished between local proposals and neural proposals, where local proposals generate local moves, e.g. through random displacement with Gaussian increments, and neural proposals connect different metastable states with global moves.

5.2. Theory

Assume that the configuration space Γ is decomposed into a number of non-overlapping subsets called cores $\{\Omega_\alpha\}_{\alpha\leqslant N} \subset \Gamma$, with $\cup_\alpha \Omega_\alpha = \Gamma$, each representing one of the N metastable states. The neural proposal kernel $K_{\alpha\beta}$ is defined as the density that proposes transitions from the core Ω_α to the core Ω_β. Assuming the system is in state $\mathbf{x} \in \Omega_\alpha$, the probability to select the neural proposal $K_{\alpha\beta}$ is denoted as $p_{\alpha\beta}(\mathbf{x})$. Once $K_{\alpha\beta}$ has been selected, a state $\mathbf{y} \in \Omega_\beta$ is drawn from the selection probability $p^{\alpha\beta}_{\text{prop}}(\mathbf{x} \to \mathbf{y})$.

A neural proposal $K_{\alpha\beta}$ can only be selected within the core Ω_α and with constant probability $p_{\alpha\beta}(\mathbf{x}) = p_{\alpha\beta}\chi_{\Omega_\alpha}(\mathbf{x})$, where $\chi_\Omega(\mathbf{x})$ denotes the characteristic function, Eq. (1.14). Assuming that each pair of states (α, β) is only connected by one neural proposal $K_{\alpha\beta}$ and that there exists an inverse proposal $K_{\beta\alpha}$ connecting β with α, a proposed move starting in Ω_α with selected neural proposal $K_{\alpha\beta}$ fulfills detailed balance if it is accepted with probability

$$p_{\text{acc}}^{\alpha\beta}(\mathbf{x} \to \mathbf{y}) = \min\left\{1, \frac{\mu(\mathbf{y})\, p_{\beta\alpha}\, p_{\text{prop}}^{\beta\alpha}(\mathbf{y} \to \mathbf{x})}{\mu(\mathbf{x})\, p_{\alpha\beta}\, p_{\text{prop}}^{\alpha\beta}(\mathbf{x} \to \mathbf{y})}\right\}, \tag{5.1}$$

which follows from the Metropoli-Hastings criterion Eq. (4.13). The neural proposal $K_{\alpha\beta}$ and its inverse $K_{\beta\alpha}$ connecting the cores Ω_α and Ω_β are furthermore parametrized as a bijective function $f_{\alpha\beta}(\cdot)$ pairing the states defined in the two cores, i.e. $\mathbf{y} = f_{\alpha\beta}(\mathbf{x})$, $f_{\alpha\beta}^{-1}(\mathbf{y}) = \mathbf{x}$, $\forall \mathbf{x} \in \Omega_\alpha$, where $\mathbf{x} \in \Omega_\alpha$, $\mathbf{y} \in \Omega_\beta$. Thus for each pair of different cores $(\Omega_\alpha, \Omega_\beta)$ a bijective function $f_{\alpha\beta}(\cdot)$ is defined. The probability distribution of neural proposals is then represented with Dirac delta distributions and the acceptance specifies to

$$p_{\text{acc}}^{\alpha\beta}(\mathbf{x} \to \mathbf{y}) = \min\left\{1, \frac{\mu(\mathbf{y})\, p_{\alpha\beta}\, \delta(\mathbf{x} - f_{\alpha\beta}^{-1}(\mathbf{y}))}{\mu(\mathbf{x})\, p_{\beta\alpha}\, \delta(\mathbf{y} - f_{\alpha\beta}(\mathbf{x}))}\right\}. \tag{5.2}$$

Using the change of variable formula in the Dirac distribution $\delta(\mathbf{x} - f_{\alpha\beta}^{-1}(\mathbf{y})) = \left|\det J_{f_{\alpha\beta}}(\mathbf{x})\right| \delta(\mathbf{y} - f_{\alpha\beta}(\mathbf{x}))$, with the Jacobian $J_{f_{\alpha\beta}}(\mathbf{x})$ of the function $f_{\alpha\beta}$, the acceptance probability for neural moves can be simplified to

$$p_{\text{acc}}^{\alpha\beta}(\mathbf{x} \to \mathbf{y}) = \min\left\{1, \frac{\mu(\mathbf{y})\, p_{\alpha\beta}}{\mu(\mathbf{x})\, p_{\beta\alpha}} \left|\det J_{f_{\alpha\beta}}(\mathbf{x})\right|\right\}. \tag{5.3}$$

In case the local proposal ($\alpha = \beta$) is selected, the inverse move is only possible with another local proposal $K_{\alpha\alpha}$. Note that a local move may leave the current core and the proposal probability for the inverse move might change. Thus the acceptance probability for a local move reduces to

$$p_{\text{acc}}^{\alpha\alpha}(\mathbf{x} \to \mathbf{y}) = \min\left\{1, \frac{\mu(\mathbf{y}) \sum_\beta \chi_{\Omega_\beta} p_{\beta\beta}}{\mu(\mathbf{x}) p_{\alpha\alpha}}\right\}. \tag{5.4}$$

In order to ensure ergodicity, there needs to be a finite probability of selecting the local proposal in all cores. Figure 5.1 gives a graphical description of the sampling scheme and Algorithm 4 summarizes the Neural MJMC sampling scheme.

5.3. Optimal proposal density

To achieve fast decorrelation of the Markov chain, the neural proposal functions $f_{\alpha\beta}$ should maximize the acceptance in both directions. This is quantified by maximizing the

Algorithm 4: Neural MJMC sampling scheme

```
input  : $l_s = [\,]$ : empty list for samples
         $\{p_{\alpha\beta}\}$ : proposal selection probabilities
         $\{f_{\alpha\beta}\}$: proposal functions
         $\mathbf{x} \leftarrow \mathbf{x}_0$ : starting point of sampling
         $N_{\text{iterations}}$ : number of generated samples
         $\sigma_{\text{local}}$ : standard deviation of local moves
while $i \leqslant N_{\text{iterations}}$ :
    draw proposal density $K_{\alpha\beta}$ from $\{p_{\alpha\beta}\}$
    if $\alpha = \beta$ :                                   // propose local move
        $\mathbf{w} \leftarrow$ sample from $\mathcal{N}(0, \mathbb{1})$
        $\mathbf{y} \leftarrow \mathbf{x} + \mathbf{w} \cdot \sigma_{\text{local}}$
        $p_{\text{acc}} \leftarrow p_{\text{acc}}^{\alpha\alpha}(\mathbf{x} \to \mathbf{y})$ (Eq. 5.4)
    else:                                                  // propose neural move
        $\mathbf{y} = f_{\alpha\beta}(\mathbf{x})$
        $p_{\text{acc}} \leftarrow p_{\text{acc}}^{\alpha\beta}(\mathbf{x} \to \mathbf{y})$ (Eq. 5.3)
    if $r \sim \mathcal{U}(0,1) < p_{\text{acc}}$ :
        $\mathbf{x} \leftarrow \mathbf{y}$
    $l_{\text{s}}$.append $(\mathbf{x})$
    $i \leftarrow i + 1$
output : list of samples $l_{\text{s}}$
```

expected log probability that the moves proposed by $f_{\alpha\beta}$ are accepted in both directions

$$\begin{aligned}\max_{f_{\alpha\beta}} \log \mathbb{E}_{\mathbf{x}\sim\Omega_\alpha}\left[p_{\text{acc}}(\mathbf{x} \to \mathbf{y}) p_{\text{acc}}(\mathbf{y} \to \mathbf{x})\right] &\geqslant \max \mathbb{E}\left\{\log\left[p_{\text{acc}}(\mathbf{x} \to \mathbf{y}) p_{\text{pacc}}(\mathbf{y} \to \mathbf{x})\right]\right\} \\ = \max_{f_{\alpha\beta}} \mathbb{E}\left[\min\left(0, \log g\right) + \min(0, -\log g)\right] &= \max_{f_{\alpha\beta}} \mathbb{E}\left[\min(\log g, -\log g)\right] \\ &= \max_{f_{\alpha\beta}} \mathbb{E}\left[-\left|\log g\right|\right], \quad (5.5)\end{aligned}$$

where $g := \frac{\mu(\mathbf{y})\, p_{\alpha\beta}}{\mu(\mathbf{x})\, p_{\beta\alpha}} \left|\det J_{f_{\alpha\beta}}(\mathbf{x})\right|$ and using Jensen's inequality [172] in the first line. Using the stationary distribution in the canonical ensemble $\mu(\mathbf{x}) \propto \exp(-\beta U(\mathbf{x}))$, with the thermal energy $\beta^{-1} = k_{\text{B}}T$, the potential energy $U(\mathbf{x})$ of the system under consideration and assuming that $f_{\alpha\beta}$ is a bijection between the cores (α, β), the equation above can be rewritten to find

$$\min_{f_{\alpha\beta}} \mathbb{E}\left[\beta \left|\Delta U_{\alpha\beta}(\mathbf{x}) + k_B T \log\left|\det J_{f_{\alpha\beta}}(x)\right| + \Delta R_{\alpha\beta}\right|\right], \tag{5.6}$$

with the potential difference $\Delta U_{\alpha\beta}(\mathbf{x}) := U(\mathbf{x}) - U(f_{\alpha\beta}(\mathbf{x}))$ and the log selection ratio $\Delta R_{\alpha\beta} := -k_B T \log p_{\alpha\beta}/p_{\beta\alpha}$. Note that the term inside the modulus is equivalent to the Kulback-Leibler divergence between the transformed distribution $f_{\alpha\beta}(\Omega_\alpha)$ and the target distribution Ω_β as found in [159].

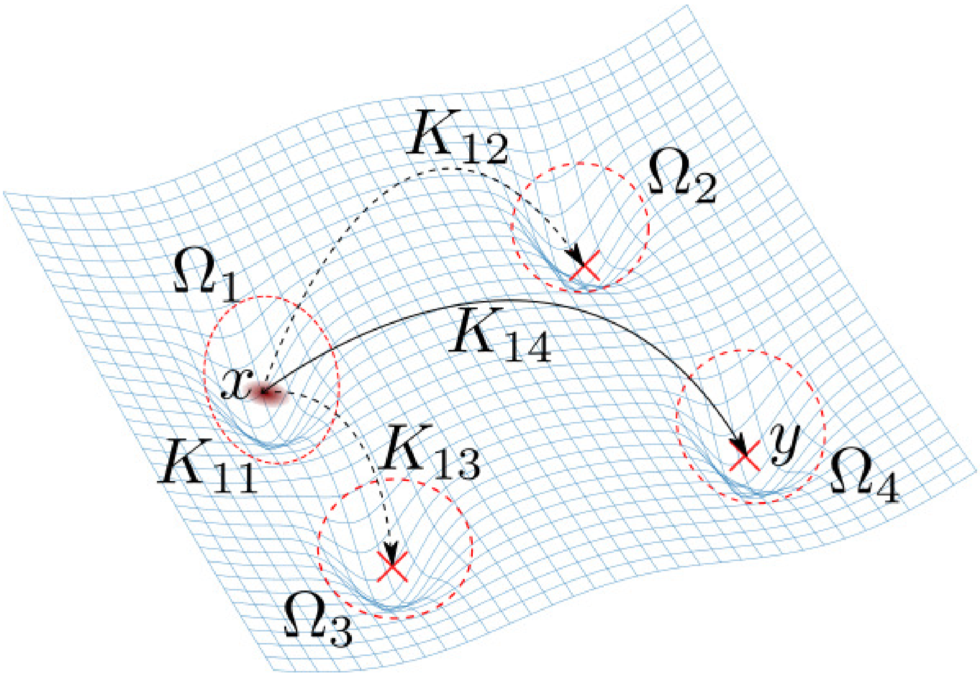

Figure 5.1.: Schematic figure of Neural MJMC scheme. Given configuration $\mathbf{x}$ in core Ω_1 there are three neural and one local proposals available, as denoted by arrows. One of these is selected and a new state $\mathbf{y}$ is proposed. Reprinted from *The Journal of Chemical Physics* "Neural mode jump Monte Carlo", Sbailò, Dibak, and Noé, 2021, with the permission of AIP Publishing.

This result can be interpreted in a physically meaningful manner by applying the triangular inequality $\mathbb{E}\left[-\left|\log g\right|\right] \geqslant -\left|\mathbb{E}\left[\log g\right]\right|$, identifying $\Delta S = -k_B \mathbb{E}\left[\log\left|\det J_{f_{\alpha\beta}}(x)\right|\right]$ as the change of differential entropy (see Appendix B for details) and $\Delta U = \mathbb{E}\left[\Delta U_{\alpha\beta}(\mathbf{x})\right]$ as the change of internal energy under the transformation $f_{\alpha\beta}(\mathbf{x})$. It is observed that the expected log acceptance is lower bound by the absolute change in free energy $\Delta F = \Delta U - T\Delta S$ under the transformation $f_{ij}(\cdot)$ divided by thermal energy

$$\mathbb{E}\left\{\log\left[p_{\mathrm{acc}}(\mathbf{x}\to\mathbf{y})p_{\mathrm{acc}}(\mathbf{y}\to\mathbf{x})\right]\right\} \geqslant -\beta\left|\Delta F + \Delta R_{\alpha\beta}\right|. \tag{5.7}$$

This results shows that the proposal selection ratio R_{ij} can be used to tune the bidirectional acceptance probability.

5.4. Neural network architecture and training

As neural moves are defined with invertible functions, it is necessary to choose neural networks that are invertible. Furthermore, the Jacobian of the transformation needs to be computable at a reasonable cost. A class of neural network architectures designed for this task are are *flows* [157, 161–163, 173]. Section 4.5.1 gives an introduction to flows and introduces some common variants thereof.

In order to ensure that the output of the network $f_{\alpha\beta}(\cdot)$ is in the correct well, a harmonic bias potential centered in the target core is added during training

$$U_{\mathrm{bias}}(\mathbf{x}) = \begin{cases} k\left(\mathbf{x}-\mathbf{x}_\alpha\right)^2 & \mathbf{x} \in \Omega_\alpha \\ k\left(\mathbf{x}-\mathbf{x}_\beta\right)^2 & \mathbf{x} \in \Omega_\beta \end{cases}, \tag{5.8}$$

where $\mathbf{x}_\alpha$ is the reference configuration in core α and the spring constant k is a hyperparameter, resulting in the biased system $\tilde{U}(\mathbf{x}) = U(\mathbf{x}) + U_{\mathrm{bias}}(\mathbf{x})$ used during training. The network is trained in several stages, gradually lowering the strength of the bias potential. To find the reference configurations $\mathbf{x}_\alpha$ k-means clustering [174] is run on samples generated from local MCMC sampling in either well. Training sets of both of the wells are generated for a set of gradually decreasing bias strengths $\{k_i\}_{i\leqslant N_k}$. After convergence of the training at k_i, the training set is exchanged and training is restarted with $k_{i+1} \leqslant k_i$. This allows for a slowly expanding training set, which enables the network to learn how to generate meaningful moves on a gradually more complex set of training data. The loss that is to be minimized during training is given by the negative bi-directional acceptance Eq. (5.6):

$$\mathcal{L}_{\mathrm{acc}} = \mathbb{E}_{\mathbf{x}\sim\Omega_\alpha}\left\{\left[\Delta\tilde{U}_{\alpha\beta}(\mathbf{x}) + k_B T \log\left|\det J_{f_{\alpha\beta}}(\mathbf{x})\right|\right]^2\right\}, \tag{5.9}$$

where the square of the norm is used to penalize high energies. Training is performed in the forward and backward direction and the same loss applies to samples from core Ω_β with exchanged labels $\alpha \leftrightarrow \beta$.

5.5. Numerical experiments

Neural MJMC is demonstrated on two examples: a two-dimensional potential landscape with three minima and a system consisting of two dimer particles which are suspended in a bath of repulsive particles. The functional form of the potentials and detailed training parameters are given in Appendix A.2.

As a good compromise between computational cost and expressiveness, RNVP [162] are used as network structure for these examples. See Section 4.5.1 for a detailed description of this flow architecture.

5.5.1. Gaussian triple well

As an example for a system with multiple states, Neural MJMC is demonstrated on a two-dimensional potential landscape consisting of 3 Gaussian shaped wells. The 3 cores are defined by a Voronoi tessellation [175] of the configurations space for which the minima of the Gaussians are used as centers. The functional form and its parameters is given in Appendix A.2.1.

Each flow $f_{\alpha\beta}$ is trained independently on configurations sampled from the minima.

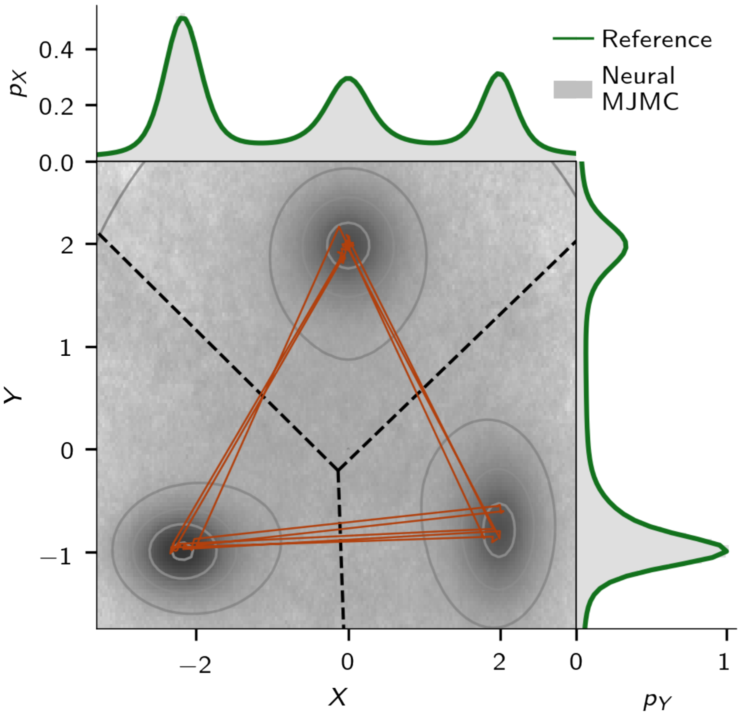

Figure 5.2.: Two-dimensional histogram (**center**) of samples from the 2D Gaussian triple well potential generated by Neural MJMC with a short section of the Markov chain (orange solid line) and marginal distributions p_X (**top**) and p_Y (**right**). The black dashed line depicts the border between the states which are defined by a Voronoi tessellation. Convergence to the correct Boltzmann distribution can be observed from the histograms of the marginal distributions, where the green line is the reference solution from numerical integration of the system's Boltzmann distribution. Reprinted from *The Journal of Chemical Physics* "Neural mode jump Monte Carlo", Sbailò, Dibak, and Noé, 2021, with the permission of AIP Publishing.

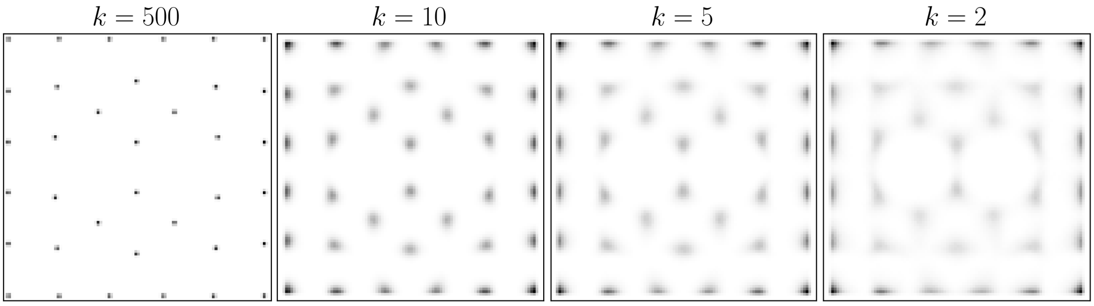

Figure 5.3.: 2D histograms of the solvent particles of the training data that was used at different strengths of the bias potential. At a high bias strength k, the particles are highly localized which enables the neural network to learn the mean position of the particles. The particles are less localized with a decrease in the strength of the bias potential k, thus enabling to learn an approximation to the true particle distribution.

In the sampling step, 100 independent trajectories of length 10^5 steps are generated and averaged. The marginal distributions p_X and p_Y which are the projections of the Boltzmann distribution on the X and Y axes are compared and great agreement is observed to results from numerical integration of the Boltzmann distribution, Fig. 5.2.

5.5.2. Dimer in repulsive Lennard Jones bath

As a bigger challenge, Neural MJMC is applied to a two-dimensional system composed of a bistable dimer immersed in a bath of strongly repelling particles and confined to a box. The bistable dimer potential has a minimum in the closed and open configurations, which are separated by a high energy barrier, Fig. 5.4 (right bottom). Opening and closing of the dimer requires a concerted motion of the solvent particles, that makes it difficult to sample the physical path connecting the two configurations. Appendix A.2.2 gives a detailed description of the system.

The open and closed configuration serve as cores [Fig. 5.4 (right top)] in Neural MJMC and are distinguished by the distance between the dimer particles. The neural network is trained on states sampled independently in the closed and open configuration at four different bias strengths with 10^5 samples for each well and bias. Figure 5.3 shows histograms of the solvent particles of the open configuration at the different bias strengths k_i used for training. As the system is invariant under permutation of solvent and dimer particles, neural moves would have to be learned independently for each permutation of the system That is clearly infeasible as the number of permutation scales factorially with the number of particles. This problem is circumvented by permutation reduction, i.e. exchanging the labels of the particles such that the labeling is consistent. This is accomplished by re-labeling the particles such that the distance to a reference configuration in each metastable state is minimized. To this end the Hungarian algorithm [176] is used to find the permutation of particle indices that minimizes the distance to the reference

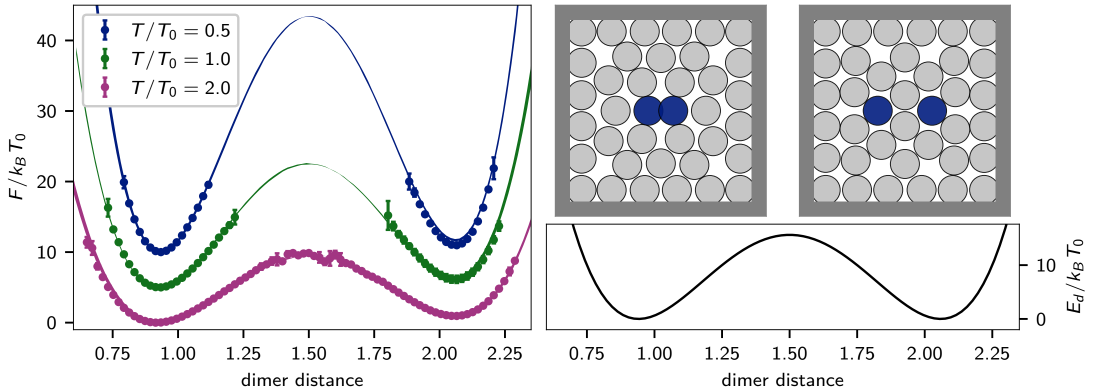

Figure 5.4.: **Left** Free energy along the distance between the dimer particles. The corresponding bands represent reference values obtained by umbrella sampling with the standard error given by their thickness. The neural network has been trained at temperature $T = T_0$, then simulations at different temperatures have been performed using Neural MJMC. Simulations are run for 1.5×10^7 steps, and error bars are generated from several sampling runs. In this figure, it is observed that Neural MJMC correctly samples the free energy along the reaction coordinate of the system at different temperatures. **Right top:** Reference configurations in the closed (left) and open (right) dimer configuration. The dimer particles are displayed in blue, and solvent particles in gray. The strongly repulsive potential does not allow for significant overlap between particles at equilibrium. **Right bottom:** Dimer interaction potential E_d as a function of the dimer distance. Reprinted from *The Journal of Chemical Physics* “Neural mode jump Monte Carlo”, Sbailò, Dibak, and Noé, 2021, with the permission of AIP Publishing.

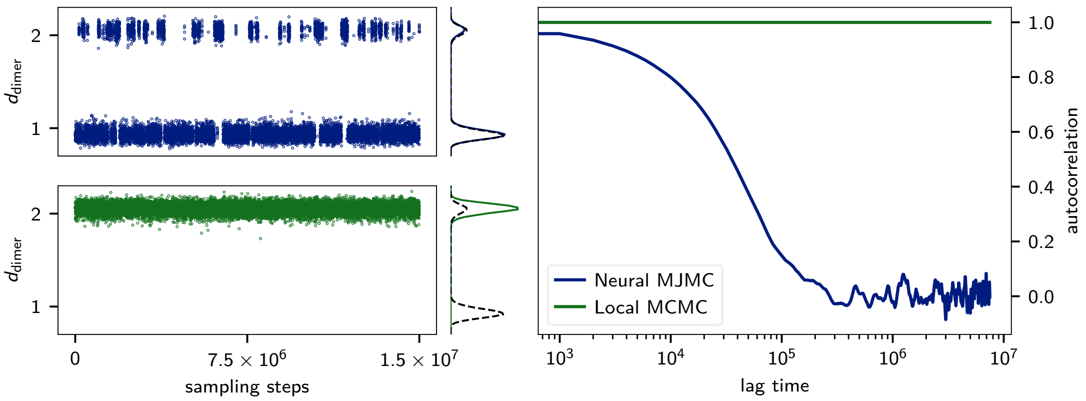

Figure 5.5.: **Left:** Dimer distance over a single realization using Neural MJMC (top), and using local MCMC (bottom), with histograms of the dimer distance obtained by the displayed trajectory on the right, where the reference value is displayed by the black dashed line. Spontaneous transitions with local MCMC are not observed at this time-scale. Neural MJMC explores both metastable states in the trajectory multiple times and correctly reproduces the distribution of dimer distances. **Right:** Autocorrelation of the dimer distance. Neural moves allow for a fast exploration of both metastable states, accelerating the production of uncorrelated samples. In this figure, it is evident that Neural MJMC frequently generates uncorrelated samples, and short trajectories are sufficient to reconstruct the right distribution. In contrast, configurations generated with local MCMC are highly correlated, as they do not cross the energy barrier. Reprinted from *The Journal of Chemical Physics* “Neural mode jump Monte Carlo”, Sbailò, Dibak, and Noé, 2021, with the permission of AIP Publishing.

configuration.

Each neural network in the RNVP architecture consists of three hidden layer with 76 nodes. The transformation consists of a total of 20 RNVP layers and contains approximately 1.4×10^6 trainable parameters. Neural MJMC is used to generate a single trajectory with 1.5×10^7 steps, where the probability of neural moves is set to $1\,\%$. In terms of computational performance, sampling with Neural MJMC is approximately a factor of four slower than MCMC with local displacements for this system. This slow down arises from the evaluation of the network and the remapping of particles. As a reference value, umbrella sampling [149] is used to sample the free energy along the dimer distance. To this end, 20 umbrellas are used along the dimer distance and the biased configurations are sampled with MCMC. The free energy is calculated using the multistate Bennet acceptance ratio [177] method.

Neural moves cause direct transitions between the two metastable states and thus a rapid exploration of the configuration space. The convergence to the Boltzmann distribution is observed as shown in Fig. 5.4 (left). An estimate for the crossing time with only local moves can be found to be at the order of 10^{12} sampling steps at $T = 1$ from the Kramers problem (Section 1.2) which make exhausting simulations, using local moves only, infeasible. In Neural MJMC many crossings of the energy barrier can be observed, Fig. 5.5 (left). This is also reflected in the autocorrelation function where samples generated with local MCMC remain highly correlated, while it decays in Neural MJMC simulations on a scale of approximately 10^5 sampling steps, Fig. 5.5 (right). Thus generating the desired uncorrelated samples of the equilibrium distribution.

5.6. Conclusion and outlook

This chapter presented Neural Mode Jump Monte Carlo (Neural MJMC), a novel method that allows for efficient sampling of the Boltzmann distribution of complex systems composed of metastable states. The method uses invertible neural networks architectures, called flows, in order to parametrize bijections between metastable regions in phase space and optimizes these for bi-directional acceptance probability. To this end, samples from either metastable state are required during training, which can e.g. be generated using local MCMC moves only. To ensure the ergodicity of the scheme, it combines short steps given by random displacements and large jumps between metastable states. This allows the method to converge quickly to the Boltzmann distribution. This approach is especially advantageous in systems where large potential barriers are providing obstacles to the convergence of other methods, which is often observed in molecular systems. The method is demonstrated on two toy examples, one with several bijections in two dimensions and a high dimensional system consisting of a particle dimer in a bath of Lennard-Jones particles, which exhibits a highly rigid free energy surface. On both toy systems, it is shown that samples generated by Neural MJMC follow the correct Boltzmann distribution. Further, it is observed that the method is able to generate equilibrium samples in systems, where the standard MCMC method with local displacements fails.

The sampling efficiency of this method highly depends on the quality of the generated samples, which again crucially depends on the network architecture that is used in the implementation. Considering the great attention the field of invertible network structures is currently facing, it would be no surprise to see dramatic improvements in terms of computational efficiency and expressiveness of such invertible networks, which in turn would allow studying more complex problems. An important step to further generalization is the incorporation of physical symmetries into the network architecture. Many recent advances when applying machine learning to physical problems make use of a symmetric representation [178–180]. This has also been applied to the field of invertible networks [181] and using an architecture that preserves symmetries appears promising to allow for scaling of the proposed method to even bigger systems.

To utilize this method to study molecules, a suitable representation of the molecule is of high significance. One approach is to represent the molecule in terms of bonds and angles. This has recently been shown to generate promising results in other methods that rely on the generation of configurations using flows [146, 160] and an extension of the method proposed here is straightforward.

One drawback that the method in its current implementation faces is that it requires training a new neural network for each combination of states. For systems with many metastable states, this might be a hurdle as it requires training $\mathcal{O}(N^2)$ networks to fully parametrize the transition kernels between all combinations of metastable states. A possible way to solve this problem would be to train a transformation f_i to the latent space for each of the metastable states, where the structure of the latent space is given by e.g. a Gaussian distribution as in Boltzmann Generators [159], but restricted to one metastable state Ω_i, resulting in a set of flows $\{f_i\}$, that share the same latent space structure. When using these flows to sample the system, a proposal $\mathbf{y}$ could then be generated by composing a forward and backward transformation $\mathbf{y} = f_j^{-1} \circ f_i(\mathbf{x})$, where $\mathbf{x} \in \Omega_i$ and j is drawn randomly from transition probabilities p_{ij}. This would require training only one flow per metastable state, thus requiring only N such invertible functions.

6. Temperature-steerable flows

The results of this chapter have been published in:

> Manuel Dibak, Leon Klein, and Frank Noé. "Temperature Steerable Flows and Boltzmann Generators". 2021. arXiv: `2108.01590`

Parts of text and figures have been adopted unchanged in this chapter. The manuscript has been submitted and is currently under review.
Manuel Dibak (MD) and Leon Klein (LK) contributed to this work equally. In particular the contributions were as follows: All contributors conceptualized the work. MD and LK layed out the theory. MD implemented the sampling algorithms. MD and LK performed the numerical experiments. MD visualized the results. All contributors wrote and approved the manuscript.

Summary

Boltzmann generators (BG) approach the sampling problem in many-body physics by combining a normalizing flow and a statistical reweighting method to generate samples of a physical system's equilibrium density. The equilibrium distribution is usually defined by an energy function and a thermodynamic state, such as a given temperature. Accessing multiple thermodynamic states with a single BG would highly improve the applicability of the method. In this chapter temperature-steerable flows (TSF) are proposed which are a type of BG able to generate a family of probability densities parametrized by a choosable temperature parameter. To this end, a temperature scaling condition is derived, which provides a condition to the scaling behavior of the flow. Two different flow architectures are proposed which follow this scaling condition to a good approximation. TSFs can be embedded in a generalized ensemble sampling framework such as parallel tempering in order to sample a physical system across multiple temperatures. The temperature scaling property of the proposed flows is verified on toy systems and an application of either of the proposed flows is given on physically relevant systems, namely the XY model and the alanine dipeptide molecule, showing that temperature-dependent quantities can accurately be recovered in the first system and that the sampling efficiency can be enhanced using MCMC moves proposed by the TSF at different temperatures.

6.1. Introduction

In the canonical ensemble the stationary distribution of configurations $\mathbf{x}$ is given by the Boltzmann distribution

$$\mu^{\tau}(\mathbf{x}) \propto \exp(-U(\mathbf{x})/\tau), \tag{6.1}$$

where the thermal energy $\tau = k_{\mathrm{B}}T$ is proportional to the temperature T and k_B is the Boltzmann constant. In order to reach all thermodynamic states, it would be desirable to sample a family of densities parameterized by the thermodynamic control variables τ.

Recently, there has been a lot of interest to train normalizing flows [161, 162, 166–168, 170, 182] to sample densities of many-body physics systems such as Eq. (6.1) directly without having to run long, correlated simulation chains. Normalizing flows transform an easy to sample prior distribution $p_Z(\mathbf{z})$, e.g. a multivariate normal distribution, via a transformation $\mathbf{x} = f(\mathbf{z})$ to the output distribution $p_X(\mathbf{x})$. If $f(\mathbf{z})$ is invertible, $p_X(\mathbf{x})$ can be computed by the change of variable formula

$$p_X(\mathbf{x}) = p_Z(\mathbf{z}) \left|\det J_f(\mathbf{z})\right|^{-1}, \tag{6.2}$$

where $\left|\det J_f(\mathbf{z})\right|^{-1}$ is the inverse of the Jacobian. Boltzmann Generators (BGs) (Section 4.5.2) combine normalizing flows to minimize the distance between Eq. (6.1) and Eq. (6.2) with a statistical reweighting or resampling method to generate unbiased samples from Eq. (6.1). This and similar approaches have been used to sample configurations of molecular and condensed matter systems [146, 159], spin models [183, 184], and gauge configuration in lattice quantum chromodynamics [185, 186].

Constructing flow architectures whose samples change in a predictable manner in response to changes at the input are important for physics and other disciplines. Equivariant flows maintain group transformations such as rotation and permutation throughout the flow [181, 187, 188]. Here, the focus lies on a specific form of equivariant flow which is better described as steerable, as its output is controlled by a parameter. The aim is to learn a flow whose output distribution follows

$$p_X^{\tau}(\mathbf{x}) \propto \exp\left(-U(\mathbf{x})/\tau\right), \tag{6.3}$$

where the control parameter τ corresponds to the thermal energy.

Specifically, TSFs are introduced, which correctly parametrize the output distribution p_X of a normalizing flow f by a temperature variable τ. The following section introduces a condition for such flows and uses it to develop two different realizations of TSFs. The temperature steering property of the proposed flow architectures is verified on simple toy systems and an application of each TSF is provided on more complex systems. One architecture is applied to the XY model, which shares similarities with the Ising model, but utilizes continuous spins. In this system, it finds the correct temperature dependence of the magnetization. The other architecture is applied to the alanine dipeptide molecule, where it is shown that the TSF is capable of producing samples close to equilibrium at different temperatures. Furthermore, it can be turned into a multi-temperature BG by using it as a proposal density in a PT framework, reducing autocorrelation times significantly compared to conventional PT.

6.2. Temperature-steerable flows

Up to a normalization constant, a change to temperature τ' of the Boltzmann distribution corresponds to raising it by the power of $\kappa = \tau/\tau'$, $\mu^{\tau'}(\mathbf{x}) \propto [\mu^{\tau}(\mathbf{x})]^{\kappa}$. Using Eq. (6.2) it is observed that the output distribution of a flow scales equivalently to a Boltzmann distribution, if for any two temperatures τ, τ'

$$p_Z^{\tau'}(\mathbf{z}) \left|\det J_{f_{\tau'}}(\mathbf{z})\right|^{-1} \propto \left[p_Z^{\tau}(\mathbf{z}) \left|\det J_{f_{\tau}}(\mathbf{z})\right|^{-1}\right]^{\kappa}. \tag{6.4}$$

Thus, flows are considered to be temperature scaling, if they preserve this scaling condition. Two instances of flows that preserve this proportionality can be found by either keeping the Jacobian constant and preserving the scaling condition in the prior, or selecting a constant prior and fulfilling the scaling condition in the flow.

6.2.1. Temperature-steerable flows by volume preservation

The proportionality in the prior distribution can be matched by selecting a Gaussian prior $p_Z^{\tau}(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \tau)$, which fulfills $p_Z^{\tau'}(\mathbf{z}) \propto [p_Z^{\tau}(\mathbf{z})]^{\kappa}$. This results in a condition on the Jacobian of the flow

$$\left|\det J_{f_{\tau}}(\mathbf{z})\right|^{\kappa} \propto \left|\det J_{f_{\tau'}}(\mathbf{z})\right|. \tag{6.5}$$

This condition is met by volume-preserving flow layers, i.e. $\left|\det J_{f_{\tau'}}(\mathbf{z})\right| = 1$, such as NICE [182], or flows with constant Jacobian, i.e. $\left|\det J_{f_{\tau'}}(\mathbf{z})\right| = \text{const}$. The latter are considered more expressive, as they are able to correct for difference in entropy between the prior and the target distribution.

The recently proposed stochastic normalizing flows (SNF) [146] can be used to construct TSFs. Here, a type of SNF is used which is motivated by HMC and utilizes operations in an augmented space, where the auxiliary momenta are distributed according to a distribution p_A (see Section 4.4.4 for details on HMC). However, in contrast to the original HMC method [150], the system is not propagated by Hamiltonian dynamics, but rather a learned deterministic flow $\mathcal{D}$ which takes this role. This flow shares similarities to other augmented flows, such as Hamiltonian flows [189, 190] and Augmented Normalizing Flows (ANFs) [191]. To sample from the model, i.e. $\mathbf{x} = f_{\tau}(\mathbf{z}); \mathbf{z} \sim p_Z^{\tau}(\mathbf{z})$, the flow consists of three consecutive steps

(i) Sample the latent space variable $\mathbf{z} \sim p_Z^{\tau}(\mathbf{z})$

(ii) Sample auxiliary momenta $\mathbf{q} \sim p_A^{\tau}(\mathbf{q})$, defining the point in latent phase space $\upsilon = (\mathbf{z}, \mathbf{q})$

(iii) Propagate the point to phase space by the dynamics $\gamma = (\mathbf{x}, \mathbf{p}) = \mathcal{D}(\upsilon)$

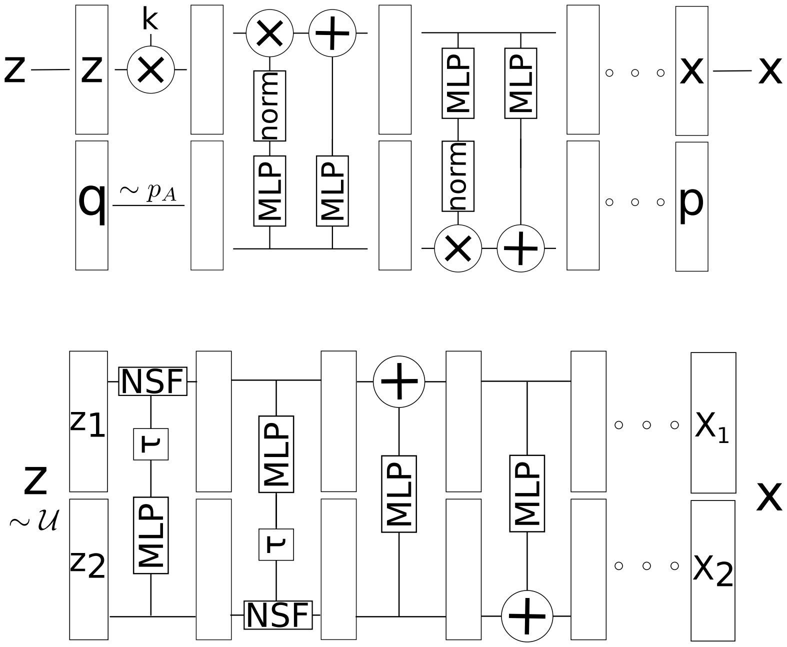

Figure 6.1.: Schematic figures of the two proposed TSF architectures. **Top:** Volume preserving augmented flow. Auxiliary momenta $\mathbf{q}$ are drawn to define a point in the augmented space $(\mathbf{z}, \mathbf{q})$. A volume-preserving flow is used to transform the augmented variable to the augmented configuration space, where it is projected onto the configuration space variable $\mathbf{x}$. In these flows, the outputs of the multi-layer perceptron (MLP) used to generate the scaling variables are normalized. The first layer multiplies the latent space coordinates $\mathbf{z}$ with a scalar factor k, which adjusts for difference in entropy between latent and phase space. **Bottom:** Temperature steerable NSF architecture with uniform prior. Samples from the uniform distribution are split into two channels which condition the NSF transformation of the other channel. The parameters for the NSF are transformed to the given temperature. This is followed by several layers of NICE transformations in order to model correlations in the system while preserving the volume.

(iv) Project onto the configuration variables $\mathbf{x}$

With the flow described above, a TSF can be constructed by choosing $p_A^\tau(\mathbf{q}) = \mathcal{N}(\mathbf{q}; 0, \tau)$ and choosing a dynamics $\mathcal{D}$ with a constant Jacobian, thus satisfying the temperature scaling condition Eq. (6.4). A convenient way of constructing volume preserving dynamics, i.e. $|\det J_D(\upsilon)| = 1$, is obtained by altering the RNVP network structure (Section 4.5.1), such that the product of the outputs of the scaling layers is equal to unity. This is done by subtracting the mean of the log outputs from each scaling layer, similar to [192]. In addition to the volume preserving RNVP layers, the latent space coordinates are scaled by a trainable scalar, which allows to adjust for entropy difference between the prior and the target. The resulting Jacobian factor is a constant and, hence, still fulfills the scaling condition. The flow architecture is shown in Fig. 6.1 (top).

This architecture can also be viewed as an instance of an ANF, where the augmented prior distribution $p_\Upsilon^\tau(\upsilon) = p_Z^\tau(\mathbf{z}) p_A^\tau(\mathbf{q})$ is mapped to the joint output distribution $p_\Gamma^\tau(\gamma)$ via the invertible dynamics D. The joint target distribution is given by $\mu^\tau(\mathbf{x}) p_A^\tau(\mathbf{p})$. As the flow fulfills the temperature scaling condition, a temperature change of the prior, i.e. $\tau \to \tau'$, will change the output accordingly. In the case of a factorized output distribution $p_\Gamma^\tau(\gamma) = p_X^\tau(\mathbf{x})\, p_A^\tau(\mathbf{p})$, the marginal distribution $p_X^\tau(\mathbf{x})$ is scaled correctly as well. This is ensured if the joint target distribution is matched correctly.

6.2.2. Temperature-steerable flows with uniform prior

In order to fulfill the temperature steering condition in the flow, it has to be taken into account, that changing the Jacobian will also change the output of the flow, i.e. $f_\tau^{-1}(\mathbf{x}) \neq f_{\tau'}^{-1}(\mathbf{x})$. This can be circumvented by selecting a prior for which $p_Z(\mathbf{z}) = \text{const.}$, namely the uniform distribution $\mathcal{U}$. This approach only allows using one of these layers, as the same problem arises with subsequent layers.

While finding a flow architecture that precisely reproduces the temperature scaling property is difficult, a good approximation can be found using Neural Spline Flows (NSF) [163].

A spline flow uses coupling layers as introduced in Section 4.5.1. The transformation of each element of the input vector is given by an invertible scalar function $y_i = h(x_i)$ defined on the unit interval $h : [0, 1] \to [0, 1]$, s.t. $h(0) = 0$ and $h(1) = 1$. For this transformation to be invertible, it needs to be monotonous. It can be interpreted as the cumulative distribution function corresponding of a probability density function p defined on the unit interval. In neural spline flows, this function $h(x)$ is approximated by a spline $s(x)$, a piecewise defined function, which is invertible in each interval. The spline is parametrized by N nodes and slopes $(x^{(i)}, y^{(i)}, \delta^{(i)})$ which are the function's values and derivatives $y^{(i)} = h(x^{(i)})$, $\delta^{(i)} = p(x^{(i)})$. It then interpolates the function values between these coordinates. There exists a whole range of different ways to define a spline. The recently proposed NSF [163] use quadratic rational splines [193] and have been shown to perform best in a series of tasks. Thus, these have been utilized in this work.

A quadratic rational spline can alternatively be defined by splitting the unit interval $[0, 1]$ into N bins. The corresponding spline is then parametrized by N bin-widths, bin-heights and slopes (w_i, b_i, δ_i), with $\sum_i w_i = 1$ and $\sum_i b_i = 1$. Assuming that the spline is an approximation of $h(x)$, it is found that $b_i = \int_{x_i}^{x_i+w_i} p(x)dx$, with $x_i = \sum_{j=1}^{i-1} w_i$. For the temperature scaling to hold, it is assumed that $p(x)$ is of the form $p(x) = \exp(\beta v(x))$ for some continuous function $v(x)$.

To parametrize the transformation at temperature τ by a spline, its slopes need to be given by $\delta_i = \exp(\beta v_i)$, with $v_i = v(x_i)$. The bin-heights can be approximated by the mean value theorem $b_i = \int_{x_i}^{x_i+w_i} p(x)dx \approx \exp(\beta \tilde{v}_i)w_i$, with $\tilde{v}_i = v(\xi)$ for some $\xi \in [x_i, x_{i+1}]$. With these assumptions, the temperature steerable spline flow is parametrized by N values of $(\tilde{w}_i, v_i, \tilde{v}_i)$ from which the bin widths, heights and slopes at temperature τ are computed as

$$w_i^\tau = \tilde{w}_i / \sum_i \tilde{w}_i \tag{6.6}$$

$$b_i^\tau = \exp(\tau \tilde{v}_i) w_i / \sum_i \exp(\tau \tilde{v}_i) w_i, \tag{6.7}$$

$$\delta_i^\tau = \exp(\tau v_i). \tag{6.8}$$

Thus the transformation is fully parametrized by the set of N values $(w_i, v_i, \tilde{v}_i)$ per dimension of the transformed channel, and the temperature parameter τ. These parameters are generated by a NN.

With this type of flow it is possible to adjust the parameters given the temperature, s.t. the temperature scaling is approximately correct. The NSF are especially useful when transforming a quantity with circular symmetry, such as angles, as they can easily be adjusted to satisfy the periodicity of the variables [194]. In practice, a one layer temperature scaling NSF will not be able to capture potentially complex correlations in the system. This can be achieved by adding volume preserving flows, i.e. NICE [157] to obtain a more expressive transformation, see Fig. 6.1 (bottom).

6.2.3. Training procedure

As in [159], the TSFs are trained by a convex combination of a maximum-likelihood (forward KL) and energy-based (reverse KL) loss $\mathcal{L} = (1-\lambda)\mathcal{L}_{ML} + \lambda \mathcal{L}_{KL}$, where the mixing parameter λ can be changed during training. The maximum likelihood loss minimizes the negative log likelihood (nll), Eq. (4.34), which requires samples from the target distribution $\mu^\tau(\mathbf{x})$. Available samples typically do not exactly follow this distribution, as otherwise the sampling problem would be solved already. Instead suppose the samples follow the empirical sampling distribution $\varrho_X(\mathbf{x})$, which can, for example, be obtained from short simulations that have not fully converged yet. The forward KL-divergence is given by

$$\mathrm{nll} = \mathcal{L}_{ML} = \mathbb{E}_{\mathbf{x}\sim\varrho_X(\mathbf{x})}\left[-\log p_Z(f^{-1}(\mathbf{x})) - \log\left|\det J_{f^{-1}}(\mathbf{x})\right|\right] \tag{6.9}$$

In the case of the volume preserving TSF described in Section 6.2.2, the flow needs to match between the augmented distribution and the ML-loss can be expressed as

$$\mathcal{L}_{ML} = \mathbb{E}_{\gamma \sim \varrho_\Gamma(\gamma)} \left[-\log p_\Upsilon(\mathcal{D}^{-1}(\gamma)) - \log \left|\det J_{\mathcal{D}^{-1}}(\gamma)\right| \right]. \tag{6.10}$$

As the target energy $U(x)$ is defined by the physical system of interest, optimization can also be performed with energy-based training which minimizes the reverse KL divergence via the variational free energy

$$\mathcal{L}_{KL} = \mathbb{E}_{\mathbf{z} \sim p_Z(\mathbf{z})} \left[U(f_\tau(\mathbf{z})) - \log \left|\det J_{f_\tau}(\mathbf{z})\right| \right] \tag{6.11}$$

$$\mathcal{L}_{KL} = \mathbb{E}_{\upsilon \sim p_\Upsilon(\upsilon)} \left[U(\mathcal{D}(\upsilon)) - \log \left|\det J_{\mathcal{D}}(\upsilon)\right| \right], \tag{6.12}$$

where the last line shows the loss for the augmented TSF.

6.3. Unbiased sampling: importance weights and latent Monte Carlo

One way to generate samples using the TSF is by directly sampling from the prior and transforming the samples to configuration space at the desired temperature. As the TSF is only an approximation to the Boltzmann distribution, these samples will generally be biased.

6.3.1. Reweighting scheme

In order to generate unbiased samples, two different methods are used here: Firstly, the generated distribution is used as a proposal distribution for a reweighting scheme as described in Section 4.1. This involves computing a weight w_i for each sample $\mathbf{x}_i$. An observable can then be computed as a weighted average $\langle O \rangle = \sum_i w_i O(\mathbf{x}_i) / \sum_i w_i$, with unnormalized weights $w_i = \exp\left[-\beta U(\mathbf{x}_i) - \log p_X(\mathbf{x}_i)\right]$. This only works if the energetic overlap between the generated distribution and the Boltzmann distribution is sufficiently large. In the case of a TSF, the (unnormalized) weights take the form

$$w_i = \exp\left[-\kappa \left(u(\mathbf{x}_i) + \log p_Z^{\tau_0}(\mathbf{z}_i) + \log \left| \det J_{f_{\tau_0}}(\mathbf{z}_i) \right| \right) \right], \tag{6.13}$$

with $u(\mathbf{x}) = U(\mathbf{x})/\tau_0$ and the relative temperature κ factors out in the exponential due to the temperature scaling condition.

6.3.2. Temperature steerable flows as proposal density

Another approach for producing unbiased samples from Eq. (6.1) with the augmented TSF is MCMC. In a sampling scheme similar to HMC (see Section 4.4.4), a proposal $\mathbf{x}'$ is generated from configuration $\mathbf{x}$ by the following steps:

(i) sample auxiliary momenta $\mathbf{p} \sim p_A^\tau(\mathbf{p})$ to define $\gamma = (\mathbf{x}, \mathbf{p})$

(ii) apply the inverse dynamics $\upsilon = \mathcal{D}^{-1}(\gamma)$

(iii) add a random displacement $\upsilon' = \upsilon + \xi$ in latent space, with $\xi \sim \mathcal{N}(0, \sigma^2)$ with some variance σ

(iv) transform back into configuration space $(\mathbf{x}', \mathbf{p}') = D(\upsilon')$.

Accepting such a step with acceptance probability

$$p^{\tau}_{\text{acc}}\left((\mathbf{x}, \mathbf{p}) \to (\mathbf{x}', \mathbf{p}')\right) = \min\left\{1, \exp\left[-\beta\left(U(\mathbf{x}') - U(\mathbf{x}) + \left\|\mathbf{p}'\right\|^2/2 - \left\|\mathbf{p}\right\|^2/2\right)\right]\right\} \tag{6.14}$$

enforces detailed balance in configuration space and thus ensures convergence to the Boltzmann distribution. This can be seen from Eq. (4.26) by noting that the dynamics $\mathcal{D}$ is volume preserving and the random displacement in the latent space is symmetric.

As the TSF is able to generate distributions at several temperatures, the MCMC moves can be combined with temperature based enhanced sampling techniques, such as the ones discussed in Section 4.4.1. Here PT is selected for such simulations. In PT, sampling is performed at a set of temperatures in parallel. Additionally to TSF-MCMC steps, samples can randomly be exchanged between two randomly chosen temperatures. This allows the sampler to overcome energy barriers quickly at high temperatures while still preserving details of the wells at lower temperatures. Following Eq. (4.22), for a random swap of temperatures between two samples $\mathbf{x}_i$ and $\mathbf{x}_j$ at temperatures τ_i and τ_j, detailed balance is preserved by the acceptance probability

$$p_{\text{acc}}\left((\tau_i, \tau_j) \to (\tau_j, \tau_i)\right) = \min\left\{1, \exp\left[\left(U(\mathbf{x}_i) - U(\mathbf{x}_j)\right)\left(1/\tau_i - 1/\tau_j\right)\right]\right\}. \tag{6.15}$$

A summary of the sampling algorithm is given in Algorithm 5.

6.4. Experiments

The instantiations of the TSFs introduced in Section 6.2 are demonstrated on different test models. Firstly, the temperature scaling is demonstrated for both proposed architectures on fundamental test systems. Furthermore, it is demonstrated that both instantiations are able to produce meaningful results in more complex applications. Namely, it is demonstrate that the TSF is able to reproduce the equilibrium properties in a spin system and it is shown that it is able to generate samples of the alanine dipeptide molecule close to the Boltzmann distribution. Furthermore, it is demonstrated that the MCMC scheme results in unbiased samples with fast decorrelation times.

6.4.1. Temperature steered inversion sampler

To demonstrate the temperature steering capability of the temperature steerable NSF, the method is applied to two test system, namely the Prinz potential [9] and the asymmetric

Algorithm 5: Parallel tempering combined with TSF sampling algorithm.

```
input  : $l_s = [\,]$ : empty list for samples
         $\tau$ : list of $N_T$ temperatures
         state $\leftarrow$ init state : $N_\tau$ initial configurations
         $\mathcal{D}$: volume preserving TSF
         $N_{\text{iterations}}$: number of generated samples
         $n_{\text{propagate}}$: number of propagation steps
         $n_{\text{swap}}$: number of temperature swaps
for $i \leftarrow 1$ to $N_{\text{iterations}}$ do
    for $j \leftarrow 1$ to $n_{\text{propagate}}$ do
        for $k \leftarrow 1$ to $N_\tau$ do
            $\mathbf{x} \leftarrow \text{state}_k$
            $\mathbf{p} \leftarrow$ sample from $p_A^{\tau_k}$
            $(\mathbf{z}, \mathbf{q}) \leftarrow \mathcal{D}^{-1}((\mathbf{x}, \mathbf{p}))$
            $\omega \leftarrow$ sample from $\mathcal{N}(0, \sigma\mathbb{1})$
            $(\mathbf{z}', \mathbf{q}') \leftarrow (\mathbf{z}, \mathbf{q}) + \omega$
            $(\mathbf{x}', \mathbf{p}') \leftarrow \mathcal{D}((\mathbf{z}', \mathbf{q}'))$
            $p_{\text{acc}} \leftarrow p_{\text{acc}}^{\tau_k}((\mathbf{x}, \mathbf{p}) \to (\mathbf{x}', \mathbf{p}'))$, Eq. (6.14)
            if $r \sim \mathcal{U}(0,1) < p_{\text{acc}}$ :
                $\text{state}_k \leftarrow \mathbf{x}'$
    for $j \leftarrow 1$ to $n_{\text{swap}}$ do
        randomly select $\alpha, \beta \leqslant N_\tau$, $\alpha \neq \beta$
        $\mathbf{x} \leftarrow \text{state}_\alpha$
        $\mathbf{y} \leftarrow \text{state}_\beta$
        $p_{\text{acc}} \leftarrow p_{\text{acc}}((\tau_\alpha, \tau_\beta) \to (\tau_\beta, \tau_\alpha))$, Eq. (6.15)
        if $r \sim \mathcal{U}(0,1) < p_{\text{acc}}$ :
            $\text{state}_\alpha \leftarrow \mathbf{y}$
            $\text{state}_\beta \leftarrow \mathbf{x}$
    $l_{\text{s}}$.append (state)
output : list of samples $l_{\text{s}}$
```

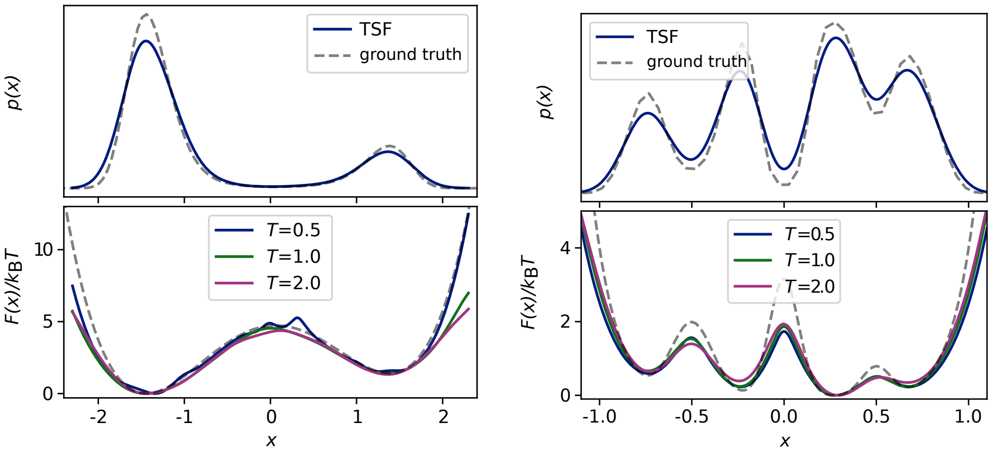

Figure 6.2.: Temperature steerable spline flow as a trained inverse sampler for 1D densities. The two figures on the **left** show an application to the asymmetric double well potential and the two on the **right** to the Prinz potential [9]. The free energies are given in units of $k_{\mathrm{B}}T$. In these units, the free energy curves nearly coincide for different temperatures, demonstrating the temperature steering property.

double well potential. The explicit energies of these systems are given in Appendix A.1.1. For these simple 1D densities only one transformation with fixed weights is used. The training is performed by purely energy based training [Eq. (4.35)] at temperature T_0. A generally good fit to the ground truth is observed, Fig. 6.2 (top row). The temperature steering property is evident by observing, that the free energy curves closely coincide at the three different temperatures when expressed in terms of their respective thermal energy, Fig. 6.2 (bottom row).

6.4.2. Mixture of multi dimensional double wells

The TSF architecture introduced in Section 6.2.1 is applied to a toy system consisting of a mixture of multi-dimensional double wells which are mixed via a correlation matrix. The energy of this system is given by $U(x) = U_{\mathrm{dw}}(A\mathbf{x})$, with $U_{\mathrm{dw}}(\mathbf{x}) = \sum_i^d a_i x_i + b x_i^2 + c {x_i}^4$. The parameter vector $\mathbf{a}$ and correlation matrix A are chosen at random. The system is studied at dimensions $d = \{5, 20\}$. In order to compare the capabilities of generating distributions at different temperatures, the TSF is compare with a RNVP flow with comparable numbers of trainable parameters.

Fig. 6.3 (left) shows the marginal density of the first 5 coordinates of the 20d system. The Boltzmann Generators are trained at $T = T_0$ and then analyzed by comparing the nll [Eq. (6.10)] of equilibrium samples, which were generated by Gaussian increment MCMC in combination with PT at 100 temperatures in the range $T = 0.1T_0$ to $T = 10T_0$, Fig. 6.3 (right). For the 5d system it is observed that both network structures perform equally well in the close vicinity of the training temperature, however the TSF exhibits

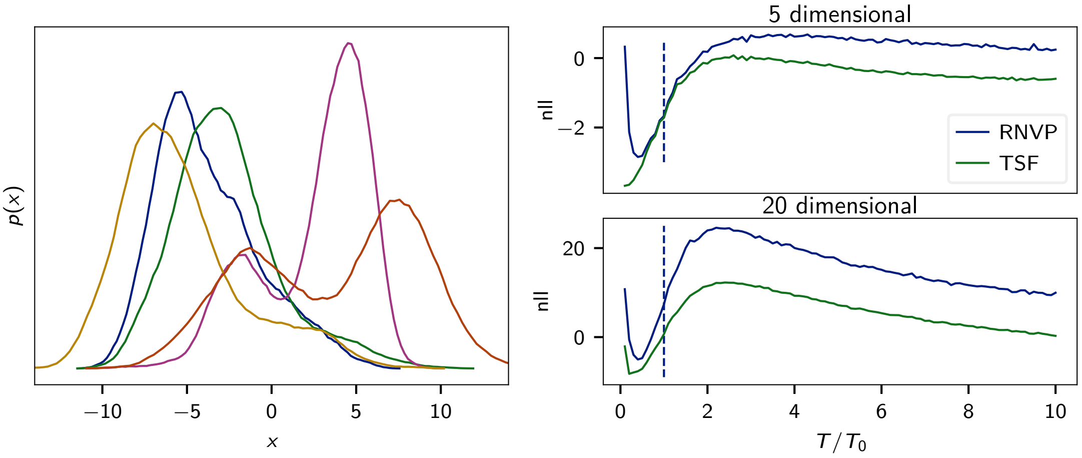

Figure 6.3.: Results for the multi-dimensional double well system. **Left:** marginal density of the first 5 coordinates of the 20d system. **Right:** comparison of the nll as a function of the temperature between the TSF and the RNVP network structure in the 5 dimensional (**top**) and 20 dimensional case (**bottom**). The dashed line indicates the training temperature $T = 1$

a significantly lower nll at temperatures further away from the ones it was trained on. In the case of the 20d system, the TSF consistently outperforms the RNVP even at the training temperature. This indicates that the TSF is a more expressive network structure with stronger temperature scaling.

6.4.3. XY-model

As an example with angular symmetry, the XY model [195] is investigated. It can be considered a version of the Ising model with continuous spins on the unit circle. For the experiments, the two dimensional lattice is chosen to be quadratic with $N \times N$ spins. The spins $\mathbf{s}_i = (\cos\theta_i, \sin\theta_i)^T$ are represented by their angle $\theta_i \in [-\pi; \pi]$ relative to the X-axes. A configuration is denoted as $\mathbf{S} = (\theta_1, \ldots, \theta_{N^2})^T$. Similar to the Ising model, each spin interacts with its four nearest neighbors and an additional external field $\mathbf{h} = (h, 0)$.

The Hamiltonian of the system is given by

$$\begin{aligned}\mathcal{H}(\mathbf{S}) &= -J \sum_{<ij>} \mathbf{s}_i \cdot \mathbf{s}_j - \sum_i \mathbf{h} \cdot \mathbf{s}_i \\ &= -J \sum_{<ij>} \cos(\theta_i - \theta_j) - \sum_i h\cos(\theta_i),\end{aligned} \tag{6.16}$$

where $\sum_{<ij>}$ denotes the sum over all pairs of nearest neighbor with periodic boundaries and J is the interaction constant. For the experiments, the interaction energy is chosen to be $J = h = k_B T_0$ and a lattice size of 16×16 spins is selected. Figure 6.4 (bottom)

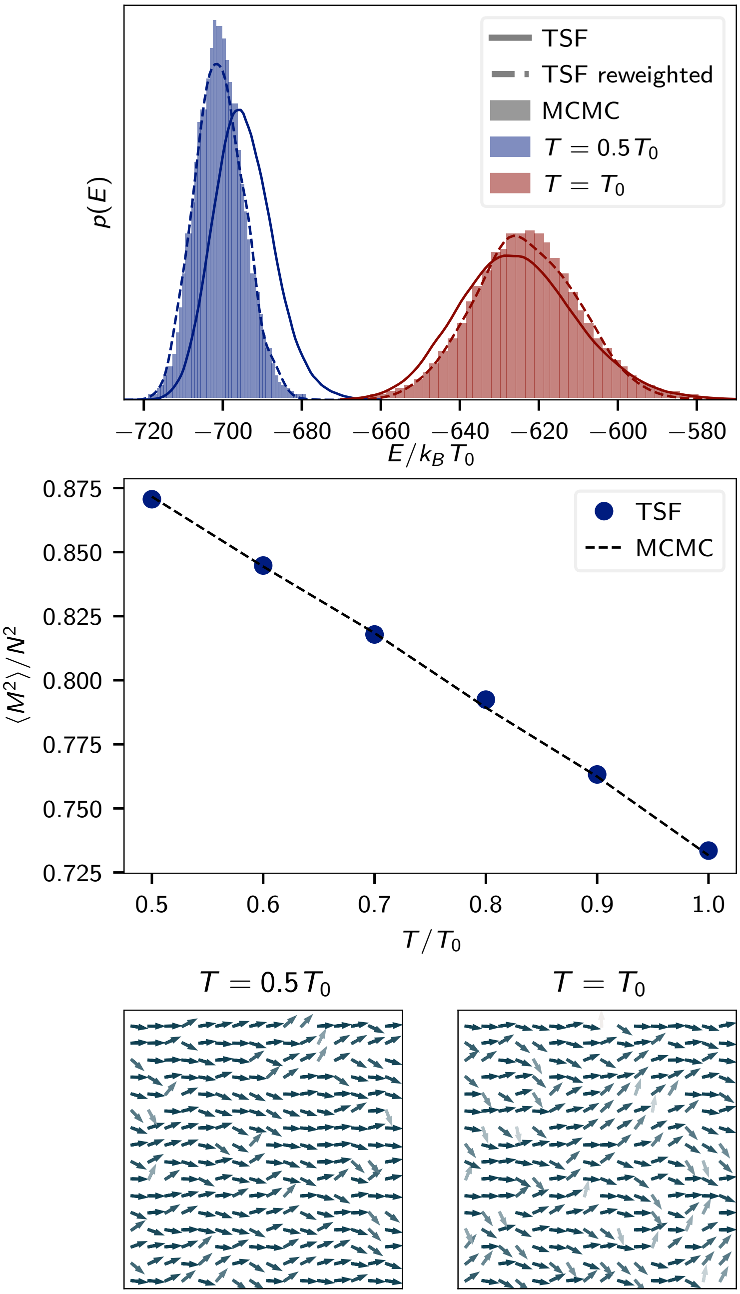

Figure 6.4.: Results of the TSF trained on the two dimensional XY model. **Top:** Distribution of energies at the training temperature T_0 and the samples temperature $T = 0.5T_0$ obtained by MCMC, TSF and TSF with reweighting. **Middle:** Magnetization as a function of the temperature compared between the TSF and MCMC samples. **Bottom:** Exemplary samples of the system at two different temperatures. The higher degree of order is evident from these samples.

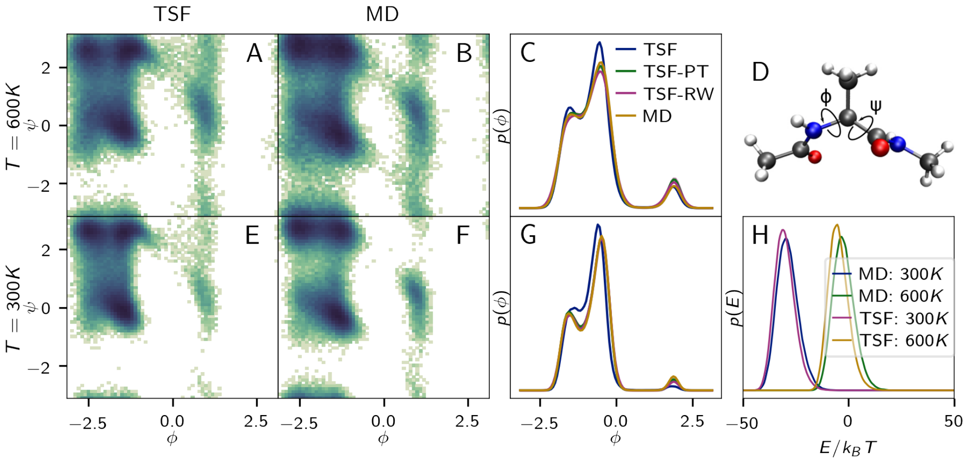

Figure 6.5.: Results for Alanine Dipeptide (**D**) in implicit solvent. **A, B, E, F:** Density of ϕ, ψ variables from different methods. **A:** TSF samples generated at the training temperature of $T = 600$ K (ground truth is **B**). **E:** Samples generated with the same TSF at $T = 300$ K to be compared (ground truth is **F**). Histograms of the dihedral angle ϕ are shown in **C** for $T = 600$ K and in **G** for $T = 300$ K. **H:** Energy histograms at the different temperatures. **D:** Alanine Dipeptide molecule.

shows example configurations as $0.5T_0$ and T_0. As temperature dependent observable, the mean squared magnetization per spin

$$\frac{\langle M^2 \rangle}{N^2} = \frac{1}{N^2} \left\langle \sum_i \mathbf{s}_i \cdot \mathbf{s}_i \right\rangle$$

is compared between the different methods. Reference configurations are generated using Glauber dynamics [196], which can be considered a MCMC scheme with single random spin perturbations. The TSF consists of a uniform prior, a temperature scaled NSF, followed by 7 blocks of circular NICE, Fig. 6.1 (bottom). The details of the network are given in Appendix A.1. Training is performed solely with the energy based loss Eq. (4.34) at a convex combination of temperatures $T = \{0.5\ T_0, T_0, 1.3T_0\}$. After training, samples are generated at temperatures ranging $0.5T_0, 0.6T_0, \dots, T_0$. The benchmark for the computation of the means squared magnetization is generated by long runs of Glauber dynamics. The distribution of energies, observed in the reference simulation, is precisely recovered by the TSF through reweighting at temperatures $T = T_0$ and $T = 0.5\ T_0$, Fig. 6.4 (top). Furthermore, the measurements of the mean squared magnetization exhibits excellent agreement between Glauber dynamics and the TSF, Fig. 6.4 (middle).

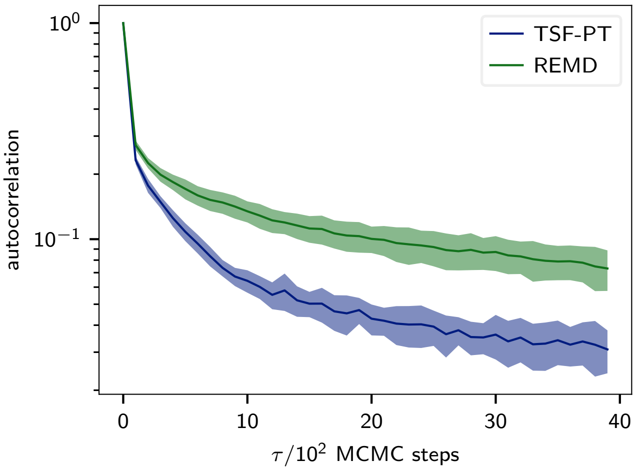

Figure 6.6.: Autocorrelation of the ϕ angle as a function of underlying MCMC/MD steps. The autocorrelation function decays more rapidly in the TSF-PT method. This hints toward this method being more sampling efficient. Non-traversing trajectories of the REMD method (as observed in Fig. 6.7) were excluded in the calculation of the autocorrelation.

6.4.4. Alanine Dipeptide

The augmented TSF is further tested on the Alanine Dipeptide molecule in an implicit solvent model. For this system, an invertible transformation to internal coordinates is used in order to represent the molecule in terms of distances and angles of the molecule's bonds. In order to investigate the temperature scaling aspect, sets of samples at temperatures $T = \{300\,\mathrm{K}, 600\,\mathrm{K}\}$ are generated using MD. The goal of this study is to use the samples at $T = 600\,\mathrm{K}$ to train the TSF and then use the TSF to sample at $T = 300\,\mathrm{K}$. The MD set at $T = 300\,\mathrm{K}$ serves as ground truth for comparison. The TSF is used to generate samples in configuration space and compare the Ramachandran plots and distributions of the ϕ angle (Fig. 6.5). Good agreement is observed at the training temperature (Fig. 6.5 (A to C)). At $T = 300\,\mathrm{K}$ (Fig. 6.5 (E to G)) the TSF still finds the major minima at around $\phi \approx -2$, but under-samples the minimum at $\phi \approx 1$. This deviation from the target distribution is likely stemming from limited expressivity of the flow. The correct distribution of the ϕ angle is recovered when using the Monte Carlo scheme in a PT fashion (Fig. 6.5 (C and G)). To this end parallel tempering is run with 8 temperatures in the range 300 K to 600 K. Furthermore, the output of the TSF closely recovers the distribution of energies (Fig. 6.5 (H)) at both temperatures.

To assess the efficiency of the sequential sampling procedure, it is compared to the replica exchange molecular dynamics simulation (REMD) at the same temperatures. It is observed that within 10×10^6 steps of the REMD simulation, only 7 out of 10 independent runs transition between the metastable states, whereas all 10 independent runs transition with the TSF-PT method, which additionally only consists of 4×10^6 steps, Fig. 6.7.

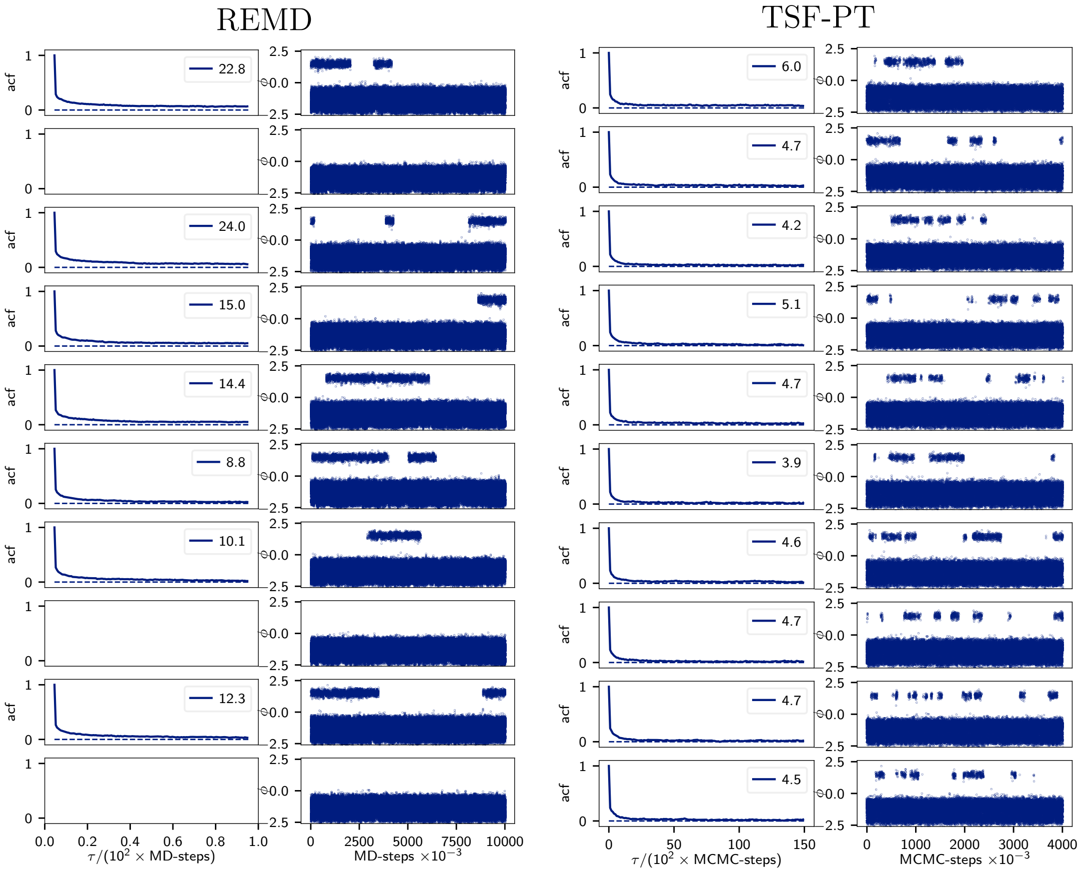

Figure 6.7.: Trace plots and autocorrelation of the ϕ angle in the sampling runs of the alanine dipeptide molecule at $T = 300\,\mathrm{K}$. Left for replica exchange molecular dynamics (REMD), right for TSF-PT. The numbers in the legends of the acf plots are the integrated autocorrelation times τ_{int}. These are omitted in the 3 trajectories of the REMD method that do not show any transition between the states. In the TSF-PT method, more frequent transitions between the metastable states at $\phi < 0$ and $\phi > 0$ are observed. This is also reflected in the integrated autocorrelation times, which are considerably lower for TSF-PT.

Table 6.1.: Efficiency as the number of effective steps per underlying sampling step for different sampling methods.

Method:	TSF-PT	REMD	MD(600K)	MD(300K)
$\eta \times 10^4$	1.36	0.32	0.38	0.02

Furthermore, the autocorrelation of the slowest process [Fig. 6.6 (bottom)], which are the transitions along the ϕ angle, decays considerably faster in the TSF-PT method. Additionally, the sampler efficiency $\eta = N_{\text{eff}}/N$ according to Eqs. (4.18) and (4.23) is compared between the different methods, where N is the number of underlying MCMC steps and N_{eff} is the number of effective sample size. The comparison Table 6.1 shows that the TSF-PT method produces independent samples at about four times the rate of the REMD method and significantly faster than MD at the target temperature.

The correct distribution can also be recovered by the importance sampling method (Section 4.1). However, this approach relies on weights being clipped appropriately, as otherwise the weight of a single sample will outweigh all the others, resulting in a large bias as discussed in Section 4.1.1. In this scenario the weight clipping scheme proposed in Ref. [197] is used in which a truncation value of the weights has to be defined and weights higher than that value are clipped. As the results heavily depend on the clip value, the MCMC method is considered to be preferable.

6.5. Discussion

Temperature-steerable flows (TSF) constitute a novel type of flow architecture, that correctly scales the output distribution of a Boltzmann Generator (BG) with temperature. In this chapter, a condition for such flows was formulated and two different realizations for a TSF were introduced. It was shown that this type of flow can be used to train a BG at one temperature and generate distributions at lower temperatures. The temperature steering property of the two proposed flow architectures was verified on simple test systems and possible applications were shown on more complex systems. For the XY model, the method was able to predict the correct temperature dependence of the magnetization after only using energy-based training on a set of temperatures. A second example demonstrated a use case of a TSF: A TSF that generates samples of the alanine dipeptide molecule was trained at a high temperature, where samples could easily be generated by conventional MD. The trained TSF was then utilized to generate samples at a considerably lower temperature. To this end, the TSF was embedded in a parallel tempering framework, where the MCMC proposals were generated with the TSF at different temperatures. With this method, the Boltzmann distribution could be accurately recovered. Furthermore, it was shown, that the sampling efficiency of this combined TSF-PT sampler outperforms conventional REMD in the given system.

Appendix A.

Detailed description of the systems and networks

A.1. Temperature-steerable flows

A.1.1. One dimensional test systems

The Prinz potential [9] is a one dimensional test system defined by the energy

$$U_{\mathrm{P}}(x) = 4\left[x^8 + 0.8\mathrm{e}^{-80x^2} + 0.2\mathrm{e}^{-80(x-0.5)^2} + 0.5\mathrm{e}^{-40(x+0.5)^2}\right]. \tag{A.1}$$

The energy of the one dimensional asymmetric double well is given by

$$U_{\mathrm{DW}}(x) = a\ x + b\ x^2 + c\ x^4, \tag{A.2}$$

with $a = 0.5$, $b = -4$ and $c = 1$.

A.1.2. Alanine dipeptide system

The following table gives the force fields and parameters used for the simulations of the alanine dipeptide molecule in implicit solvent.

Force Fields	`Amber ff99SB-ILDN` `Amber ff99-OBC`
Number Atoms	22
Total simulation time 600K	10 ns
Total simulation time 300K	1 ms
PT temperatures [K]	300.0, 331.2, 365.7, 403.8, 445.8, 492.2, 543.4, 600.0

A.1.3. Network parameters

Training of the networks was performed using Adam [198]. The examples using the volume preserving layers utilize the ReLU[199] activation function in the translation blocks, and a tanh activation in the scaling blocks. The examples NNs in the neural spline flows use Swish[200] as activation function. The parameters of each of the neural

Table A.1.: Network parameters used in the examples of Chapter 6

System	parameters	coupling blocks	hidden dimensions
DW 5	22k	20	[10, 20, 20, 10]
DW 20	330k	50	[20, 40, 40, 20]
Ala2	1.9M	50	[49, 98, 98, 49]
Prinz	30	1	-
Asym DW	30	1	-
XY-model	8.5M	7	[256, 256, 256]

networks used in the examples of the chapter are given in Table A.1.

A.2. Neural Mode Jump Monte Carlo

Here, the exact functional form of the potentials used to demonstrate the Neural MJMC method in Chapter 5 are provided

A.2.1. Triple well potential

Triple well potential is a 2D potential surface given by

$$V(\mathbf{x}) = \sum_i -a \exp\left[-\left(\mathbf{x}-\mathbf{m}_i\right)^T \Sigma_i \left(\mathbf{x}-\mathbf{m}_i\right)\right] + b\left\|\mathbf{x}\right\|^2, \tag{A.3}$$

with $b = 0.1 k_{\mathrm{B}}T/l$, $a = 5 k_{\mathrm{B}}T$ and other parameters given in table A.2, where l is an (arbitrary) length-scale of the system.

Table A.2.: Parameters of the triple well potential. All quantities are given in units of the (arbitrary) length scale l

i	$(\Sigma_i)_{11}$	$(\Sigma_i)_{22}$	$(\mathbf{m}_i)_1$	$(\mathbf{m}_i)_2$
1	0.5	0.3	-2.2	-1
2	0.5	0.4	0	2
3	0.4	0.5	2	-0.8

A.2.2. Dimer in a Lennard Jones bath

The dimer system is adapted from Ref. [159]. It consists of $N = 38$ particles in 2 dimensions, two of which are considered the *dimer*, which additional interact via a double well interaction U_d. Its potential energy is given by

$$U(\mathbf{x}_1, \ldots, \mathbf{x}_N) = U_\mathrm{box}(\mathbf{x}_1, \ldots, \mathbf{x}_N) + U_\mathrm{rep}(\mathbf{x}_1, \ldots, \mathbf{x}_N) \tag{A.4}$$

$$+ U_\mathrm{d}(\mathbf{x}_1, \mathbf{x}_2), \quad \text{with} \tag{A.5}$$

$$U_\mathrm{box}(\mathbf{x}_1, \ldots, \mathbf{x}_N) = k_\mathrm{box} \sum_{i=1}^{N} \sum_{d=1}^{2} \theta\left(|x_{i,d}| - l_\mathrm{box}\right)\left(|x_{i,d}| - l_\mathrm{box}\right)^2, \tag{A.6}$$

$$U_\mathrm{rep}(\mathbf{x}_1, \ldots, \mathbf{x}_N) = \varepsilon \sum_{j \leqslant N} \sum_{\substack{i<j \\ (i,j) \neq (1,2)}} \left(r_\mathrm{m}^{-1} \left\|\mathbf{x}_i - \mathbf{x}_j\right\|_2\right)^{-12}, \tag{A.7}$$

$$U_\mathrm{d}(\mathbf{x}, \mathbf{y}) = a d\left(\mathbf{x}, \mathbf{y}\right) + b d\left(\mathbf{x}, \mathbf{y}\right)^2 + c d\left(\mathbf{x}, \mathbf{y}\right)^4 \tag{A.8}$$

$$+ k_d \left[\left(x_1 + y_1\right)^2 + x_2^2 + y_2^2\right], \tag{A.9}$$

with $d\left(\mathbf{x}, \mathbf{y}\right) = 2\left(\|\mathbf{x} - \mathbf{y}\|_2^2 - d_0\right)$ and the Heaviside step function $\theta(x)$. The specific parameters used in this paper are given in table A.3.

Table A.3.: Parameters of the particle dimer system. Basic unit of energy and length are given as $k_\mathrm{B}T$ and r_m respectively.

Quantity	Symbol	Value	Unit
Number of particles	N	38	1
Box spring constant	k_box	100.0	$k_\mathrm{B}T/r_\mathrm{m}^2$
Box edge length	l_box	3.0	r_m
Repulsion strength	ε	1.	$k_\mathrm{B}T$
Mean squared dimer distance	d_0	1.5	r_m^2
Dimer potential slope	a	0	$k_\mathrm{B}T/r_\mathrm{m}^2$
Dimer quadratic term	b	-25	$\left(k_\mathrm{B}T/r_\mathrm{m}^2\right)^2$
Dimer quartic term	c	10	$\left(k_\mathrm{B}T/r_\mathrm{m}^2\right)^4$
Dimer spring constant	k_d	20	$k_\mathrm{B}T/r_\mathrm{m}^2$

A.2.3. Network parameters

The RNVP network consists of several subsequent blocks as depicted in Fig. 4.1. Each of these blocks consists of 4 independent networks, two for scaling and two for translation. All networks use leaky ReLU in each hidden layer. The output of the scaling networks uses a hyperbolic tangent scaled by a trainable scalar. The output of the translation networks is linear. Adam [198] is used as optimizer with standard parameters and a learning rate depending on the system. Table A.4 gives an overview of the exact network architectures and hyperparameters used in the experiments.

Table A.4.: Parameters of the RNVP networks used in the experiments

	DW	particles
Number of RNVP blocks	10	20
Hidden dimensions	$[20, 20, 20]$	$[76, 76, 76]$
Number of parameters	3.6×10^4	1.4×10^6
Number of training samples per bias and core	1×10^5	1×10^5
Bias strengths $/k_B T/r_\mathrm{m}^2$	$[10, 0]$	$[500, 10, 5, 2]$
Learning rate	10^{-3}	$[10^{-3}$,10^{-4},10^{-4},$10^{-5}]$
Batchsize	2000	8192

Appendix B.

Differential entropy difference

The differential entropy of a distribution p_X defined on a region $\Omega \subset \mathbb{R}^N$ is defined as

$$S_X = -k_B \int_\Omega p_X(x) \log p_X(x) \mathrm{d}x. \tag{B.1}$$

For a bijective function $y = f(x)$, the change of variable formula [Eq. (4.29)] can be applied to compute the change in differential entropy under the transformation. With the transformed density being $p_Y(y) = p_X(f^{-1}(y)) \left|\det J_{f^{-1}}(y)\right|$ this yields

$$\begin{aligned} S_X = -k_\mathrm{B} \int_{f(\Omega)} p_Y(y) \log p_Y(y) \left|\det J_{f^{-1}}(y)\right| \mathrm{d}y \\ = S_Y - k_\mathrm{B} \int_{f(\Omega)} p_Y(y) \log \left|\det J_{f^{-1}}(y)\right| \mathrm{d}y. \end{aligned} \tag{B.2}$$

Thus the difference in differential entropy under the transformation $f(\cdot)$ is given as

$$\Delta S = S_Y - S_X = -k_B \mathbb{E}_{x \sim p_\Omega} \left[\log \left|\det J_f(x)\right|\right], \tag{B.3}$$

where the inverse function theorem was used to compute the Jacobian.

Bibliography

[1] Miles D. Houslay. “Underpinning compartmentalised cAMP signalling through targeted cAMP breakdown”. In: *Trends in Biochemical Sciences* 35.2 (Feb. 2010), pp. 91–100. DOI: `10.1016/j.tibs.2009.09.007`.

[2] Fabian Paul et al. “Protein-peptide association kinetics beyond the seconds timescale from atomistic simulations”. In: *Nature Communications* 8.1 (2017), p. 1095. DOI: `10.1038/s41467-017-01163-6`.

[3] Duncan E. Scott et al. “Small molecules, big targets: Drug discovery faces the protein-protein interaction challenge”. In: *Nature Reviews Drug Discovery* 15.8 (2016), pp. 533–550. DOI: `10.1038/nrd.2016.29`.

[4] Nuria Plattner et al. “Complete protein–protein association kinetics in atomic detail revealed by molecular dynamics simulations and Markov modelling”. In: *Nature Chemistry* 9.10 (June 2017), pp. 1005–1011. DOI: `10.1038/nchem.2785`.

[5] Jacqueline Burré, Manu Sharma, and Thomas C. Südhof. “α-Synuclein assembles into higher-order multimers upon membrane binding to promote SNARE complex formation”. In: *Proceedings of the National academy of Sciences of the United States of America* 111.40 (Oct. 2014), E4274–E4283. DOI: `10.1073/pnas.1416598111`.

[6] Johannes Schöneberg et al. “Lipid-mediated PX-BAR domain recruitment couples local membrane constriction to endocytic vesicle fission”. In: *Nature Communications* 8.5 (2017), p. 15873. DOI: `10.1038/ncomms15873`.

[7] Marek Wieczorek et al. “Major histocompatibility complex (MHC) class I and MHC class II proteins: conformational plasticity in antigen presentation”. In: *Frontiers in Immunology* 8 (2017), p. 292. DOI: `10.3389/fimmu.2017.00292`.

[8] Robert Zwanzig. “From classical dynamics to continuous time random walks”. In: *Journal of Statistical Physics* 30.2 (1983), pp. 255–262. DOI: `10.1007/BF01012300`.

[9] Jan-Hendrik Prinz et al. “Markov models of molecular kinetics: Generation and Validation”. In: *The Journal of Chemical Physics* 134.17 (2011), p. 174105. DOI: `10.1063/1.3565032`.

[10] Dominique J Bicout and Attila Szabo. “Entropic barriers, transition states, funnels, and exponential protein folding kinetics: a simple model”. In: *Protein Science* 9.3 (2000), pp. 452–465. DOI: `10.1110/ps.9.3.452`.

[11] Hermann Grabert, Peter Hänggi, and Peter Talkner. “Microdynamics and nonlinear stochastic processes of gross variables”. In: *Journal of Statistical Physics* 22.5 (1980), pp. 537–552. DOI: `10.1007/BF01011337`.

[12] Robert Zwanzig. *Nonequilibrium statistical mechanics.* Oxford University Press, 2001. ISBN: 978-0195140187.

[13] Robert Zwanzig. “Nonlinear generalized Langevin equations”. In: *Journal of Statistical Physics* 9.3 (1973), pp. 215–220. DOI: `10.1007/BF01008729`.

[14] M. von Smoluchowski. “Versuch einer mathematischen Theorie der Koagulationskinetik kolloider Lösungen”. In: *Zeitschrift fur Physikalische Chemie* 92 (1917), pp. 129–168. DOI: `10.1007/BF01427232`.

[15] Hendrik Anthony Kramers. “Brownian motion in a field of force and the diffusion model of chemical reactions”. In: *Physica* 7.4 (1940), pp. 284–304. DOI: `10.1016/S0031-8914(40)90098-2`.

[16] Peter Debye. “Näherungsformeln für die Zylinderfunktionen für große Werte des Arguments und unbeschränkt veränderliche Werte des Index”. In: *Mathematische Annalen* 67.4 (1909), pp. 535–558. DOI: `10.1007/BF01450097`.

[17] Jacobus Henricus Van't Hoff and Jacobus Henricus Hoff. *Etudes de dynamique chimique.* F. Muller & Company, 1884.

[18] Svante Arrhenius. “Über die Reaktionsgeschwindigkeit bei der Inversion von Rohrzucker durch Säuren”. In: *Zeitschrift für physikalische Chemie* 4.1 (1889), pp. 226–248. DOI: `10.1515/zpch-1889-0416`.

[19] C. Schütte et al. “A Direct Approach to Conformational Dynamics based on Hybrid Monte Carlo”. In: *Journal of Computational Physics* 151 (1999), pp. 146–168. DOI: `10.1006/jcph.1999.6231`.

[20] B. Trendelkamp-Schroer et al. “Estimation and uncertainty of reversible Markov models”. In: *The Journal of Chemical Physics* 143 (2015), p. 174101. DOI: `10.1063/1.4934536`.

[21] Brooke E Husic and Vijay S Pande. “Markov state models: From an art to a science”. In: *Journal of the American Chemical Society* 140.7 (2018), pp. 2386–2396. DOI: `10.1021/jacs.7b12191`.

[22] Stefanie Winkelmann and Christof Schütte. *Stochastic Dynamics in Computational Biology.* Springer, 2020. ISBN: 978-3-030-62387-6. DOI: `10.1007/978-3-030-62387-6`.

[23] Felix Höfling and Thomas Franosch. “Anomalous transport in the crowded world of biological cells”. In: *Reports on Progress in Physics* 76.4 (2013), p. 046602. DOI: `10.1088/0034-4885/76/4/046602`.

[24] Frank C. Collins and George E. Kimball. “Diffusion-controlled reaction rates”. In: *Journal of Colloid Science* 4.4 (1949), pp. 425–437. DOI: `10.1016/0095-8522(49)90023-9`.

[25] D. Shoup and A. Szabo. “Role of diffusion in ligand binding to macromolecules and cell-bound receptors”. In: *Biophysical Journal* 40.1 (Oct. 1982), pp. 33–39. DOI: `10.1016/S0006-3495(82)84455-X`.

[26] Attila Szabo, Klaus Schulten, and Zan Schulten. “First passage time approach to diffusion controlled reactions”. In: *The Journal of Chemical Physics* 72.8 (1980), pp. 4350–4357. DOI: `10.1063/1.439715`.

[27] Margarita Kostré et al. “Coupling particle-based reaction-diffusion simulations with reservoirs mediated by reaction-diffusion PDEs”. 2020. arXiv: `2006.00003`.

[28] P. Debye. “Reaction Rates in Ionic Solutions”. In: *Journal of the Electrochemical Society* 82.1 (1942), p. 265. DOI: `10.1149/1.3071413`.

[29] Frank C. Collins and George E. Kimball. “Diffusion-Controlled Reactions in Liquid Solutions”. In: *Industrial and Engineering Chemistry Research* 41.November (1949), pp. 2551–2553. DOI: `10.1021/ie50479a040`.

[30] Ei Teramoto and Nanako Shigesada. “Theory of Bimolecular Reaction Processes in Liquids”. In: *Progress of Theoretical Physics* 37.1 (1967), pp. 29–51. DOI: `10.1143/PTP.37.29`.

[31] Masao Doi. “Theory of diffusion-controlled reactions between non-simple molecules. I”. In: *Chemical Physics* ll.1 (1975), pp. 107–113. DOI: `10.1016/0301-0104(75)80043-7`.

[32] Masao Doi. “Stochastic theory of diffusion-controlled reaction”. In: *Journal of Physics A: Mathematical and Theoretical* 9 (1976). DOI: `10.1088/0305-4470/9/9/009`.

[33] Masao Doi. “Theory of diffusion-controlled reactions between non-simple molecules. II”. In: *Chemical Physics* 3 (1975), pp. 115–121. DOI: `10.1016/0301-0104(75)80044-9`.

[34] Daniel T. Gillespie. “Stochastic Simulation of Chemical Kinetics”. In: *Annual Review of Physical Chemistry* 58.1 (2007), pp. 35–55. DOI: `10.1146/annurev.physchem.58.032806.104637`.

[35] CM Guldberg and P Waage. *Studier over affiniteten.* Forhandlinger I Videnskabsselskabet I Christiania, 1864.

[36] Leonor Michaelis and Maud L Menten. “Die kinetik der invertinwirkung”. In: *Biochemische Zeitschrift* 49.333-369 (1913), p. 352.

[37] Christoph Fröhner and Frank Noé. “Reversible Interacting-Particle Reaction Dynamics”. In: *The Journal of Physical Chemistry B* 122.49 (2018), pp. 11240–11250. DOI: `10.1021/acs.jpcb.8b06981`.

[38] Johannes Schöneberg, Alexander Ullrich, and Frank Noé. “Simulation tools for particle-based reaction-diffusion dynamics in continuous space”. In: *BMC Biophysics* 7.1 (2014), p. 11. DOI: `10.1186/s13628-014-0011-5`.

[39] Manuel Dibak et al. “MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations”. In: *The Journal of Chemical Physics* 148.21 (2018), p. 214107. DOI: `10.1063/1.5020294`.

[40] A. Warshel and M. Levitt. “Theoretical studies of enzymic reactions: Dielectric, electrostatic and steric stabilization of the carbonium ion in the reaction of lysozyme”. In: *Journal of Molecular Biology* 103 (1976), pp. 227–249. DOI: `10.1016/0022-2836(76)90311-9`.

[41] M. Shirts and V. S. Pande. “Screen Savers of the World Unite!” In: *Science* 290 (2000), pp. 1903–1904. DOI: `10.1126/science.290.5498.1903`.

[42] I. Buch et al. “High-throughput all-atom molecular dynamics simulations using distributed computing”. In: *Journal of Chemical Information and Modeling* 50 (2010), pp. 397–403. DOI: `10.1021/ci900455r`.

[43] D. E. Shaw et al. “Atomic-Level Characterization of the Structural Dynamics of Proteins”. In: *Science* 330 (2010), pp. 341–346. DOI: `10.1126/science.1187409`.

[44] S. Doerr et al. “HTMD: High-Throughput Molecular Dynamics for Molecular Discovery”. In: *Journal of Chemical Theory and Computation* 12 (2016), pp. 1845–1852. DOI: `10.1021/acs.jctc.6b00049`.

[45] G. R. Bowman, V. S. Pande, and F. Noé, eds. *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation.* Vol. 797. Advances in Experimental Medicine and Biology. Springer Heidelberg, 2014.

[46] M. Sarich and C. Schütte. *Metastability and Markov State Models in Molecular Dynamics.* Courant Lecture Notes. American Mathematical Society, 2013.

[47] Frank Noé et al. “Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations”. In: *Proceedings of the National Academy of Sciences* 106.45 (Nov. 2009), pp. 19011–19016. DOI: `10.1073/pnas.0905466106`.

[48] G. R. Bowman et al. “Progress and challenges in the automated construction of Markov state models for full protein systems.” In: *The Journal of Chemical Physics* 131 (2009), p. 124101. DOI: `10.1063/1.3216567`.

[49] K. Lindorff-Larsen et al. “How Fast-Folding Proteins Fold”. In: *Science* 334 (2011), pp. 517–520. DOI: `10.1126/science.1208351`.

[50] K. J. Kohlhoff et al. “Cloud-based simulations on Google Exacycle reveal ligand modulation of GPCR activation pathways”. In: *Nature Chemistry* 6 (2014), pp. 15–21. DOI: `10.1038/nchem.1821`.

[51] I. Buch, T. Giorgino, and G. De Fabritiis. “Complete reconstruction of an enzyme-inhibitor binding process by molecular dynamics simulations”. In: *Proceedings of the National Academy of Sciences of the United States of America* 108 (2011), pp. 10184–10189. DOI: `10.1073/pnas.1103547108`.

[52] D.-A. Silva et al. “A Role for Both Conformational Selection and Induced Fit in Ligand Binding by the LAO Protein”. In: *PLoS Computational Biology* 7 (2011), e1002054. DOI: `10.1371/journal.pcbi.1002054`.

[53] N. Plattner and F. Noé. “Protein conformational plasticity and complex ligand binding kinetics explored by atomistic simulations and Markov models”. In: *Nature Communications* 6 (2015), p. 7653. DOI: `10.1038/ncomms8653`.

[54] D. De Sancho et al. "Identification of mutational hot spots for substrate diffusion: Application to myoglobin". In: *Journal of Chemical Theory and Computation* 11.4 (2015), pp. 1919–1927. DOI: `10.1021/ct5011455`.

[55] Adam Kubas et al. "Mechanism of O2 diffusion and reduction in FeFe hydrogenases". In: *Nature Chemistry* 9.1 (Aug. 2016), pp. 88–95. DOI: `10.1038/nchem.2592`.

[56] N. Plattner et al. "Complete protein–protein association kinetics in atomic detail revealed by molecular dynamics simulations and Markov modelling". In: *Nature Chemistry* 9 (2017), pp. 1005–1011. DOI: `10.1038/nchem.2785`.

[57] H. Wu et al. "Statistically optimal analysis of state-discretized trajectory data from multiple thermodynamic states". In: *The Journal of Chemical Physics* 141 (2014), p. 214106. DOI: `10.1063/1.4902240`.

[58] E. Rosta and G. Hummer. "Free energies from dynamic weighted histogram analysis using unbiased Markov state model". In: *Journal of Chemical Theory and Computation* 11 (2015), pp. 276–285. DOI: `10.1021/ct500719p`.

[59] Hao Wu et al. "Multiensemble Markov models of molecular thermodynamics and kinetics". In: *Proceedings of the National academy of Sciences of the United States of America* 113.23 (2016), E3221–E3230. DOI: `10.1073/pnas.1525092113`.

[60] A. S. J. S. Mey, H. Wu, and F. Noé. "xTRAM: Estimating equilibrium expectations from time-correlated simulation data at multiple thermodynamic states". In: *Physical Review X* 4 (2014), p. 041018. DOI: `10.1103/PhysRevX.4.041018`.

[61] R. Casasnovas et al. "Unbinding kinetics of a p38 MAP kinase type II inhibitor from metadynamics simulations". In: *Journal of the American Chemical Society* 139 (2017), pp. 4780–4788. DOI: `10.1021/jacs.6b12950`.

[62] H. Qian. "Cellular biology in terms of stochastic nonlinear biochemical dynamics: Emergent properties, isogenetic variations and chemical system inheritability". In: *The Journal of Statistical Physics* 141.6 (2010), pp. 990–1013. DOI: `10.1007/s10955-010-0093-7`.

[63] Radek Erban and S Jonathan Chapman. "Stochastic modelling of reaction–diffusion processes: algorithms for bimolecular reactions". In: *Physical Biology* 6.4 (2009), p. 046001. DOI: `10.1088/1478-3975/6/4/046001`.

[64] D. Fange et al. "Stochastic reaction-diffusion kinetics in the microscopic limit". In: *Proceedings of the National Academy of Sciences of the United States of America* 107.46 (2010), pp. 19820–19825. DOI: `10.1073/pnas.1006565107`.

[65] K. Takahashi, S. Tănase-Nicola, and P. R. ten Wolde. "Spatio-temporal correlations can drastically change the response of a MAPK pathway". In: *Proceedings of the National Academy of Sciences of the United States of America* 107.6 (2010), pp. 2473–2478. DOI: `10.1073/pnas.0906885107`.

[66] S. S. Andrews and D. Bray. “Stochastic simulation of chemical reactions with spatial resolution and single molecule detail”. In: *Physical Biology* 1.3 (2004), p. 137. DOI: `10.1088/1478-3967/1/3/001`.

[67] J. Biedermann et al. “ReaDDyMM: Fast interacting particle reaction-diffusion simulations using graphical processing units”. In: *Biophysical Journal* 108 (2015), pp. 457–461. DOI: `10.1016/j.bpj.2014.12.025`.

[68] A. Donev et al. “A first-passage kinetic Monte Carlo algorithm for complex diffusion–reaction systems”. In: *Journal of Computational Physics* 229.9 (2010), pp. 3214–3236. DOI: `10.1016/j.jcp.2009.12.038`.

[69] Moritz Hoffmann, Christoph Fröhner, and Frank Noé. “ReaDDy 2: Fast and flexible software framework for interacting-particle reaction dynamics”. In: *PLoS Computational Biology* 15.2 (2019), e1006830. DOI: `10.1371/journal.pcbi.1006830`.

[70] Aleksandar Donev, Chiao-Yu Yang, and Changho Kim. “Efficient reactive Brownian dynamics”. In: *The Journal of Chemical Physics* 148.3 (2018), p. 034103. DOI: `10.1063/1.5009464`.

[71] J. Hattne, D. Fange, and J. Elf. “Stochastic Reaction-Diffusion Simulation with MesoRD”. In: *Bioinformatics* 21.12 (2005), pp. 2923–2924. DOI: `10.1093/bioinformatics/bti431`.

[72] J. Schöneberg and F. Noé. “ReaDDy - a software for particle based reaction diffusion dynamics in crowded cellular environments.” In: *PLOS ONE* 8.e74261 (2013). DOI: `10.1371/journal.pone.0074261`.

[73] J. S. van Zon and P. R. Ten Wolde. “Green’s-function reaction dynamics: A particle-based approach for simulating biochemical networks in time and space”. In: *The Journal of Chemical Physics* 123.23 (2005), p. 4910. DOI: `10.1063/1.2137716`.

[74] J. S. Van Zon and P. R. ten Wolde. “Simulating biochemical networks at the particle level in time and space: Green’s Function Reaction Dynamics”. In: *Physical Review Letters* 94 (2005), p. 128103. DOI: `10.1103/PhysRevLett.94.128103`.

[75] P. Mereghetti et al. “Diffusion and association processes in biological systems: theory, computation and experiment”. In: *BMC biophysics* 4.1 (2011), p. 2. DOI: `10.1186/2046-1682-4-2`.

[76] Mauricio J. del Razo et al. “Multiscale molecular kinetics by coupling Markov state models and reaction-diffusion dynamics”. In: *The Journal of Chemical Physics* 155.12 (2021), p. 124109. DOI: `10.1063/5.0060314`.

[77] Luigi Sbailò and Frank Noé. “An efficient multi-scale Green’s function reaction dynamics scheme”. In: *The Journal of Chemical Physics* 147.18 (2017), p. 184106. DOI: `10.1063/1.5010190`.

[78] A. Vijaykumar, P. G. Bolhuis, and P. R. ten Wolde. “Combining molecular dynamics with mesoscopic Green’s function reaction dynamics simulations”. In: *The Journal of Chemical Physics* 143.21 (2015), p. 214102. DOI: `10.1063/1.4936254`.

[79] L. W. Votapka et al. “SEEKR: Simulation Enabled Estimation of Kinetic Rates, A Computational Tool to Estimate Molecular Kinetics and Its Application to Trypsin–Benzamidine Binding”. In: *The Journal of Physical Chemistry B* 121 (2017), pp. 3597–3606. DOI: `10.1021/acs.jpcb.6b09388`.

[80] Anton K Faradjian and Ron Elber. “Computing time scales from reaction coordinates by milestoning”. In: *The Journal of Chemical Physics* 120.23 (2004), pp. 10880–10889. DOI: `10.1063/1.1738640`.

[81] Mauricio J. del Razo and Hong Qian. “A discrete stochastic formulation for reversible bimolecular reactions via diffusion encounter”. In: *Communications in Mathematical Sciences* 14.6 (2016), pp. 1741–1772. DOI: `10.4310/CMS.2016.v14.n6.a13`.

[82] M. B. Flegg, S. J. Chapman, and R. Erban. “The two-regime method for optimizing stochastic reaction–diffusion simulations”. In: *Journal of the Royal Society Interface* 9.70 (2012), pp. 859–868. DOI: `10.1098/rsif.2011.0574`.

[83] M. B. Flegg, S. Hellander, and Radek R. Erban. “Convergence of methods for coupling of microscopic and mesoscopic reaction–diffusion simulations”. In: *Journal of Computational Physics* 289 (2015), pp. 1–17. DOI: `10.1016/j.jcp.2015.01.030`.

[84] S. Donnini et al. “Constant pH molecular dynamics in explicit solvent with lambda-dynamics”. In: *Journal of Chemical Theory and Computation* 7 (2011), pp. 1962–1978. DOI: `10.1021/ct200061r`.

[85] F. Noé and F. Nüske. “A variational approach to modeling slow processes in stochastic dynamical systems”. In: *Multiscale Modeling and Simulation* 11 (2013), pp. 635–655. DOI: `10.1137/110858616`.

[86] F. Nüske et al. “Variational Approach to Molecular Kinetics”. In: *Journal of Chemical Theory and Computation* 10 (2014), pp. 1739–1752. DOI: `10.1021/ct4009156`.

[87] G. Perez-Hernandez et al. “Identification of slow molecular order parameters for Markov model construction”. In: *The Journal of Chemical Physics* 139 (2013), p. 015102. DOI: `10.1063/1.4811489`.

[88] C. R. Schwantes and V. S. Pande. “Improvements in Markov State Model Construction Reveal Many Non-Native Interactions in the Folding of NTL9”. In: *Journal of Chemical Theory and Computation* 9 (2013), pp. 2000–2009. DOI: `10.1021/ct300878a`.

[89] F. Noé and C. Clementi. “Collective variables for the study of long-time kinetics from molecular trajectories: theory and methods”. In: *Current Opinion in Structural Biology.* 43 (2017), pp. 141–147. DOI: `10.1016/j.sbi.2017.02.006`.

[90] S. Klus et al. “Data-driven model reduction and transfer operator approximation”. In: *Journal of Nonlinear Science* (2017), pp. 1–26. DOI: `10.1007/s00332-017-9437-7`.

[91] F. Noé and C. Clementi. “Kinetic distance and kinetic maps from molecular dynamics simulation”. In: *Journal of Chemical Theory and Computation* 11 (2015), pp. 5002–5011. DOI: `10.1021/acs.jctc.5b00553`.

[92] F. Noé, R. Banisch, and C. Clementi. “Commute maps: separating slowly-mixing molecular configurations for kinetic modeling”. In: *Journal of Chemical Theory and Computation* 12 (2016), pp. 5620–5630. DOI: `10.1021/acs.jctc.6b00762`.

[93] M. K. Scherer et al. “PyEMMA 2: A software package for estimation, validation and analysis of Markov models”. In: *Journal of Chemical Theory and Computation* 11 (2015), pp. 5525–5542. DOI: `10.1021/acs.jctc.5b00743`.

[94] M. P. Harrigan et al. “MSMBuilder: Statistical Models for Biomolecular Dynamics”. In: *Biophysical Journal* 112 (2017), pp. 10–15. DOI: `10.1101/084020`.

[95] N. V. Buchete and G. Hummer. “Coarse Master Equations for Peptide Folding Dynamics”. In: *The Journal of Physical Chemistry B* 112 (2008), pp. 6057–6069. DOI: `10.1021/jp0761665`.

[96] S. Kube and M. Weber. “A coarse graining method for the identification of transition rates between molecular conformations”. In: *The Journal of Chemical Physics* 126 (2007), p. 024103. DOI: `10.1063/1.2404953`.

[97] G. Hummer and A. Szabo. “Optimal Dimensionality Reduction of Multistate Kinetic and Markov-State Models”. In: *The Journal of Physical Chemistry B* 119 (2015), pp. 9029–9037. DOI: `10.1021/jp508375q`.

[98] S. Orioli and P. Faccioli. “Dimensional reduction of Markov state models from renormalization group theory”. In: *The Journal of Chemical Physics* 145 (2016), p. 124120. DOI: `10.1063/1.4963196`.

[99] F. Noé et al. “Projected and Hidden Markov Models for calculating kinetics and metastable states of complex molecules”. In: *The Journal of Chemical Physics* 139 (2013), p. 184114. DOI: `10.1063/1.4828816`.

[100] A. Mardt et al. “VAMPnets for deep learning of molecular kinetics”. In: *Nature Communications* (2018), p. 5. DOI: `10.1038/s41467-017-02388-1`.

[101] C. Schütte et al. “Markov state models based on milestoning”. In: *The Journal of Chemical Physics* 134.20 (2011), 05B609. DOI: `10.1063/1.3590108`.

[102] J. Schluttig et al. “Dynamics of protein-protein encounter: A Langevin equation approach with reaction patches”. In: *The Journal of Chemical Physics* 129 (2008), p. 155106. DOI: `10.1063/1.2996082`.

[103] A. Vijaykumar et al. “Multiscale simulations of anisotropic particles combining Brownian Dynamics and Green’s Function Reaction Dynamics”. In: *The Journal of Chemical Physics* 146.11 (2017), p. 114106. DOI: `10.1063/1.4977515`.

[104] Doris A Schuetz et al. “Kinetics for drug discovery: an industry-driven effort to target drug residence time”. In: *Drug Discovery Today* 22.6 (2017), pp. 896–911. DOI: `10.1016/j.drudis.2017.02.002`.

[105] D. J. Higham. “An algorithmic introduction to numerical simulation of stochastic differential equations”. In: *SIAM review* 43.3 (2001), pp. 525–546. DOI: `10.1137/S0036144500378302`.

[106] P. Leopardi. “A partition of the unit sphere into regions of equal area and small diameter”. In: *Electronic Transactions on Numerical Analysis* 25.12 (2006), pp. 309–327.

[107] M. Ester et al. “A density-based algorithm for discovering clusters in large spatial databases with noise.” In: *Kdd*. Vol. 96. 34. 1996, pp. 226–231. DOI: `10.5120/739-1038`.

[108] D. L. Wise and G. Houghton. “Diffusion coefficients of neon, krypton, xenon, carbon monoxide and nitric oxide in water at 10–60 °C”. In: *Chemical Engineering Science* 23.10 (1968), pp. 1211–1216. DOI: `0.1142/S0217979216502052`.

[109] T. E. Carver et al. “Analysis of the kinetic barriers for ligand binding to sperm whale myoglobin using site-directed mutagenesis and laser photolysis techniques.” In: *Journal of Biological Chemistry* 265.32 (1990), pp. 20007–20020. DOI: `10.1016/S0021-9258(17)45475-5`.

[110] Thomas C. Südhof. “Neurotransmitter Release: The Last Millisecond in the Life of a Synaptic Vesicle”. In: *Neuron* 80 (2013), pp. 675–690. DOI: `10.1016/j.neuron.2013.10.022`.

[111] Manuel Dibak et al. “Diffusion-influenced reaction rates in the presence of pair interactions”. In: *The Journal of Chemical Physics* 151.16 (2019), p. 164105. DOI: `10.1063/1.5124728`.

[112] Johannes Schöneberg and Frank Noé. “ReaDDy–a software for particle-based reaction-diffusion dynamics in crowded cellular environments.” In: *PLOS ONE* 8.9 (2013), e74261. DOI: `10.1371/journal.pone.0074261`.

[113] Gordon D Smith. *Numerical solution of partial differential equations: finite difference methods*. Oxford university press, 1985.

[114] Peter Poier et al. “An Anisotropic Effective Model for the Simulation of Semiflexible Ring Polymers”. In: *Macromolecules* 48.14 (2015), pp. 4983–4997. DOI: `10.1021/acs.macromol.5b00603`.

[115] Ignasi Buch, S Kashif Sadiq, and Gianni De Fabritiis. “Optimized potential of mean force calculations for standard binding free energies”. In: *Journal of Chemical Theory and Computation* 7.6 (2011), pp. 1765–1772. DOI: `10.1021/ct2000638`.

[116] Xiao Xu et al. “Interaction of Proteins with Polyelectrolytes: Comparison of Theory to Experiment”. In: *Langmuir* 35.16 (2019), pp. 5373–5391. DOI: `10.1021/acs.langmuir.8b01802`.

[117] D. Hinzke and U. Nowak. “Monte Carlo simulation of magnetization switching in a Heisenberg model for small ferromagnetic particles”. In: *Computer Physics Communications* 121-122 (Sept. 1999), pp. 334–337. DOI: `10.1016/s0010-4655(99)00348-3`.

[118] K. Binder, H. Rauch, and V. Wildpaner. "Monte Carlo calculation of the magnetization of superparamagnetic particles". In: *Journal of Physics and Chemistry of Solids* 31.2 (Feb. 1970), pp. 391–397. DOI: `10.1016/0022-3697(70)90119-8`.

[119] Deborah K. Jones-Hertzog and William L. Jorgensen. "Binding Affinities for Sulfonamide Inhibitors with Human Thrombin Using Monte Carlo Simulations with a Linear Response Method". In: *Journal of Medicinal Chemistry* 40.10 (May 1997), pp. 1539–1549. DOI: `10.1021/jm960684e`.

[120] Jonathan W. Essex et al. "Monte Carlo Simulations for Proteins: Binding Affinities for Trypsin-Benzamidine Complexes via Free-Energy Perturbations". In: *The Journal of Physical Chemistry B* 101.46 (Nov. 1997), pp. 9663–9669. DOI: `10.1021/jp971990m`.

[121] Steve Brooks et al. *Handbook of markov chain monte carlo*. CRC press, 2011. ISBN: 978-1-4200-7941-8.

[122] W. K. Hastings. "Monte Carlo sampling methods using Markov chains and their applications". In: *Biometrika* 57.1 (1970), pp. 97–109. URL: `http://www.jstor.org/stable/2334940`.

[123] Jun S Liu. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008.

[124] S. Agapiou et al. "Importance Sampling: Intrinsic Dimension and Computational Cost". In: *Statistical Science* 32.3 (2017), pp. 405–431. DOI: `10.1214/17-STS611`.

[125] Thomas Bengtsson, Peter Bickel, Bo Li, et al. "Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems". In: *Probability and statistics: Essays in honor of David A. Freedman*. Institute of Mathematical Statistics, 2008, pp. 316–334. DOI: `10.1214/193940307000000518`.

[126] Esa Nummelin. *General irreducible Markov chains and non-negative operators*. 83. Cambridge University Press, 2004.

[127] Nicholas Metropolis et al. "Equation of state calculations by fast computing machines". In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092. DOI: `10.1063/1.1699114`.

[128] Alan Sokal. "Monte Carlo methods in statistical mechanics: foundations and new algorithms". In: *Functional integration*. Springer, 1997, pp. 131–192. DOI: `10.1007/978-1-4899-0319-8_6`.

[129] K Kikuchi et al. "Metropolis Monte Carlo method as a numerical technique to solve the Fokker—Planck equation". In: *Chemical Physics Letters* 185.3-4 (1991), pp. 335–338. DOI: `10.1016/S0009-2614(91)85070-D`.

[130] E Marinari and G Parisi. "Simulated Tempering: A New Monte Carlo Scheme". In: *EPL* 19.6 (July 1992), pp. 451–458. DOI: `10.1209/0295-5075/19/6/002`.

[131] Charles J Geyer and Elizabeth A Thompson. “Annealing Markov chain Monte Carlo with applications to ancestral inference”. In: *Journal of the American Statistical Association* 90.431 (1995), pp. 909–920. DOI: `10.1080/01621459.1995.10476590`.

[132] Robert H. Swendsen and Jian-Sheng Wang. “Replica Monte Carlo Simulation of Spin-Glasses”. In: *Physical Review Letters* 57 (21 Nov. 1986), pp. 2607–2609. DOI: `10.1103/PhysRevLett.57.2607`.

[133] Charles J Geyer. “Markov chain Monte Carlo maximum likelihood”. In: *Computing Science and Statistics: Proceedings of the 23rd Symposiumon on the Interface*. Interface Foundation of North America, 1991. DOI: `11299/58440`.

[134] Koji Hukushima and Koji Nemoto. “Exchange Monte Carlo method and application to spin glass simulations”. In: *Journal of the Physical Society of Japan* 65.6 (1996), pp. 1604–1608. DOI: `10.1143/JPSJ.65.1604`.

[135] Andrew J. Ballard and Christopher Jarzynski. “Replica exchange with nonequilibrium switches”. In: *Proceedings of the National Academy of Sciences of the United States of America* 106.30 (2009), pp. 12224–12229. DOI: `10.1073/pnas.0900406106`.

[136] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi. “Optimization by Simulated Annealing”. In: (1987). Ed. by Martin A. Fischler and Oscar Firschein, pp. 606–615. DOI: `10.1016/B978-0-08-051581-6.50059-3`.

[137] K. Hukushima and Y. Iba. “Population Annealing and Its Application to a Spin Glass”. In: *AIP Conference Proceedings*. Vol. 690. 1. American Institute of Physics. 2003, pp. 200–206. DOI: `10.1063/1.1632130`.

[138] Wenlong Wang, Jonathan Machta, and Helmut G Katzgraber. “Population annealing: Theory and application in spin glasses”. In: *Physical Review E* 92.6 (2015), p. 063307. DOI: `10.1063/1.1632130`.

[139] Christopher Amey and Jonathan Machta. “Analysis and optimization of population annealing”. In: *Physical Review E* 97 (3 Mar. 2018), p. 033301. DOI: `10.1103/PhysRevE.97.033301`.

[140] Hiroaki Fukunishi, Osamu Watanabe, and Shoji Takada. “On the Hamiltonian replica exchange method for efficient sampling of biomolecular systems: Application to protein structure prediction”. In: *The Journal of Chemical Physics* 116.20 (2002), pp. 9058–9067. DOI: `10.1063/1.1472510`.

[141] Harry A. Stern. “Molecular simulation with variable protonation states at constant pH”. In: *The Journal of Chemical Physics* 126.16 (2007), p. 164112. DOI: `10.1063/1.2731781`.

[142] Yunjie Chen and Benoît Roux. “Constant-pH Hybrid Nonequilibrium Molecular Dynamics Monte Carlo Simulation Method”. In: *Journal of Chemical Theory and Computation* 11.8 (2015). PMID: 26300709, pp. 3919–3931. DOI: `10.1021/acs.jctc.5b00261`.

[143] Jerome P. Nilmeier et al. “Nonequilibrium candidate Monte Carlo is an efficient tool for equilibrium simulation”. In: *Proceedings of the National Academy of Sciences of the United States of America* 108.45 (2011), E1009–E1018. DOI: `10.1073/pnas.1106094108`.

[144] Daan Frenkel. “Speed-up of Monte Carlo simulations by sampling of rejected states”. In: *Proceedings of the National Academy of Sciences* 101.51 (2004), pp. 17571–17575. DOI: `10.1073/pnas.0407950101`.

[145] Gavin E. Crooks. “Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences”. en. In: *Physical Review E* 60.3 (Sept. 1999), pp. 2721–2726. DOI: `10.1103/PhysRevE.60.2721`.

[146] Hao Wu, Jonas Köhler, and Frank Noé. “Stochastic Normalizing Flows”. 2020. arXiv: `2002.06707`.

[147] Alessandro Laio and Michele Parrinello. “Escaping free-energy minima”. In: *Proceedings of the National academy of Sciences of the United States of America* 99.20 (2002), pp. 12562–12566. DOI: `10.1073/pnas.202427399`.

[148] Jun Zhang, Yi Isaac Yang, and Frank Noé. “Targeted Adversarial Learning Optimized Sampling”. In: *The Journal of Physical Chemistry Letters* 10.19 (2019). PMID: 31522495, pp. 5791–5797. DOI: `10.1021/acs.jpclett.9b02173`.

[149] Glenn M Torrie and John P Valleau. “Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling”. In: *Journal of Computational Physics* 23.2 (1977), pp. 187–199. DOI: `10.1016/0021-9991(77)90121-8`.

[150] Simon Duane et al. “Hybrid monte carlo”. In: *Physics letters B* 195.2 (1987), pp. 216–222. DOI: `10.1016/0370-2693(87)91197-X`.

[151] Roger W Hockney and James W Eastwood. *Computer simulation using particles.* CRC Press, 2021. DOI: `10.1201/9780367806934`.

[152] Ioan Andricioaei, John E. Straub, and Arthur F. Voter. “Smart Darting Monte Carlo”. In: *The Journal of Chemical Physics* 114.16 (2001), pp. 6994–7000. DOI: `10.1063/1.1358861`.

[153] K. Roberts, R. Sebsebie, and E. Curotto. “A rare event sampling method for diffusion Monte Carlo using smart darting”. In: *The Journal of Chemical Physics* 136.7 (2012), p. 074104. DOI: `10.1063/1.3685453`.

[154] Lionel Walter and Marcus Weber. *ConfJump: a fast biomolecular sampling method which drills tunnels through high mountains.* eng. Tech. rep. 06-26. Takustr. 7, 14195 Berlin: ZIB, 2006.

[155] Huitao Shen, Junwei Liu, and Liang Fu. “Self-learning Monte Carlo with deep neural networks”. In: *Physical Review B* 97 (20 May 2018), p. 205140. DOI: `10.1103/PhysRevB.97.205140`.

[156] Raza Habib and David Barber. “Auxiliary Variational MCMC”. In: *International Conference on Learning Representations.* 2019. URL: `https://openreview.net/forum?id=r1NJqsRctX`.

[157] Jiaming Song, Shengjia Zhao, and Stefano Ermon. "A-NICE-MC: Adversarial Training for MCMC". In: *Advances in Neural Information Processing Systems 30*. Ed. by I. Guyon et al. Curran Associates, Inc., 2017, pp. 5140–5150. URL: `http://papers.nips.cc/paper/7099-a-nice-mc-adversarial-training-for-mcmc.pdf`.

[158] Daniel Levy, Matt D. Hoffman, and Jascha Sohl-Dickstein. "Generalizing Hamiltonian Monte Carlo with Neural Networks". In: *International Conference on Learning Representations*. 2018. URL: `https://openreview.net/forum?id=B1n8LexRZ`.

[159] Frank Noé et al. "Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning". In: *Science* 365.6457 (2019). DOI: `10.1126/science.aaw1147`.

[161] Danilo Jimenez Rezende and Shakir Mohamed. "Variational Inference with Normalizing Flows". 2016. arXiv: `1505.05770`.

[162] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. "Density estimation using Real NVP". 2017. arXiv: `1605.08803`.

[163] Conor Durkan et al. "Neural Spline Flows". In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach et al. Vol. 32. Curran Associates, Inc., 2019. URL: `https://proceedings.neurips.cc/paper/2019/file/7ac71d433f282034e088473244df8c02-Paper.pdf`.

[164] Durk P Kingma and Prafulla Dhariwal. "Glow: Generative Flow with Invertible 1x1 Convolutions". In: *Advances in Neural Information Processing Systems 31*. Ed. by S. Bengio et al. Curran Associates, Inc., 2018, pp. 10215–10224. URL: `http://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions.pdf`.

[165] Conor Durkan et al. "Neural Spline Flows". In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach et al. Curran Associates, Inc., 2019, pp. 7511–7522. URL: `http://papers.nips.cc/paper/8969-neural-spline-flows.pdf`.

[166] Esteban G. Tabak and Eric Vanden-Eijnden. "Density estimation by dual ascent of the log-likelihood". In: *Communications in Mathematical Sciences* 8 (2010), pp. 217–233. DOI: `10.4310/CMS.2010.v8.n1.a11`.

[167] George Papamakarios et al. "Normalizing Flows for Probabilistic Modeling and Inference". 2021. arXiv: `1912.02762`.

[168] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. "Normalizing flows: An introduction and review of current methods". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020). DOI: `10.1109/TPAMI.2020.2992934`.

[169] Solomon Kullback and Richard A Leibler. "On information and sufficiency". In: *The annals of mathematical statistics* 22.1 (1951), pp. 79–86. DOI: `10.1214/aoms/1177729694`.

[170] Luigi Sbailò, Manuel Dibak, and Frank Noé. "Neural mode jump Monte Carlo". In: *The Journal of Chemical Physics* 154.7 (2021), p. 074101. DOI: `10.1063/5.0032346`.

[171] Daniel L Parton et al. "Ensembler: Enabling high-throughput molecular simulations at the superfamily scale". In: *PLoS computational biology* 12.6 (2016), e1004728. DOI: `10.1371/journal.pcbi.1004728`.

[172] Johan Ludwig William Valdemar Jensen. "Sur les fonctions convexes et les inégalités entre les valeurs moyennes". In: *Acta Mathematica* 30.1 (1906), pp. 175–193. DOI: `10.1007/BF02418571`.

[173] Ricky T. Q. Chen et al. "Neural Ordinary Differential Equations". In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio et al. Vol. 31. Curran Associates, Inc., 2018. URL: `https://proceedings.neurips.cc/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf`.

[174] James MacQueen et al. "Some methods for classification and analysis of multivariate observations". In: *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*. Vol. 1. 14. Oakland, CA, USA. 1967, pp. 281–297.

[175] M. G. Voronoi. "Nouvelles applications des parametres continus a la theorie des formes quadratiques". In: *Journal für die reine und angewandte Mathematik* 134 (1908), pp. 198–287. DOI: `10.1515/crll.1908.133.97`.

[176] H. W. Kuhn. "The Hungarian method for the assignment problem". In: *Naval Research Logistics* 2.12 (1955), pp. 83–97. DOI: `10.1002/nav.3800020109`.

[177] Michael R Shirts and John D Chodera. "Statistically optimal analysis of samples from multiple equilibrium states". In: *The Journal of Chemical Physics* 129.12 (2008), p. 124105. DOI: `10.1063/1.2978177`.

[178] Jan Hermann, Zeno Schätzle, and Frank Noé. "Deep-neural-network solution of the electronic Schrödinger equation". In: *Nature Chemistry* 12.10 (Sept. 2020), pp. 891–897. DOI: `10.1038/s41557-020-0544-y`.

[179] Kristof Schütt et al. "Schnet: A continuous-filter convolutional neural network for modeling quantum interactions". In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS'17. Long Beach, California, USA, 2017, pp. 991–1001.

[180] Nathaniel Thomas et al. "Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds". 2018. arXiv: `1802.08219`.

[181] Jonas Köhler, Leon Klein, and Frank Noé. "Equivariant flows: exact likelihood generative learning for symmetric densities". In: *International Conference on Machine Learning*. PMLR. 2020, pp. 5361–5370. URL: `https://proceedings.mlr.press/v119/kohler20a.html`.

[182] Laurent Dinh, David Krueger, and Yoshua Bengio. "NICE: Non-linear Independent Components Estimation". 2015. arXiv: `1410.8516`.

[183] S. Li and Lei Wang. "Neural Network Renormalization Group". In: *Physical Review Letters* 121 26 (2018), p. 260601. DOI: `10.1103/PhysRevLett.121.260601`.

[184] Kim A. Nicoli et al. "Asymptotically unbiased estimation of physical observables with neural samplers." In: *Physical Review E* 101 2-1 (2020), p. 023304. DOI: `10.1103/PhysRevE.101.023304`.

[185] MS Albergo, G Kanwar, and PE Shanahan. "Flow-based generative models for Markov chain Monte Carlo in lattice field theory". In: *Physical Review D* 100.3 (2019), p. 034515. DOI: `10.1103/PhysRevD.100.034515`.

[186] Denis Boyda et al. "Sampling using SU(N) gauge equivariant flows". In: *Physical Review D* 103 (7 Apr. 2021), p. 074504. DOI: `10.1103/PhysRevD.103.074504`.

[187] Danilo Jimenez Rezende et al. "Equivariant Hamiltonian Flows". 2019. arXiv: `1909.13739`.

[188] Linfeng Zhang, Weinan E, and Lei Wang. "Monge-Ampère Flow for Generative Modeling". 2018. arXiv: `1809.10188`.

[189] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. "Hamiltonian neural networks". In: *Advances in Neural Information Processing Systems.* 2019, pp. 15379–15389.

[190] Peter Toth et al. "Hamiltonian Generative Networks". 2020. arXiv: `1909.13789`.

[191] Chin-Wei Huang, Laurent Dinh, and Aaron Courville. "Augmented Normalizing Flows: Bridging the Gap Between Generative Flows and Latent Variable Models". 2020. arXiv: `2002.07101`.

[192] Peter Sorrenson, Carsten Rother, and Ullrich Köthe. "Disentanglement by Nonlinear ICA with General Incompressible-flow Networks (GIN)". 2020. arXiv: `2001.04872`.

[193] J. A. Gregory and R. Delbourgo. "Piecewise Rational Quadratic Interpolation to Monotonic Data". In: *IMA Journal of Numerical Analysis* 2.2 (Apr. 1982), pp. 123–130. DOI: `10.1093/imanum/2.2.123`.

[194] Danilo Jimenez Rezende et al. "Normalizing flows on tori and spheres". In: *International Conference on Machine Learning.* PMLR. 2020, pp. 8083–8092.

[195] Elliott Lieb, Theodore Schultz, and Daniel Mattis. "Two soluble models of an antiferromagnetic chain". In: *Annals of Physics* 16.3 (Dec. 1961), pp. 407–466. DOI: `10.1016/0003-4916(61)90115-4`.

[196] Roy J Glauber. "Time-dependent statistics of the Ising model". In: *Journal of Mathematical Physics* 4.2 (1963), pp. 294–307. DOI: `10.1063/1.1703954`.

[197] Eugenia Koblents and Joaquín Míguez. "A population Monte Carlo scheme with transformed weights and its application to stochastic kinetic models". In: *Statistics and Computing* 25.2 (2015), pp. 407–425. DOI: `10.1007/s11222-013-9440-2`.

[198] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". 2017. arXiv: `1412.6980`.

[199] Vinod Nair and Geoffrey E. Hinton. “Rectified Linear Units Improve Restricted Boltzmann Machines”. In: *ICML*. 2010, pp. 807–814. URL: `https://icml.cc/Conferences/2010/papers/432.pdf`.

[200] Prajit Ramachandran, Barret Zoph, and Quoc V. Le. “Searching for Activation Functions”. 2017. arXiv: `1710.05941`.