

Overcoming the timescale barrier in molecular dynamics: Transfer operators, variational principles and machine learning

Christof Schütte

*Zuse Institute Berlin and Freie Universität Berlin,
14195 Berlin, Germany
E-mail: schuette@zib.de*

Stefan Klus

*Heriot–Watt University, Edinburgh EH14 4AS, UK
E-mail: s.klus@hw.ac.uk*

Carsten Hartmann

*Brandenburgische Technische Universität Cottbus-Senftenberg,
03046 Cottbus, Germany
E-mail: carsten.hartmann@b-tu.de*

One of the main challenges in molecular dynamics is overcoming the ‘timescale barrier’: in many realistic molecular systems, biologically important rare transitions occur on timescales that are not accessible to direct numerical simulation, even on the largest or specifically dedicated supercomputers. This article discusses how to circumvent the timescale barrier by a collection of transfer operator-based techniques that have emerged from dynamical systems theory, numerical mathematics and machine learning over the last two decades. We will focus on how transfer operators can be used to approximate the dynamical behaviour on long timescales, review the introduction of this approach into molecular dynamics, and outline the respective theory, as well as the algorithmic development, from the early numerics-based methods, via variational reformulations, to modern data-based techniques utilizing and improving concepts from machine learning. Furthermore, its relation to rare event simulation techniques will be explained, revealing a broad equivalence of variational principles for long-time quantities in molecular dynamics. The article will mainly take a mathematical perspective and will leave the application to real-world molecular systems to the more than 1000 research articles already written on this subject.

2020 Mathematics Subject Classification: Primary 37M10, 37M25, 82C31, 82M37
Secondary 47D07, 60J35, 60J60

© The Author(s), 2023. Published by Cambridge University Press.

This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

CONTENTS

1	Introduction	518
2	Dynamical systems in molecular dynamics	522
3	Statistical mechanics of slow processes	534
4	Numerical analysis of transfer operators	573
5	Data-driven methods	589
6	Rare event simulation	617
7	Concluding remarks	658
	References	659

1. Introduction

Rare but important transition events between long-lived states are a key feature of many systems arising in physics, chemistry, biology and many other fields, and particularly in *molecular dynamics* (MD). MD simulations describe the dynamical behaviour of realistic molecular systems in atomistic resolution. However, in many realistic molecular systems, biologically important rare transitions occur on timescales that are not accessible to direct numerical simulation, even on dedicated supercomputers, and will still remain inaccessible on emerging exascale machines, with more than 18 orders of magnitude between a typical simulation time step ($\sim 1 \text{ fs} = 10^{-15} \text{ s}$) and slow biologically relevant processes such as protein–ligand or protein–protein association (10^1 – 10^3 s and beyond). This severely limits the MD-based analysis of many biological processes: the average waiting time between the rare transition events of interest is orders of magnitude longer than the timescale of the transition characterizing the event itself. Therefore, performing direct numerical simulation of the system until a reasonable number of events has been observed is impractically excessive for most interesting systems. As a consequence of this *timescale barrier*, rare event simulation and estimation are among the most challenging topics in molecular dynamics, despite the fact that MD simulation methods have seen significant improvement since their inception in the late 1950s. Constraints of simulation size and duration that once impeded the field have lessened with the advent of better theory and algorithms, faster processors and parallel computing. With newer computational techniques and hardware available, MD simulations of more biologically relevant timescales can now sample a broader range of conformational and dynamical changes in realistic molecular systems.

This article discusses how to overcome the timescale barrier in MD by a collection of techniques that have emerged from dynamical systems theory, numerical mathematics and machine learning over the last two decades. We will focus on transfer operator-based approaches to the approximation of dynamical behaviour on long timescales, and review its introduction into molecular dynamics and the respective theory, its algorithmic development from the early numerics-based

methods to modern data-based techniques utilizing and improving concepts from machine learning, and its complementation by rare event simulation techniques.

The transfer operator approach to molecular dynamics started with Deuffhard, Dellnitz, Junge and Schütte (1999), Schütte (1998) and Schütte, Fischer, Huisinga and Deuffhard (1999), arising from the two main insights that the rare transition events of importance in MD can be seen as transitions between metastable sets that can be understood as almost invariant subsets of state space in dynamical systems theory, and the determination of these metastable subsets and the transition between them from spatially discretized transfer operators. The first component – the idea of discretizing transfer operators using Ulam’s method – had been known for a fairly long time already, and its connection to almost invariant sets of dynamical systems stems from Dellnitz and Junge (1998, 1999), as far as we are aware. The other component – the relation between metastable subsets and dominant spectral elements of the transfer operator – was also known already, but in a completely different and fairly general setting, e.g. in Davies (1982a,b).

Putting these two components together started a growing theme in the molecular dynamics literature, which had its first peak around the year 2010 with the introduction of *Markov state models* (MSMs) (Bowman, Pande and Noé 2014) and its application to many realistic timescale barrier problems in MD, including protein folding (Chodera, Swope, Pitera and Dill 2006, Noé *et al.* 2009, Bowman, Volez and Pande 2011), kinetic fingerprinting and spectroscopic observables (Keller, Prinz and Noé 2011, Prinz, Keller and Noé 2011), RNA (Huang *et al.* 2010, Pina-monti *et al.* 2017), protein–peptide association (Paul *et al.* 2017), ligand binding, rebinding and multivalency (Weber and Fackeldey 2014, Ge and Voelz 2021), as well as numerical recipes (Pande, Beauchamp and Bowman 2010) and software implementations (Senne *et al.* 2012, Scherer *et al.* 2015, Beauchamp *et al.* 2011), to name just a few examples.

Based on the ideas and the theory behind MSMs, the years since 2010 have seen an enormous activity in extending the transfer operator approach by leaving the idea of a spatially discretized transfer operator behind in favour of more general low-dimensional representations via generalized ansatz spaces and variational formulations. The same period saw a recombination of transfer operator-based methods in molecular dynamics and advanced Koopman operator-based techniques in dynamical systems theory. This recombination process has led to considerable progress in both fields, not least because of the insight that many methods developed in the MD community, such as the *variational approach to conformation dynamics* (VAC) (Noé and Nüske 2013) or *time-lagged independent component analysis* (TICA) (Perez-Hernandez *et al.* 2013), are equivalent or very similar to methods such as *extended dynamic mode decomposition* (EDMD) (Williams, Kevrekidis and Rowley 2015a) originating in the dynamical systems community; see Klus *et al.* (2018b) and Wu *et al.* (2017) for details of the relations.

During the last ten years it has become ever clearer that there is a large family of dynamical quantities, such as expected hitting times, (dominant) relaxation

times, committor functions, moment generating functions, etc., which, although theoretically computable by solving very high-dimensional linear problems that involve transfer operators and their generators, admit variational formulations that can be used to design algorithms that are robust in high dimensions.

In recent years this process has been further enriched by means of advanced data-based approaches which utilize the immense progress in machine learning while being inspired by the availability of loss functionals originating from the variational formulations of many problems. Progress has been achieved, for example, by using kernel-based methods (Williams, Rowley and Kevrekidis 2015*b*, Schwantes and Pande 2015, Klus, Bittracher, Schuster and Schütte 2018*a*), utilizing reproducing kernel Hilbert spaces (RKHSs) (Klus, Schuster and Muandet 2019*b*) and deep learning approaches to molecular kinetics such as VAMPnet (Mardt, Pasquali, Wu and Noé 2018, Wu and Noé 2020), and their joint software implementation (Hoffmann *et al.* 2021). The enormous progress in data-based approaches has also inspired a variety of new theoretical and algorithmic approaches to the old problem of finding good reaction coordinates for molecular systems, ranging from its theoretical underpinning via transfer operators (Bittracher *et al.* 2018, Bittracher, Mollenhauer, Koltai and Schütte 2021) to the time-lagged autoencoder approach (Wehmeyer and Noé 2018) or the *learning the effective dynamics* (LED) approach (Vlachas, Zavadlav, Praprotnik and Koumoutsakos 2022) and many others; see Section 5.5 for more details.

The aim of this article is to review this development in Sections 2–5, mainly from a mathematical perspective. Section 2 sets the stage and notation and gathers the basics on Markov processes required subsequently. We also introduce the main kinds of (stochastic) dynamical systems typically considered in molecular dynamics. Section 3 outlines some of the main theoretical insights, such as the relation between the dominant spectrum of the transfer operator and long relaxation and autocorrelation timescales, expected passage times and transition rates from one metastable set to another, the relation between small exit rates and metastable sets, optimal metastable decompositions, committor functions, and the effective dynamics induced by optimal reaction coordinates. Section 4 takes the numerical perspective and reviews different techniques for finding optimal low-dimensional representations of the transfer operator by spatial discretization (theory of MSMs) or Galerkin projections to finite ansatz spaces (projected transfer operators). Finally, Section 5 adopts the dynamical systems perspective and considers data-based approaches to transfer/Koopman operators and generators such as EDMD and extensions, including data-based approaches to finding optimal reaction coordinates.

However, the large body of transfer operator-based theory and algorithms is just one side of the story regarding rare events in molecular systems. The molecular dynamics literature on rare event simulations is extremely rich. Since the 1930s, *transition state theory* (TST) (Eyring 1935, Wigner 1938) and its extensions based on reactive flux formalism have provided the main theoretical framework for the description of transition events. TST is based on partitioning the state space into

two sets with a dividing surface in between, leaving set A on one side and the target set B on the other, and the theory only tells us how this surface is crossed during the reaction. It is often difficult to choose a suitable dividing surface, and a bad choice will lead to a very poor estimate of the rate. The TST estimate is then extremely difficult to correct, especially if the rare event is of the diffusive type where many different reaction channels co-exist. Therefore many techniques have been proposed that try to go beyond TST.

In recent years, a wide variety of such computational techniques, jointly called *advanced sampling techniques*, have been developed to systematically capture and characterize rare transitions in molecular systems. Most notable among these techniques are the following.

- (1) The well-known family of techniques based on *transition path sampling* (TPS) (Bolhuis, Chandler, Dellago and Geissler 2002, Bolhuis and Swenson 2021).
- (2) The so-called string methods (Ren and Vanden-Eijnden 2002, Roux 2021) and optimal path approaches (Beccara, Skrbic, Covino and Faccioli 2012, Faccioli, Lonardi and Orland 2010, Pinski and Stuart 2010), and variants thereof.
- (3) Techniques that follow the progress of the transition through interfaces, such as *forward-flux simulation* (FFS) (Allen, Frenkel and Ten Wolde 2006, Hussain and Haji-Akbari 2020), *transition interface sampling* (TIS) (Moroni, van Erp and Bolhuis 2004, Swenson and Bolhuis 2014) and the milestoning techniques (Faradjian and Elber 2004, Bello-Rivas and Elber 2015, Berezhkovskii and Szabo 2019).
- (4) Methods that drive the molecular system by external forces, with the aim of making the required transition more frequent while still allowing computation of the exact rare event statistics for the unforced system, for example based on Jarzynski's and Crook's identity (Jarzynski 1997, Crooks 1999) such as (adaptive) steered MD (Zhuang, Bureau, Quirk and Hernandez 2021), or via accelerated molecular dynamics (Hamelberg, Mongan and McCammon 2004, Yang *et al.* 2022).

In this article we cannot even attempt to review these methods, all of which have their merits. Instead we will describe an alternative approach based on stochastic control techniques that reveals a (dual) variational formulation of the transfer operator approach to the timescale barrier problem. We explain why the direct numerical simulation of rare events in molecular simulation (such as protein folding) is often infeasible. Based on this, we will then explain in Section 6 how rare event statistics can be efficiently estimated using a combination of ideas from stochastic analysis, optimal control and statistical mechanics. The approach is based on the insight that direct rare event simulation mainly suffers from the fact that the sample variance of an estimator is generally much larger than the quantity of interest.

It turns out that a reformulation of the estimation problem by means of a Legendre-type duality argument offers some insight into how to reduce the sample variance in rare event simulation by a suitable change of the underlying probability measure. (Theoretically, it is even possible to obtain a zero variance estimator.) This *reweighting* can be seen as an importance sampling in path space, but also, equivalently, as an optimal control problem for MD. Then, utilizing the well-known Feynman–Kac theorem, it is shown that the Hamilton–Jacobi–Bellman (HJB) equation related to the optimal control problem is equivalent to a *linear equation* where the two equations are related by a logarithmic transformation. Our approach has some similarities to the adaptive importance sampling techniques developed by Dupuis, Wang and collaborators (e.g. Dupuis and Wang 2004, 2007), but it is specifically tailored to simulate rare events far from the asymptotic regime of large deviations (e.g. far from the limit of vanishing noise or temperature). Understanding these connections opens the door to a rather general insight into connections and equivalence of different perspectives (optimal reweighting, adaptive importance sampling in path space, optimal stochastic control, linear boundary value problems: see Theorem 6.22 for details) and their variational formulation which, in turn, opens the door to utilizing deep learning strategies in rare event simulation.

The theories and algorithms discussed in this article all aim to overcome or circumvent the timescale gap by a mixture of (i) reformulation of quantities characterizing long timescales in MD as linear problems, (ii) using equivalences of these linear problems to other, particularly variational formulations, and (iii) designing efficient algorithms by computing optimal dimension-reduced problems, which then are solved by means of the data-based approaches where MD simulations are used to produce the required data.

2. Dynamical systems in molecular dynamics

In this section we will introduce mathematical descriptions of molecular systems as well as transfer operators and their generators. The key notation is summarized in Table 2.1.

2.1. Markov processes

The literature describing the dynamical behaviour of molecular systems is extremely rich: it ranges from classical deterministic Hamiltonian models that try to cover the actual motion of each single molecule in the system to stochastic descriptions such as Langevin dynamics or iterative schemes such as most Markov chain Monte Carlo approaches. In this section we will introduce the mathematical framework that subsumes both stochastic and deterministic approaches.

Basic concept and notation. Consider the state space $\mathbb{X} \subset \mathbb{R}^{\mathcal{N}}$ for some $\mathcal{N} \in \mathbb{N}$ equipped with the Borel σ -algebra \mathcal{A} on \mathbb{X} . The evolution of a single microscopic system is assumed to be given by a homogeneous Markov process $X_t = \{X_t\}_{t \in \mathbb{T}}$ in

Table 2.1. Overview of the key notation.

X_t	stochastic process
$\mathbb{X} \subset \mathbb{R}^N$	state space
P^t	transfer operator or Perron–Frobenius operator with lag time t
P_*^t	adjoint transfer operator with lag time t
K^t	Koopman operator with lag time t
\mathcal{L} or L	generator of the transfer or Perron–Frobenius operator
\mathcal{L}^*	generator of the adjoint transfer operator or Koopman operator
$p(t, x, y)$ or $p^t(x, y)$	transition function of the process X_t
$\Lambda(t)$	eigenvalues of the transfer operator P^t
λ	eigenvalues of the generator under consideration
μ	invariant measure
$\langle \cdot, \cdot \rangle_\mu$	μ -weighted inner product
$\rho_x(A)$	hitting time of set A if starting at x
$\tau_x(A)$	exit time from set A if starting at x
τ	characteristic intermediate lag time
ξ	reaction coordinate
$\mathbb{Z} = \xi(\mathbb{X})$	reaction coordinate space
\mathcal{L}	likelihood ratio (Radon–Nikodým derivative)

continuous or discrete time $\mathbb{T} = \mathbb{R}$ or $\mathbb{T} = \mathbb{N}$, respectively. We write $X_0 \sim \mu$ if the Markov process X_t is initially distributed according to the probability measure μ , i.e. if $\mathbb{P}[X_0 \in A] = \mu(A)$ for every $A \subset \mathbb{X}$. We use $X_0 = x$ if $X_0 \sim \delta_x$, where δ_x denotes the Dirac measure centred at x . The motion of X_t is given in terms of the *stochastic transition function* $p: \mathbb{T} \times \mathbb{X} \times \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$ according to

$$p(t, x, A) = \mathbb{P}[X_{s+t} \in A \mid X_s = x], \tag{2.1}$$

for every $s, t \in \mathbb{T}$, $x \in \mathbb{X}$ and $A \subset \mathbb{X}$. Hence $p(t, x, A)$ describes the probability that the system moves from state x into the subset A within time t . The relation between a stochastic transition function and a homogeneous Markov process is one-to-one (Meyn and Tweedie 1993, Chapter 3). In the special case where $p(t, x, A) = \delta_{\Phi^t(x)}(A) = \mathbf{1}_A(\Phi^t(x))$ with the indicator function $\mathbf{1}_A$ of the set A , the Markov process is in fact a deterministic process, whose evolution is defined by the flow map $\Phi^t(x)$ in state space. As well as some more technical properties, the stochastic transition function fulfils the so-called Chapman–Kolmogorov equation

$$p(s + t, x, A) = \int_{\mathbb{X}} p(t, x, dz) p(s, z, A), \tag{2.2}$$

which holds for every $s, t \in \mathbb{T}$, $x \in \mathbb{X}$ and $A \subset \mathbb{X}$, and represents the semigroup property of the Markov process. As a consequence, in the discrete-time case $\mathbb{T} = \mathbb{N}$ it suffices to specify $p(x, dy) = p(1, x, dy)$.

We say that the Markov process X_t admits an *invariant probability measure* μ , or μ is invariant with respect to X_t if

$$\int_{\mathbb{X}} p(t, x, A) \mu(dx) = \mu(A) \quad (2.3)$$

for every $t \in \mathbb{T}$ and $A \subset \mathbb{X}$ (Meyn and Tweedie 1993, Chapter 10). Note that the invariant probability measure does not need to be unique.

For a measurable function $u: \mathbb{X} \rightarrow \mathbb{R}$, we denote the expectation value with respect to the measure μ by

$$\mathbb{E}_\mu(u) = \int_{\mathbb{X}} u(x) \mu(dx).$$

A Markov process is called *ergodic* with respect to μ if, for all functions $u: \mathbb{X} \rightarrow \mathbb{R}$ with $\mathbb{E}_\mu(|u|) < \infty$, we have

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T u(X_t) dt = \mathbb{E}_\mu(u), \quad (2.4)$$

for almost all initial values $X_0 = x$. Furthermore, the process is called (uniformly) *geometrically ergodic* if the convergence in (2.4) exhibits a (uniform) geometric rate. The integral on the left-hand side of the equation has to be replaced by a sum if $\mathbb{T} = \mathbb{N}$. There are many other definitions of ergodicity; we have chosen the one that will serve us most in what follows but is not the most general. For an ergodic process (X_t) we will also consider the temporal correlation between the functions $u, v: \mathbb{X} \rightarrow \mathbb{R}$ given by

$$C_{u,v}(t) = \mathbb{E}_\mu(v(X_t)u(X_0)) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T v(X_t)u(X_0) dt. \quad (2.5)$$

Other time-related quantities of central interest here are the first *hitting time*

$$\rho_x(A) = \inf\{t \geq 0: X_t \in A, X_0 = x\} \quad (2.6)$$

of a set A conditioned on the initial state x , and its expectation value

$$m_A(x) = \mathbb{E}(\rho_x(A)),$$

which is often called the *mean first passage time*, and the *exit time*

$$\tau_x(A) = \rho_x(A^c) \quad (2.7)$$

from A , where $A^c = \mathbb{X} \setminus A$ denotes the complement of A .

A Markov process is called *reversible* with respect to an invariant probability measure μ if

$$\int_A p(t, x, B) \mu(dx) = \int_B p(t, x, A) \mu(dx) \quad (2.8)$$

for every $t \in \mathbb{T}$ and $A, B \subset \mathbb{X}$. If μ is unique, X_t is simply called reversible. For

the special case of a stochastic transition function and the invariant measure being absolutely continuous with respect to the Lebesgue measure, that is, if

$$p(t, x, A) = \int_A p(t, x, y) \mu(y) dy,$$

then reversibility reads

$$\mu(x) p(t, x, y) = \mu(y) p(t, y, x) \quad (2.9)$$

for every $t \in \mathbb{T}$ and a.e. $x, y \in \mathbb{X}$.

2.2. Sampling and dynamics

In principle, molecular dynamics aims to define a mathematical model which, given an exact initial state of the molecular system, describes the true motion of the molecule under consideration in all necessary detail. However, this is only true in principle. In fact there are two different concepts of MD. The first concept introduces different Markov processes (X_t) in order to *sample* the associated invariant measure μ . In this *sampling* concept, we want ergodic Markov processes such that the long-term simulation of (X_t) allows us to approximate the expectation value $\mathbb{E}_\mu(u)$ of an observable u or similar statistical quantities by means of running averages

$$\mathbb{E}_\mu(u) \approx \frac{1}{T} \int_0^T u(X_t) dt. \quad (2.10)$$

In contrast, the *dynamics* concept uses specific Markov processes to compute dynamical information such as autocorrelation functions or exit times. In this concept, we build molecular dynamics models that try to incorporate all necessary internal and external interactions of the molecular system, derive associated forces and describe the molecular motion by equations of motion resulting from these forces. This approach normally utilizes classical mechanics and leads to deterministic Hamiltonian systems. The corresponding initial value problem is thought to model the evolution of the state of a *single, isolated molecular system*. If, however, we consider non-isolated molecular systems that interact with a macroscopic environment (heat bath), the Hamiltonian system model is no longer appropriate, and several extensions or stochastic models are considered instead. We will discuss some of these mathematical models, but start with an introduction to molecular interactions and forces underlying the Hamiltonian model.

Molecular forces, energy landscape, canonical ensemble. Let N denote the number of atoms of the molecular system and let $\Gamma = \mathbb{R}^{\mathcal{N}}$, $\mathcal{N} = 3N$, denote the position space, that is, $q = (q_1, \dots, q_{3N}) \in \Gamma$ represents the vector of atomic position coordinates, where $r_j = (q_{3(j-1)+1}, \dots, q_{3j}) \in \mathbb{R}^3$ is the vector of the Euclidean coordinates of atom j such that we can also write $q = (r_1, \dots, r_N)$. Here the atoms of the water molecules surrounding the molecule of interest have to be included. Let $\xi \in \mathbb{R}^{3N}$ denote the vector of all conjugated momenta. The differentiable

potential energy function $V: \mathbb{R}^{3N} \rightarrow \mathbb{R}$ that is intended to describe all interactions between the atoms in the molecular system is traditionally split into 1-atom, 2-atom, 3-atom, . . . forces, and often has the following form:

$$\begin{aligned} V(q) = & \sum_{(i,j) \in B} V_B(|r_i - r_j|) + \sum_{(i,j,k) \in A} V_A(\phi(r_i, r_j, r_k)) \\ & + \sum_{(i,j,k,l) \in D} V_D(\theta(r_i, r_j, r_k, r_l)) + \sum_{i,j=1}^N V_C(|r_i - r_j|) \\ & + \sum_{i,j=1}^N V_{LJ}(|r_i - r_j|). \end{aligned}$$

Here B is the set of all (i, j) such that atoms i and j are connected by a covalent bond, and V_B is the associated potential describing the forces related to the bond, $(i, j, k) \in A$ means that atom j has covalent bonds to i and k that form a bond angle $\phi(r_i, r_j, r_k)$ which enters into the bond angle potential V_A , V_D models the dihedral angle interaction between four neighbouring atoms connected by covalent bonds, V_C is the electrostatic interaction in the form of the Coulomb potential, and V_{LJ} is the Lennard-Jones interaction. Additional potentials such as the hydrogen bond interaction potentials may be added. There is an extensive literature on the way these potentials are determined experimentally or modelled theoretically.

For many molecular systems (in particular for biomolecular ones) the resulting energy landscape has a *rich spatial multiscale structure*. That is, there is a huge number of local minima and saddle points (the number is estimated to grow combinatorially with the number of atoms of the system), but the minima are grouped into a rather small number of deep wells in the energy landscape that are separated by high energy barriers or extended energy plateaus (Wales 2005).

Since we are working with Euclidean coordinates, the Hamiltonian of the system giving its total energy is simply given by

$$H(q, \xi) = \frac{1}{2} \xi^\top M^{-1} \xi + V(q), \quad (2.11)$$

where M denotes the diagonal mass matrix of the system.

The forces F acting on the atoms in the molecular system result from the gradient of the potential, i.e. $F(q) = -\nabla_q V(q)$. According to the form of the potential just described, these forces do *not* incorporate the thermodynamic interaction of the molecular system with some external heat bath that surrounds the system. They describe the *internal* interaction forces *within* the molecular system only, that is, the above potential energy function describes an *isolated* molecular system.

In many cases – e.g. in all biologically relevant cases – we are *not* interested in isolated systems but rather in systems in contact with a (macroscopic) heat bath that defines the temperature of the system. According to statistical mechanics, a molecular system in equilibrium contact with a heat bath of temperature \mathcal{T} (with

constant volume and without exchange of particles) generates a probability density function on the state space that is known as the canonical density or *canonical ensemble*

$$\mu(x) \propto \exp(-\beta H(x))$$

for some constant $\beta = 1/(k_B \mathcal{T}) > 0$ that can be interpreted as the inverse temperature, where k_B is Boltzmann's constant. The associated measure $\mu(dx) \propto \mu(x) dx$ is called the *canonical measure*.

Hamiltonian dynamics. The most fundamental model for the dynamical behaviour of molecular systems exploits classical Hamiltonian mechanics, that is, atoms are described as mass points subject to forces that are generated by the interaction potentials V outlined above. The dynamical behaviour is described by some deterministic Hamiltonian system of the form

$$\dot{q} = M^{-1}\xi, \quad \dot{\xi} = -\nabla_q V(q), \tag{2.12}$$

defined on the state space $\mathbb{X} = \mathbb{R}^{3N} \times \mathbb{R}^{3N}$, where M again denotes the diagonal mass matrix.

Equation (2.12) models an energetically closed, i.e. *isolated system*, whose total energy, given by the Hamiltonian H , is preserved under the dynamics. For the sake of simplicity, we assume in what follows that M is the identity matrix.

Let Φ^t denote the flow associated with the Hamiltonian system (2.12), that is, the solution x_t for the initial value x_0 is given by $x_t = \Phi^t x_0$, where we denote the states by $x = (q, \xi)$. Let $\mathbf{1}_C$ denote the characteristic function of the subset $C \subset \mathbb{X}$. Then the stochastic transition function corresponding to (2.12) is given by

$$p(t, x, C) = \mathbf{1}_C(\Phi^t x) = \delta_{\Phi^t x}(C) \tag{2.13}$$

for every $t \in \mathbb{R}_0^+$ and $C \subset \mathbb{X}$. The Markov process $X_t = \{X_t\}_{t \in \mathbb{R}_0^+}$ induced by the stochastic transition function p coincides with the flow Φ^t . Hence $X_t = \Phi^t x_0$ for the initial distribution $X_0 = x_0$.

One traditional aspect of molecular dynamics is the construction of extended Hamiltonian systems that allow for sampling the canonical ensemble by means of long-term simulation. Several concepts have been discussed that all boil down to the idea of constructing a Hamiltonian system in some extended state space $\hat{\mathbb{X}}$, whose projection onto the lower-dimensional state space \mathbb{X} of positions and momenta allows us to generate such a sampling. One of the most prominent examples is defined in terms of the Nosé Hamiltonian

$$H_{\text{Nosé}}(q, \xi, s, \pi) = \underbrace{\frac{1}{2s^2} \xi^\top \xi + V(q)}_{=H_s(q, \xi)} + \frac{1}{2Q} \pi^2 + \frac{1}{\beta} \log s,$$

where s is called the thermostat with conjugate momentum π and associated artificial mass Q . Let the flow of the associated Nosé Hamiltonian system be denoted

by Ψ^t and let Π denote the projection $(q, \xi, s, \pi) \mapsto (q, \xi)$. If Ψ^t is ergodic with respect to the microcanonical measure on the associated energy cell of $H_{\text{Nosé}}$, then $\Pi\Psi^t$ is ergodic with respect to the canonical measure $\mu(dx) \propto \exp(-\beta H_{s=1}(x)) dx$, where $x = (q, \xi)$ (Bond, Benedict and Leimkuhler 1999).

Langevin molecular dynamics. The most popular model for an *open system* having stochastic interaction with its environment is the so-called Langevin system (Risken 1996)

$$\dot{q} = \xi, \quad \dot{\xi} = -\nabla_q V(q) - \gamma \xi + \sigma \dot{W}_t, \quad (2.14)$$

defined on the state space $\mathbb{X} = \mathbb{R}^{6N}$. Here $\gamma > 0$ denotes some friction constant and $F_{\text{ext}} = \sigma \dot{W}_t$ the external forcing given by a $3N$ -dimensional Brownian motion W_t . The external stochastic force is assumed to model the influence of the heat bath surrounding the molecular system. In this case the internal energy given by the Hamiltonian H , as defined in (2.11), is not preserved, but the interplay between stochastic excitation and damping balances the internal energy. Under appropriate conditions on the potential V (see Mattingly, Stuart and Higham 2002 for details) that prevent the process from being explosive, the Langevin process is geometrically ergodic and the canonical measure $\mu(dx) \propto \exp(-\beta H(x)) dx$ with $x = (q, \xi)$ is its unique invariant measure if the noise and damping constants satisfy

$$\beta = \frac{2\gamma}{\sigma^2} \quad (2.15)$$

(see Risken 1996). The Langevin process does not satisfy the standard detailed balance condition but (under some growth conditions on V : see Hérau, Hitrik and Sjöstrand 2008) instead satisfies the following *extended detailed balance condition*:

$$\mu(x) p(t, x, y) = \mu(Ry) p(t, Ry, Rx), \quad (2.16)$$

with $R(q, \xi) = (q, -\xi)$.

Diffusive molecular dynamics. Diffusive molecular dynamics can be understood as an approximation to the Langevin model in the limit of high friction $\gamma \rightarrow \infty$; see Huisinga (2001) and Schütte and Huisinga (2000) for details. While the Langevin model gives a description of molecular motion in terms of positions and momenta of all atoms in the system, the diffusion model is stated in the position space only. Moreover, in contrast to the Langevin equation it defines a *reversible* Markov process that is given by the equation

$$\gamma \dot{q} = -\nabla_q V(q) + \sigma \dot{W}_t. \quad (2.17)$$

The stochastic differential equation (2.17) defines a continuous-time Markov process Q_t on the state space $\mathbb{X} = \Gamma \subseteq \mathbb{R}^{3N}$ with states $x = q$ with invariant probability measure $\mu(q) = Z^{-1} \exp(-\beta V(q))$ (Risken 1996), where $\beta = 2\gamma/\sigma^2$. It is known that under weak conditions on the potential function V the Markov process is ergodic and reversible (Huisinga 2001, Mattingly *et al.* 2002).

Obviously, the diffusion model cannot be expected to model all details of molecular motion, but it samples the required statistical distribution correctly. Nevertheless, it has a long history of use as a simple toolkit for investigating the dynamical behaviour in complicated energy landscapes (Chandler 1998).

In the literature, a more general form of diffusion molecular dynamics is sometimes considered:

$$\dot{q} = -\beta D(q) \nabla_q V(q) + \nabla \cdot D(q) + \sqrt{2} D^{1/2}(q) \dot{W}_t, \quad (2.18)$$

where D denotes the position-dependent diffusion tensor, and for which we again have $\mu(q) = Z^{-1} \exp(-\beta V(q))$.

Markov chain Monte Carlo. Markov chain Monte Carlo (MCMC) techniques are designed to sample a given probability density, particularly in high-dimensional state spaces. MCMC is an iterative realization of a specific Markov chain, whose stochastic transition function is given by

$$p(x, dy) = q(x, y) \mu(dy) + r(x) \delta_x(dy). \quad (2.19)$$

That is, the stochastic transition function is composed of a transition kernel $q(x, y)$, which is assumed to be μ -integrable, and a rejection probability

$$r(x) = 1 - \int_X q(x, y) \mu(dy) \geq 0.$$

In almost all situations, the transition kernel q is chosen in such a way that the stochastic transition function is reversible with respect to μ .

As an illustrative example, assume that we want to sample the canonical density $\mu(x) \propto \exp(-\beta V)$. Since we know diffusion molecular dynamics $\dot{x} = -\nabla_x V(x) + \sigma \dot{W}$ to be ergodic with respect to μ if $2/\sigma^2 = \beta$, we consider its time discretization using the Euler–Maruyama discretization with step size Δt :

$$x_{k+1} = x_k - \Delta t \nabla V(x_k) + \sigma \sqrt{\Delta t} \zeta_k, \quad (2.20)$$

where the ζ_k denote independent standard normally distributed random variables. Although resulting from a consistent first-order time discretization of diffusion molecular dynamics, the resulting Markov chain $(x_k)_{k=1,2,\dots}$ will generally *not* have μ as its invariant density. Instead we use the absolutely continuous transition function $\tilde{p}(x, y)$ associated with (2.20) as the so-called *proposal step* in the following sense. A Markov chain (X_k) is determined by the following transition rule: at X_k at time k we first determine

$$Z_{k+1} = X_k - \Delta t \nabla V(X_k) + \sigma \sqrt{\Delta t} \zeta_k.$$

However, we take this to be the next state, that is, we set $X_{k+1} = Z_{k+1}$, with so-called *acceptance probability*

$$r = a(X_k, Z_{k+1}) = \min(1, \exp(\beta(V(X_k) - V(Z_{k+1}))),$$

and with probability $1 - r$ we set $X_{k+1} = X_k$. The resulting transition function

$p(x, dy)$ is of the following form:

$$p(x, dy) = \tilde{p}(x, y) \cdot a(x, y) \cdot dy + r(x)\delta_x(dy),$$

where $r(x) = \int \tilde{p}(x, y) \cdot a(x, y) \cdot dy$ is the total rejection probability in x . Therefore it is of the form (2.19). The resulting Markov chain is ergodic with respect to its invariant density μ as long as the potential satisfies some essential growth conditions; see Bou-Rabee and Vanden-Eijnden (2010).

In general, MCMC is an *artificial* dynamical model that is not intended to model the dynamics of the molecular system but is just used to sample the canonical measure. However, there are MCMC methods, such as the popular *hybrid Monte Carlo* (HMC) method, which can be understood as a special realization of the Hamiltonian system with randomized momenta.

2.3. Fokker–Planck equations, transfer operators, and generators

First we consider the diffusive molecular dynamics Markov process (X_t) given by $\gamma \dot{X}_t = -\nabla V(X_t) + \sigma \dot{W}_t$. The transition function of this process is absolutely continuous with respect to the Lebesgue measure, i.e. $p(t, x, dy) = p(t, x, y) dy$. The evolution

$$v_t(x) = P^t v_0(x) = \int p(t, y, x) v_0(y) dy \quad (2.21)$$

of a function v_0 in time t under the process is given by the Kolmogorov forward equation (Risken 1996),

$$\partial_t v_t = \mathcal{L}_{\text{Diff}} v_t, \quad v_{t=0} = v_0, \quad (2.22)$$

with the differential operator

$$\mathcal{L}_{\text{Diff}} = \frac{\sigma^2}{2\gamma^2} \Delta + \frac{1}{\gamma} \nabla V(x) \cdot \nabla + \frac{1}{\gamma} \Delta V(x),$$

such that the semigroup defined by (2.21) can be formally written as

$$P^t v_0 = \exp(t \mathcal{L}_{\text{Diff}}) v_0,$$

defining a Markov operator $P^t : L^1 \rightarrow L^1$. It is often advantageous to consider the evolution of $v_t(x) = u_t(x)\mu(x)$ weighted with the invariant measure $\mu(x) \propto \exp(-\beta V(x))$, i.e. to look at the evolution of u_t instead of v_t . Then the evolution equation takes the form

$$u_t(x) = P^t u_0(x) = \frac{1}{\mu(x)} \int p(t, y, x) u_0(y) \mu(y) dy, \quad (2.23)$$

defining the Markov operator $P^t : L^1(\mu) \rightarrow L^1(\mu)$ or, using Hölder's inequality, in the Hilbert space $L^2(\mu) \subset L^1(\mu)$. Next we introduce the inner product

$$\langle u, v \rangle_\mu = \int_{\mathbb{X}} u(x)v(x)\mu(dx), \quad (2.24)$$

weighted by the invariant measure $\mu(x) \propto \exp(-\beta V(x))$. This bilinear form can be seen as the inner product of the Hilbert space $L^2(\mu)$ or as the duality bracket of $L^1(\mu)$ and $L^\infty(\mu)$, respectively. The adjoint operator P_*^t of the propagator P^t of the semigroup (2.23) with respect to $\langle \cdot, \cdot \rangle_\mu$ then has the form

$$P_*^t u(x) = \int p(t, x, y) u(y) dy = \mathbb{E}_x [u(X_t)] = \mathbb{E}(u(X_t) \mid X_0 = x). \tag{2.25}$$

The evolution equation associated with $u_t = P_*^t u_0$ reads

$$\partial_t u_t = L_{\text{Diff}} u_t, \tag{2.26}$$

where, for twice differentiable functions u_t , the generator has the form

$$L_{\text{Diff}} = \frac{\sigma^2}{2\gamma^2} \Delta - \frac{1}{\gamma} \nabla V(x) \cdot \nabla,$$

such that formally $P_*^t = \exp(tL_{\text{Diff}})$. Equations (2.22) and (2.26), respectively, are called the *Fokker–Planck equations* of diffusive molecular dynamics. For $\gamma = 1$ and $\beta = 2/\sigma^2$, L_{Diff} can also be expressed in the following compact form:

$$L_{\text{Diff}} u = \frac{1}{\beta} e^{\beta V} \nabla \cdot (e^{-\beta V} \nabla u). \tag{2.27}$$

For the case of Langevin molecular dynamics process $(X_t) = (q_t, \xi_t)$, the transition function is again absolutely continuous with respect to the Lebesgue measure, and the literature gives the Kolmogorov forward equation

$$\partial_t v_t = \mathcal{L}_{\text{Lan}} v_t, \quad v_{t=0} = v_0,$$

with the differential operator

$$\mathcal{L}_{\text{Lan}} = \frac{\sigma^2}{2} \Delta_\xi - \xi \cdot \nabla_q + \nabla_q V \cdot \nabla_\xi + \gamma \xi \cdot \nabla_\xi + \gamma,$$

and the Fokker–Planck equation can again be considered weighted according to the invariant measure $\mu(x) = \mu(q, \xi) \propto \exp(-\beta(\xi^\top \xi/2 + V(q)))$ in the same form as above. Then the weighted evolution equation for (2.23) again reads $\partial_t u_t = L_{\text{Lan}} u_t$, this time with

$$L_{\text{Lan}} u = \left(\frac{\sigma^2}{2} \Delta_p - p \cdot \nabla_q + \nabla_q V \cdot \nabla_p - \gamma p \cdot \nabla_p \right) u.$$

We consider more general stochastic differential equations of the form

$$dX_t = b(X_t) dt + \sigma(X_t) dW_t, \tag{2.28}$$

using the ‘infinitesimal’ notation in parallel to the more traditional ODE-inspired notation $\dot{x}_t = b(x_t) + \sigma(x_t) \dot{W}_t$. Here $b: \mathbb{X} \rightarrow \mathbb{X}$ is the drift term, $\sigma: \mathbb{X} \rightarrow \mathbb{R}^{N \times s}$ the state-dependent diffusion term and W_t an s -dimensional Wiener process. The

generator is given by

$$\mathcal{L}f = - \sum_{i=1}^{\mathcal{N}} \frac{\partial(b_i f)}{\partial x_i} + \frac{1}{2} \sum_{i=1}^{\mathcal{N}} \sum_{j=1}^{\mathcal{N}} \frac{\partial^2(a_{ij} f)}{\partial x_i \partial x_j} \quad (2.29)$$

and its adjoint generator (with respect to the standard L^2 inner product) by

$$L f = b \cdot \nabla_x f + \frac{1}{2} a : \nabla_x^2 f = \sum_{i=1}^{\mathcal{N}} b_i \frac{\partial f}{\partial x_i} + \frac{1}{2} \sum_{i=1}^{\mathcal{N}} \sum_{j=1}^{\mathcal{N}} a_{ij} \frac{\partial^2 f}{\partial x_i \partial x_j}. \quad (2.30)$$

Here $a = \sigma \sigma^\top$ and ∇_x^2 denotes the Hessian.

2.3.1. Transfer operators

For Langevin and diffusion molecular dynamics, the μ -weighted evolution semi-groups (2.23) are governed by the respective Fokker–Planck equations and define *transfer operators* P^t in $L^r(\mu)$ for $1 \leq r < \infty$ which admit strong generators L_{Diff} and L_{Lan} such that the transfer operators can be written as

$$P^t = \exp(tL_{\text{Diff}}) \quad \text{and} \quad P^t = \exp(tL_{\text{Lan}}),$$

respectively. The definition (2.23) of the transfer operator can be repeated for all Markov processes in the following slightly more general form.

Definition 2.1 (transfer operators and adjoints). For the Markov process (X_t) with transition function $p(t, \cdot, \cdot)$ and unique invariant measure μ , the transfer operator $P^t : L^r(\mu) \rightarrow L^r(\mu)$, with $r = 1, 2$, is given by its action on μ -measurable subsets $C \subset \mathbb{X}$:

$$\int_C (P^t u)(x) \mu(dx) = \int_{\mathbb{X}} p(t, y, C) u(y) \mu(dy). \quad (2.31)$$

We denote the adjoint operator by $P_*^t u(x) = \mathbb{E}_x[X_t(x)]$. For deterministic processes with flow map Φ^t , the transfer operator is given by the *Perron–Frobenius operator*

$$\int_C P^t u(x) \mu(dx) = \int_{\Phi^{-t}(C)} u(x) \mu(dx)$$

and its adjoint by the *Koopman operator* $K^t u(x) = P_*^t u(x) = u(\Phi^t x)$.

Remark 2.2 (Koopman and Perron–Frobenius operators). Different communities use different naming conventions. For example, P^t is called the ‘transfer operator’ in parts of the stochastic community in order to contrast it to the ‘Perron–Frobenius operator’ for deterministic systems. In turn, the dynamical systems community tends to call both $K^t u(x) = u(\Phi^t x)$ and $P_*^t u(x) = \mathbb{E}_x[X_t(x)]$ the ‘(stochastic) Koopman operator’, although the name was originally used only for

the adjoint transfer operator for deterministic systems. We will use the name that is mostly used in the respective context, that is, we will mainly use the phrase ‘transfer operator’ throughout Sections 2, 3 and 4 and switch to ‘Koopman operator’ and ‘Perron–Frobenius operator’ in Section 5.

Remark 2.3 (Hamiltonian dynamics). For Hamiltonian dynamics, the Perron–Frobenius operator has the simple form $P^t u(x) = u(\Phi^{-t}x)$.

Depending on the context, the different transfer operators are considered on the Hilbert spaces $L^2_{1/\mu}(\mathbb{X})$ and $L^2_\mu(\mathbb{X})$, or on the Banach spaces L^1 or $L^1(\mu)$. They are still well-defined non-expansive operators on these spaces.

2.3.2. Pathwise quantities

We can also consider pathwise quantities such as the non-probabilistic (i.e. not probability-conserving) semigroup

$$P^t_f g(x) = \mathbb{E} \left[g(X_t) \exp \left(- \int_0^t f(X_s, s) ds \right) \middle| X_0 = x \right], \tag{2.32}$$

for a lower-bounded function f where the expectation is taken over the realizations of the process $(X_s)_{0 \leq s \leq t}$ for fixed initial condition $X_0 = x$.

Now, let (X_t) denote the diffusive or Langevin molecular dynamics process with generator $L = L_{\text{Diff}}$ or $L = L_{\text{Lan}}$, respectively. Then the *Feynman–Kac theorem* tells us that $u(x, t) = P^t_f g(x)$ is the solution of

$$\partial_t u(x, t) = Lu(x, t) - f(x, t)u(x, t), \quad u(x, 0) = g(x). \tag{2.33}$$

In a similar fashion, in potential theory we consider the so-called *potential* ϕ associated with two real-valued functions c and g on state space, a set D and its complement $D^c = \mathbb{X} \setminus D$,

$$\phi(x) = \mathbb{E} \left[\int_0^\rho c(X_t) dt + g(X_\rho) 1_{\rho < \infty} \middle| X_0 = x \right], \tag{2.34}$$

where $\rho = \rho_x(D^c)$ is the hitting time of the complement of D . There are some restrictions to the choice of D . For example, for processes with $\mathbb{T} = \mathbb{R}$ and continuous paths, the process is stopped on the boundary ∂D of D , i.e. $X_\rho \in \partial D$, such that D should be an open set with ‘nice’ boundary. The functions c and g are regarded as *cost functions*, so that $\int_0^\rho c(X_t) dt$ is the cost for ‘wandering around’ in D , while $g(X_\rho)$ is the final cost when the process hits D^c at the boundary ∂D , and the potential ϕ can be interpreted as an *expected total cost*.

We will come back to the cost interpretation of the potential in Section 6. For now we just emphasize that many interesting quantities can be formulated in terms of potentials (Doob 1984), and we will restrict our attention to the mean first passage time. To this end, first consider a set A , choose $D = A^c$ so that $D^c = A$, and set $c = \mathbf{1}$, the constant function, and $g = 0$. Then the potential ϕ takes the following

form (Doob 1984):

$$\phi(x) = \mathbb{E} \left[\int_0^{\rho_x(A)} c(X_t) dt \right] = \mathbb{E}[\rho_x(A)] = m_A(x).$$

One of the main insights of potential theory is that the potential can be determined as the solution of a linear boundary value problem (Doob 1984),

$$\begin{aligned} L\phi &= -c && \text{in } D, \\ \phi &= g && \text{in } D^c, \end{aligned} \tag{2.35}$$

if a solution exists. In the continuous-time case, that is, for $\mathbb{T} = \mathbb{R}$ (diffusive or Langevin molecular dynamics), L denotes the generator of the underlying Markov process. If $\mathbb{T} = \mathbb{N}$ we set $L = P - \text{Id}$, where P is the one-step transfer operator associated with the process. When considering the mean first hitting time m_A of a set A with the above choices for D , g and c , we get

$$Lm_A = -\mathbf{1} \quad \text{in } A^c, \tag{2.36}$$

$$m_A = 0 \quad \text{in } A. \tag{2.37}$$

3. Statistical mechanics of slow processes

3.1. Dominant timescales, rare events and metastability

In this section we will outline the theory behind the transfer operator approach: how the dominant spectral elements of transfer operators may be used to characterize the longest relaxation timescales of the Markov process under consideration, how these spectral elements are related to metastable sets as well as large exit and hitting times, how the rare transition processes between the main metastable sets can be described, and how ‘good’ reaction coordinates can be characterized and related to the effective dynamics of the respective molecular system.

3.1.1. Properties of transfer operators

The transfer operator P^t as defined in (2.31) is a Markov operator in $L^1(\mu)$, that is, P^t conserves norm, $\|P^t v\|_{r,\mu} = \|v\|_{r,\mu}$ for $r = 1, 2$, and positivity, $P^t v \geq 0$ if $v \geq 0$. That is, $P^t v_0$ describes the transport of the function v_0 in time by the underlying dynamics given by X_t and weighted relative to μ . The spectrum of P^t is contained in the unit circle of the complex plane, i.e. $\text{spec}(P^t) \subset B_1(0) \subset \mathbb{C}$. As a consequence of the invariance of μ , the characteristic function $\mathbf{1}_{\mathbb{X}}$ of the entire state space is invariant under the action of P^t , that is,

$$P^t \mathbf{1}_{\mathbb{X}} = \mathbf{1}_{\mathbb{X}},$$

that is, $\Lambda_0 = 1$ always is an eigenvalue. Whenever $\mathbb{T} = \mathbb{R}$ the Chapman-Kolmogorov property of the transition functions makes the family $\{P^t\}_{t \in \mathbb{R}}$ a continuous semigroup. For special processes, the associated (infinitesimal) generators have already been discussed above; the general case is given next.

Definition 3.1. For the semigroup of propagators $P^t : L^r(\mu) \rightarrow L^r(\mu)$ with $1 \leq r \leq \infty$, define $\mathcal{D}(L)$ as the set of all $v \in L^r(\mu)$ such that the strong limit

$$Lv = \lim_{t \rightarrow \infty} \frac{P^t v - v}{t}$$

exists. Then the operator $L : \mathcal{D}(L) \rightarrow L^r(\mu)$ is called the infinitesimal generator corresponding to the semigroup P^t (Karatzas and Shreve 1991, Lasota and Mackey 1994).

This definition has consequences in general that we have already used above for special cases.

- (1) The extension of L to $L^2(\mu)$ satisfies $P^t = \exp(tL)$ in $L^2(\mu)$; its largest eigenvalue is $\lambda = 0$.
- (2) The function $\mathbf{1} \in L^2(\mu)$ represents the invariant measure and satisfies $L\mathbf{1} = 0$.
- (3) An eigenvalue $\Lambda(t)$ of P^t is related to an eigenvalue λ of L via

$$\Lambda(t) = \exp(t\lambda), \tag{3.1}$$

and the associated eigenfunctions are identical.

Reversibility and self-adjointness. Transfer operators P_t associated with reversible Markov processes are of particular interest since they possess additional structure on the Hilbert space $L^2(\mu)$. Such propagators will be called reversible too. A reversible P^t is μ -symmetric and, under some technical conditions, self-adjoint with respect to the inner product $\langle \cdot, \cdot \rangle_\mu$ in $L^2(\mu)$. Thus the eigenvalues of P^t are real-valued such that $\text{spec}(P^t) \subset [-1, 1]$ and the generator L is self-adjoint too, and its spectrum is contained in $(-\infty, 0]$. This is true for (the reversible Markov process related to) diffusive molecular dynamics, but not for Langevin molecular dynamics!

The Langevin molecular dynamics process does not satisfy the standard detailed balance condition but instead the extended detailed balance condition (2.16). Therefore the operator $(\mathcal{R}u)(x) = u(Rx)$ with $R(q, \xi) = (q, -\xi)$ satisfies $\mathcal{R}P^t_* = P^t \mathcal{R}$. That is, with the bilinear form on $L^2(\mu)$ defined by $\langle u, v \rangle_R = \langle \mathcal{R}u, v \rangle_\mu$, we get

$$\langle u, P^t v \rangle_R = \langle P^t u, v \rangle_R,$$

that is, an extended reversibility of the Langevin transfer operator P^t in the associated Hilbert space $L^2(\mathcal{R})$. As we will see later, this property implies that the leading eigenvalues of P^t in $L^2(\mu)$ are also real-valued in most cases of interest in molecular dynamics.

3.1.2. Dominant eigenmodes and implied timescales

Self-adjoint transfer operators. Let us first consider the simple case of a reversible process with a compact transfer operator P^t . Then classical results tell us the following about the spectral decomposition of P^t .

Theorem 3.2 (spectral decomposition of self-adjoint transfer operators). Let the transfer operator P^t be self-adjoint and compact in $L^2(\mu)$. Then there is a countable, monotonically decreasing sequence of real-valued eigenvalues $1 = \Lambda_0 \geq \Lambda_1(t) \geq \Lambda_2(t) \geq \dots$ with repetitions due to multiplicity, and associated eigenfunctions $u_i \in L^2(\mu)$ such that $P^t u_i = \Lambda_i(t) u_i$ and $\langle u_i, u_j \rangle_\mu = \delta_{ij}$, such that P^t exhibits the spectral decomposition

$$P^t = \sum_{i=0}^{\infty} \Lambda_i(t) \langle u_i, \cdot \rangle_\mu u_i.$$

Whenever P^t possesses an associated generator L such that $P^t = \exp(tL)$, then its eigenvalues λ_i are real-valued and non-positive starting with $\lambda_0 = 0$, and

$$P^t = \sum_{i=0}^{\infty} \exp(t\lambda_i) \langle u_i, \cdot \rangle_\mu u_i. \quad (3.2)$$

That is, the longest timescales exhibited by the evolution governed by P^t are induced by the dominant eigenvalues and satisfy

$$T_i = \frac{1}{|\lambda_i|} = -\frac{t}{\log(\Lambda_i(t))}. \quad (3.3)$$

This implies important relations for some dynamical quantities. For example, the time correlation $C_{u,v}(t)$ of functions u and v , as defined in (2.5), is of the form

$$C_{f,g}(t) = \langle P_*^t g, f \rangle_\mu = \sum_{i=0}^{\infty} \exp(t\lambda_i) \langle u_i, g \rangle_\mu \langle u_i, f \rangle_\mu,$$

showing that the dominant timescales T_i belong to the slowest temporal auto-correlations or relaxations in the system.

If P^t is not compact, the spectrum may be composed of discrete (isolated eigenvalues) and continuous parts. If the $m+1$ leading spectral elements are isolated eigenvalues $1 = \Lambda_0 \geq \Lambda_1(t) \geq \dots \geq \Lambda_m(t)$, $\Lambda_i(t) = \exp(t\lambda_i)$, then the rest of the spectrum is contained in $[0, e^{-tr}]$ with $|\lambda_m| < r \leq |\lambda_{m+1}|$, and the spectral decomposition reads

$$P^t = \sum_{i=0}^m \exp(t\lambda_i) \langle u_i, \cdot \rangle_\mu u_i + E^t, \quad \|E^t\|_2 \leq e^{-tr}. \quad (3.4)$$

Thus the leading implied timescales from (3.3) would still be the dominant timescales of P^t .

Non-self-adjoint transfer operators. For non-reversible processes we have to switch to singular values instead of eigenvalues. For a compact transfer operator $P = P^t$ for some time t we have the following. Consider $P: H \rightarrow F$, in general for two

separable Hilbert spaces H and F . There exists a singular value decomposition (SVD) given by

$$P = \sum_{i \in I} s_i \langle u_i, \cdot \rangle_H v_i, \tag{3.5}$$

where I is either a finite or a countably infinite ordered index set, $\{u_i, i \in I\} \subset H$ and $\{v_i, i \in I\} \subset F$ are two orthonormal systems, and $\{s_i, i \in I\} \subset \mathbb{R}^+$ is the set of singular values. As for the eigendecomposition, the sequence $\{s_i\}$ is a null sequence if I is not finite. Without loss of generality, we assume the singular values of compact operators are sorted in decreasing order, i.e. $s_i \geq s_{i+1}$. The following theorem from [Mollenhauer and Koltai \(2020\)](#) gives the relation of the SVD to eigenvalues.

Theorem 3.3 (singular value decomposition of compact transfer operators).

Let H and F be two separable Hilbert spaces, let $P: H \rightarrow F$ be compact and P_* its adjoint. Let $\{\lambda_i, i \in I\}$ denote the set of non-zero eigenvalues of P_*P counted with their multiplicities and $\{u_i\}$ the corresponding normalized eigenfunctions of P_*P . Then, with $v_i = \Lambda^{-1/2} P u_i$, the singular value decomposition of P is given by

$$P = \sum_{i \in I} \Lambda_i^{1/2} \langle u_i, \cdot \rangle_H v_i, \tag{3.6}$$

where $\langle \cdot, \cdot \rangle_H$ denotes the inner product of H .

For the situation of interest here, we have $H = F = L^2(\mu)$, and the implied dominant timescales of the process associated with P^t have to be defined by

$$T_i = -\frac{2t}{\log(\Lambda_i(t))},$$

where $\Lambda_i(t)$ denotes the i th eigenvalue of $P_*^t P^t$.

Existence of dominant eigenvalues. The existence of dominant eigenvalues requires that the essential/continuous part of the spectrum is bounded away from the dominant elements of the discrete spectrum. For the sake of simplicity, let us fix a time t and consider the transfer operator $P = P^t$ in the Hilbert space $L^2(\mu)$; for the case of L^1_μ , see [Huisinga \(2001\)](#). In order to allow for at least one discrete, isolated eigenvalue we need the essential spectral radius r_{ess} of P , i.e. the minimal radius of a circle in the complex plane around $\Lambda = 0$ that contains the essential spectrum, to be strictly bounded away from the largest eigenvalues. Moreover, we need $\Lambda = 1$ to be the only eigenvalue on the unit circle since we want uniqueness of the invariant measure. This leads to the following two *conditions on the transfer operator P*.

- (C1) The essential spectral radius of P is less than one, i.e. $r_{\text{ess}}(P) < 1$.
- (C2) The eigenvalue $\Lambda = 1$ of P is simple and dominant, i.e. $\eta \in \text{spec}(P)$ with $|\eta| = 1$ implies $\eta = 1$.

For diffusive molecular dynamics, for example, (C1) and (C2) are satisfied whenever the potential energy V grows to infinity for $x \rightarrow \infty$ and we can have an even stronger statement (Zhang, Li and Schütte 2022).

Theorem 3.4 (discrete spectrum for diffusive molecular dynamics). Let the energy function $V \in C^2(\mathbb{X})$ satisfy the growth conditions $\liminf_{x \rightarrow \infty} V(x) = \infty$ and

$$\lim_{x \rightarrow \infty} |\nabla V|(x) = \infty, \quad \liminf_{x \rightarrow \infty} ((1 - \delta)\beta |\nabla V(x)|^2 - \Delta V(x)) > 0 \quad (3.7)$$

for some $\delta \in (0, 1)$. Then the spectrum of L_{Diff} in $L^2(\mu)$ is discrete with isolated eigenvalues $\lambda \in (-\infty, 0]$.

In general, if P is self-adjoint, it satisfies conditions (C1) and (C2) in $L^2(\mu)$ if its stochastic transition function is geometrically or \mathcal{V} -uniformly ergodic; more precisely, there is an $\epsilon > 0$ such that the $\text{spec}(P) \cap (1 - \epsilon, 1]$ is a finite set and $\Lambda = 1$ is a simple eigenvalue. For non-reversible transfer operators, the statement is not true in general; for sufficient conditions see Huisinga (2001, Section 6) or Mattingly *et al.* (2002), Kontoyiannis and Meyn (2012, 2003) and Down, Meyn and Tweedie (1995). However, ergodicity may be difficult or even infeasible to check in practice. Therefore we consider the *Lebesgue decomposition* of the stochastic transition function

$$p(x, dy) = p_a(x, y)\mu(dy) + p_s(x, dy),$$

where p_a and p_s , respectively, represent the absolutely continuous and singular parts with respect to μ . Assume that

(A1) the inequality

$$\int_{\mathbb{X}} \int_{\mathbb{X}} p_a(x, y)^2 \mu(dx)\mu(dy) < \infty$$

holds, and

(A2) there exists some $\eta < 1$ such that

$$\eta = \sup p_s(x, \mathbb{X}) = 1 - \inf \int_{\mathbb{X}} p_a(x, y)\mu(dy) \quad \text{for } \mu\text{-a.e. } x \in \mathbb{X}.$$

Then the essential spectrum is uniformly bounded away from 1 and condition (C1) is fulfilled. That is, if the invariant measure μ is absolutely continuous with respect to the Lebesgue measure as for the canonical ensemble, for example, then the essential spectral radius $r_{\text{ess}}(P)$ is strictly bounded away from one if the singular part p_s is bounded away from one *and* the absolutely continuous part p_a does not grow too fast at infinity. This implies the following conclusions. For the special case of a deterministic Hamiltonian system, the absolutely continuous part p_a vanishes such that $r_{\text{ess}}(P) = 1$. Then, for the case of Langevin or diffusion molecular dynamics with smooth potentials, the singular part vanishes and the validity of condition (C1) only depends on the growth of p_a at infinity, which can be controlled by appropriate growth conditions on the potential energy V .

So, as the essential result of this discussion, let us summarize that under appropriate growth conditions on V , the transfer operators for Langevin and diffusion molecular dynamics will satisfy conditions (C1) and (C2) but the conditions are not satisfied for deterministic Hamiltonian systems, for which the essential spectral radius is one. More detailed analysis of the latter case reveals that the transfer operator is even unitary in $L^2(\mu)$; see Schütte (1998) for details.

3.1.3. Dominant eigenvalues, metastable sets and large exit times

As we have seen, the dominant eigenvalues of the transfer operator are directly related to the dominant relaxation timescales of the dynamics. In molecular dynamics, these timescales are often caused by the existence of *metastable sets*. The intuitive concept of metastable sets is based on two properties: (i) if the process starts in the set, it is likely to take a long time to get out (large exit time), and (ii) once the process has escaped, it is likely to return only after a long time.

Theoretical approaches to metastability. Metastability has been studied from different perspectives and with fairly diverse theoretical tools: quite early on, E. B. Davies provided the following insight into the relation between metastable sets and low-lying eigenvalues of the generator of a reversible Markov process (Davies 1982b, Theorem 6).

Theorem 3.5 (two dominant eigenvalues induce two metastable sets). Let L be the generator of a reversible Markov process with leading eigenvalues $0 = \lambda_0 > \lambda_1$ and $\text{spec}(L) \subset \{\lambda_0, \lambda_1\} \cup [-1, -\infty)$ in $L^2(\mu)$, and let u_1 denote the normalized eigenfunction associated with λ_1 . Then there is a Borel set $M \subset \mathbb{X}$ with complement $M^c = \mathbb{X} \setminus M$ such that, possibly by replacing u_1 with $-u_1$,

$$\left\| u_1 - \left(\sqrt{\frac{\mu(M^c)}{\mu(M)}} \mathbf{1}_M - \sqrt{\frac{\mu(M)}{\mu(M^c)}} \mathbf{1}_{M^c} \right) \right\|_{2,\mu} \leq 4|\lambda_1|^{1/2}.$$

This theorem (along with its generalizations to more than two dominant eigenvalues, also discussed in Davies 1982b) seems to indicate that the following two ‘rules’ hold.

- (R1) Dominant eigenvalues and timescales are related to the most important metastable sets.
- (R2) Dominant eigenfunctions are almost constant on these most important metastable sets.

However, Davies himself states that ‘one of the difficulties . . . is that the partition into M and M^c is not unique but only more and more nearly so as λ_1 tends to zero’. There is a large variety of articles that aim to describe the metastable sets in a more precise way and also to ‘justify’ rules (R1) and (R2). In this respect a crucial role was played by large-deviation tools inherited from Wentzell and Freidlin in their reduction procedure from continuous stochastic processes

to finite configuration-space Markov chains with exponentially small transition rates (Freidlin and Wentzell 1998), in particular for small noise; see Olivieri and Vares (2005). The relation between the spectrum of the transfer operator or generator of the stochastic process and metastability has been studied from different perspectives and linked to molecular dynamics; see Schütte (1998) and Huisinga, Meyn and Schütte (2004). Then, using potential-theoretic rather than large-deviation tools, Bovier, Eckhoff, Gayraud and Klein developed a set of general techniques to compute sharp asymptotics of the expected value of asymptotic exponential laws associated with the metastability phenomenon and revisited the relation between the spectrum of the generator of the stochastic dynamics and metastability (Bovier, Eckhoff, Gayraud and Klein 2002a,b, Bovier, Gayraud and Klein 2002c, Bovier and Den Hollander 2016). The potential-theoretic approach also led to the so-called martingale approach, which replaces the characterization of metastability with one expressed only in terms of the capacities that can be estimated using the Dirichlet and Thomson variational principles; see Beltran and Landim (2010, 2013). Many other theoretical concepts have been developed, for instance the notion of quasi-stationary distributions within a metastable set for the continuous state space Markov process to parametrize the exit event from the sets (Di Gesu, Lelièvre, Peutrec and Nectoux 2016). A measure ν_C is called the *quasi-stationary distribution* (QSD) induced by the process (X_t) in the set C if, for all measurable $A \subset C$,

$$\nu_C(A) = \frac{\int_C \mathbb{P}_x(X_t \in A, t < \tau_x(C)) \nu_C(dx)}{\int_C \mathbb{P}_x(t < \tau_x(C)) \nu_C(dx)}. \quad (3.8)$$

Thus, if the initial state $X_0 \in C$ is distributed according to ν_C , then, under the condition that the process does not exit from C in $[0, t]$, (X_t) still has law ν_C , which explains the name ‘quasi-equilibrium distribution’. For diffusive molecular dynamics the QSD is absolutely continuous, $\nu_C(dx) = \nu_C(x)\mu(dx)$, for sets C with smooth boundary ∂C , and its density ν_C satisfies

$$\begin{aligned} L_{\text{Diff}} \nu_C &= \tilde{\lambda}_0 \nu_C & \text{on } C, \\ \nu_C &= 0 & \text{on } \partial C, \end{aligned} \quad (3.9)$$

where $\tilde{\lambda}_0$ is the largest (negative) eigenvalue of L_{Diff} with Dirichlet boundary conditions on ∂C . The expected exit time $\mathbb{E}[\tau(C)]$ from C scales like $1/|\tilde{\lambda}_0|$. Moreover, the equilibration time within C towards the QSD scales like $1/(\tilde{\lambda}_0 - \tilde{\lambda}_1)$, where $\tilde{\lambda}_1 < 0$ is the second largest eigenvalue of L_{Diff} with Dirichlet boundary conditions on ∂C . Thus, in the QSD approach, the set C is called metastable if $1/|\tilde{\lambda}_0| \gg 1/(\tilde{\lambda}_0 - \tilde{\lambda}_1)$.

Metastability in molecular dynamics. In MD, especially in biomolecular systems, transitions between metastable sets seem to be the main cause of the longest timescales in the system, and these timescales are often approximately given by the

implied timescales T_i induced by the dominant eigenvalues. In fact, there are at least two different concepts for relating dominant eigenvalues and metastable sets.

- Dominant eigenvalues are related to large exit times from the dominant metastable sets (or the associated very small exit rates) or, analogously, large hitting times of a certain metastable set if starting in another.
- Dominant eigenvalues (and the associated eigenfunctions) are related to the optimal decomposition of the process's state space into metastable sets.

The underlying idea is most easily explained for diffusive molecular dynamics via Kramer's rule (Berglund 2013). The energy landscape V contains some deep wells that are separated by other wells by high energy barriers; for an overview regarding energy landscapes and their multiscale structure, see Wales (2003) and Röder and Wales (2022). When starting deep within one of these deep wells, the expected hitting time ρ of another of the deep wells scales with the energy barrier ΔV like

$$\mathbb{E}[\rho] \asymp \exp\left(\frac{2\Delta V}{\sigma^2}\right) \quad \text{as } \sigma \rightarrow 0, \quad (3.10)$$

where \asymp means asymptotic equivalence, $\lim_{\sigma \rightarrow 0} \sigma^2 \mathbb{E}[\rho] = 2\Delta V$. Moreover, $\mathbb{E}[\rho]$ is related to the corresponding dominant eigenvalue $\lambda < 0$ of the generator L of the process by

$$|\lambda| \asymp \frac{1}{\mathbb{E}[\rho]} \quad \text{as } \sigma \rightarrow 0.$$

That is, the longest timescales result from the metastability of the deepest wells in the energy landscape. In what follows, we will also see that the eigenfunctions related to the dominant eigenvalues are 'almost constant' within the main wells.

Example 3.6 (illustration for diffusive molecular dynamics). Figure 3.1 gives an illustration of the diffusive molecular dynamics process (2.17) with parameters $\gamma = 1$, $\sigma = \sqrt{2\epsilon}$ and $\epsilon = 0.25$ for the three-well potential

$$V(x) = \frac{1}{400}(x^6 - 30x^4 + 238x^2 + 56x + 100), \quad (3.11)$$

and a rugged potential with three main wells but several additional smaller wells in and between them (see Figure 3.1(a,d)). We observe that for both potentials the three main wells lead to three dominant eigenvalues $0 = \lambda_1 > \lambda_2 > \lambda_3 \gg \lambda_4$ of the associated generator L_{Diff} , and the second and third eigenfunctions are almost constant within the core segments of the main wells. For the rugged potential the additional smaller wells lead to larger eigenvalues $\lambda_4, \dots, \lambda_6$ that are still separated by a spectral gap from the three dominant ones.

Needless to say, not all metastable effects are caused by energy barriers, since there are also other strong dynamic separation phenomena not associated with energy barriers, the most prominent often called *entropic barriers*; see Example 3.21

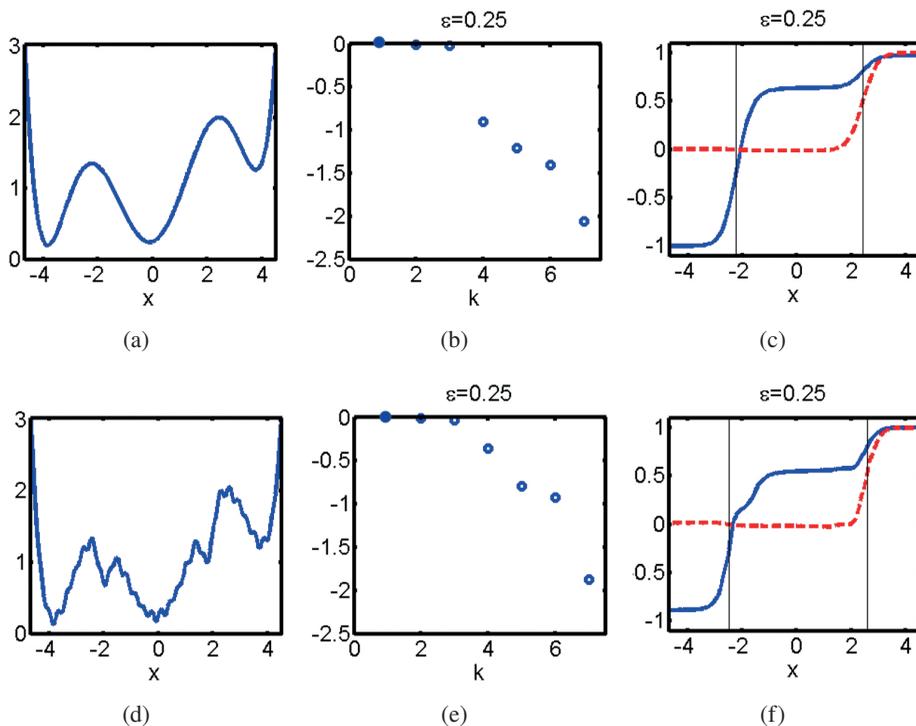


Figure 3.1. (a,d) Three-well potential (a) and rugged three-well potential (d). (b,e) Eigenvalues of the generator of the diffusive molecular dynamics process for $\gamma = 1$, $\sigma = \sqrt{2\epsilon}$ and $\epsilon = 0.25$ for the three-well potential (b) and rugged three-well potential (e). (c,f) Second (blue, solid) and third (red, dashed) eigenfunctions for the three-well potential (c) and rugged three-well potential (f). For both potentials, the third eigenfunction has a sign change close to the only zero of the second eigenfunction for $x \approx -2$. The results are based on high-accuracy finite element discretizations of L_{Diff} .

for an illustration. But even then the above ‘rules’ (R1) and (R2) seem valid in most cases by far.

Example 3.7 (Langevin dynamics). Let us now consider the Langevin molecular dynamics model (2.14) with potential energy landscape V and $\sigma = \sqrt{2\gamma}$. This process is not reversible, so we would normally utilize the singular value decomposition as detailed in Theorem 3.3, and we cannot expect the dominant eigenvalues of P^t aside from $\Lambda_0 = 1$ to be real-valued. In cases where the dominant eigenvalues have imaginary parts, we have to deal with almost periodic behaviour and cyclic probabilistic flows as discussed in Banisch, Conrad and Schütte (2015) and Conrad, Weber and Schütte (2015). However, as we will see, it is quite often the case that the dominant eigenvalues are real-valued. To this end, we consider the

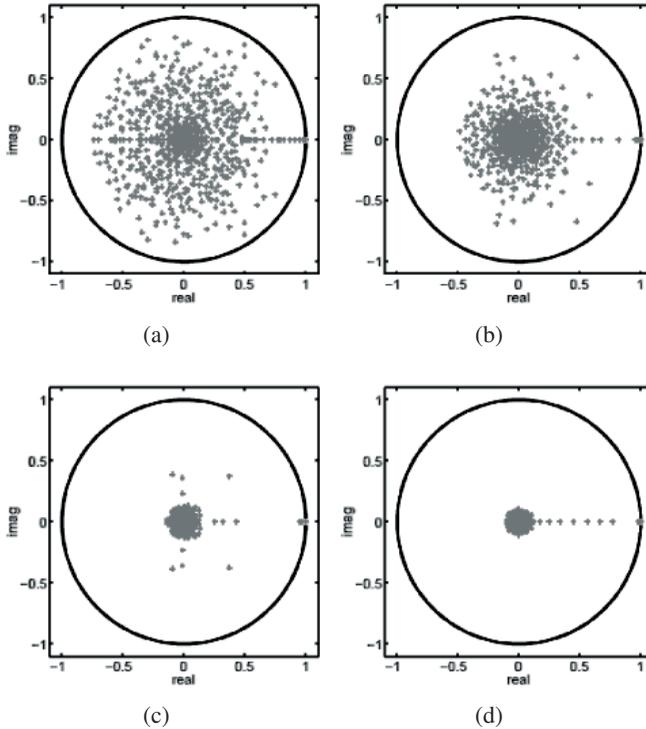


Figure 3.2. Eigenvalues (imaginary part versus real part) of the transfer operator P^τ with $\tau = 1$ for Langevin molecular dynamics (3.13) with $\beta = 2$ for an energy landscape with three wells with varying values of the friction coefficient: (a) $\gamma = 0.01$, (b) $\gamma = 0.16$, (c) $\gamma = 1.0$ and (d) $\gamma = 4.0$. For details concerning the underlying computations, see Huisinga (2001).

case $\sigma = 1$ and $\gamma = 1$ in (2.14) with energy landscape V given by the three-well potential (3.11). For $\tau = 1$, the leading eigenvalues of the associated Langevin transfer operator P^τ are given by

$$\Lambda_0 = 1.000, \quad \Lambda_1 = 0.970, \quad \Lambda_2 = 0.950, \quad \Lambda_3 = 0.440, \quad \dots$$

That is, we still observe that the leading eigenvalues are real-valued (despite the non-reversibility of the process) and that the three wells in the energy landscape result in three dominant eigenvalues with a significant gap after the third.

Figure 3.2 shows that the spectrum of the non-self-adjoint Langevin transfer operator P^τ does not in fact lie on the real line but exhibits real and imaginary components. For different values of γ and $\sigma = \sqrt{2\gamma}$ (leading to $\beta = 2$ and therefore always the same invariant measure), we observe that for small values of γ the spectrum is spread all over the unit disc, while it concentrates ever more on the interval $[0, 1]$ for larger values of γ . This reflects the fact that the Langevin

process is similar to the deterministic Hamiltonian system for $\gamma \approx 0$ (Freidlin and Wentzell 1998), while its dominant eigenvalues and eigenfunctions converge to those of diffusive molecular dynamics for large $\gamma \rightarrow \infty$ and constant $\beta = 2\gamma/\sigma^2$; see Schütte and Sarich (2014, Theorem 7).

Next we will present some rigorous results that illustrate the relation between the dominant eigenvalues and exit times and metastable decompositions. In order to reduce technicalities, we will do this for reversible processes and, where possible, without lengthy technical constructions for Langevin molecular dynamics.

Dominant eigenvalues and hitting times for small noise

We will now see how hitting times can be used to characterize the transition times between deep wells in an energy landscape and, in turn, are associated with the dominant eigenvalues of the associated transfer operator or its generator, respectively. We will consider the small noise case for diffusive and Langevin molecular dynamics. When noise is small, the transitions from one well into the others are rare events, and the basins of the wells are intuitively metastable sets of the dynamics. To make this more precise, we will compute explicit asymptotic expressions for the hitting time of the bottom of one well if the process is started at the bottom of another well. For the multi-dimensional situation we need some additional quantities and some technical assumption on V . For any disjoint sets $A, B \subset \mathbb{X}$, define the height of the saddle between A and B by

$$\hat{V}(A, B) = \inf_{\gamma \in \mathcal{P}(A, B)} \sup_{t \in [0, 1]} V(\gamma(t)),$$

where $\mathcal{P}(A, B)$ denotes the set of all continuous paths γ in \mathbb{X} with $\gamma(0) \in A$ and $\gamma(1) \in B$. Based on that, define the set of minimal points on these paths, $\mathcal{G}(A, B) = \{z \in \mathbb{X} : V(z) = \hat{V}(A, B)\}$. Moreover, we let $\mathcal{P}_{\min}(A, B)$ denote the set of minimal paths from A to B ,

$$\mathcal{P}_{\min}(A, B) = \left\{ \gamma \in \mathcal{P}(A, B) : \sup_{t \in [0, 1]} V(\gamma(t)) = \hat{V}(A, B) \right\},$$

and we let $S(A, B)$ denote the *set of saddle points* as the maximal subset of $\mathcal{G}(A, B)$ such that for every $x \in S(A, B)$ there is a minimal path $\gamma \in \mathcal{P}_{\min}(A, B)$ that goes through x .

Let the potential V be sufficiently smooth and satisfy appropriate growth conditions (see Theorem 3.4). Moreover, assume that V has finitely many minima $x \in \mathcal{M} = \{x_0, \dots, x_m\}$, and that for every minimum $x \in \mathcal{M}$ and any set $M \subset \mathcal{M} \setminus \{x\}$ of other minima, the set of saddle points $S(\{x\}, M)$ contains exactly one point of minimal energy,

$$S(\{x\}, M) = \{z(x, M)\}.$$

Further, assume that the Hessian of V in the minima x_i as well as in the saddle points $z(x, M)$ has only non-zero eigenvalues such that $\det(\nabla^2 V)$ does not vanish

in the minima or in the saddle points. Finally, define the *depth of the well* around the minimum $x \in \mathcal{M}$ regarding transitions to other minima from $M \subset \mathcal{M}$ with $x \notin M$ as

$$\Delta(x, M) = V(z(x, M)) - V(x).$$

Moreover, let B_j denote a ball of radius ϵ around the minimum x_j , set $S_k = \bigcup_{j=0}^k B_j$, and assume that ϵ is sufficiently small that $\text{dist}(z(x, M), S_k) > 0$, and that the minima can be ordered, x_0, \dots, x_m , such that the wells around the minima have different and decreasing depths:

$$\Delta(x_k, M_{k-1}) < \min_{i < k} \Delta(x_i, M_k \setminus \{x_i\}), \quad k = 1, 2, \dots,$$

where $M_k = \{x_0, \dots, x_k\}$.

Diffusive molecular dynamics. First let us consider the diffusive molecular dynamics model (2.17) with potential energy landscape V , and choose $\gamma = 1$ and $\sigma = \sqrt{2\epsilon}$,

$$\dot{x}_t = -\nabla V(x_t) + \sqrt{2\epsilon} \dot{W}_t,$$

Then the hitting times from one minimum to the other and the dominant eigenvalues of the associated generator $L_{\text{Diff}} = \epsilon \Delta_x - \nabla_x V(q) \cdot \nabla_x$ are connected and given by the following theorem Bovier *et al.* (2002c).

Theorem 3.8 (exponentially small eigenvalues for small noise diffusive MD).

Under the assumptions made above on the potential, its minima and saddle points, there are n exponentially small eigenvalues $\lambda_0 = 0 > \lambda_1 > \dots > \lambda_m$ of L_{Diff} , and there is a $\delta > 0$ such that for all $k = 1, \dots, m$ we have

$$\begin{aligned} \lambda_k &= -\frac{1}{\mathbb{E}(\rho_{x_k}(S_{k-1}))} (1 + O(e^{-\delta/\epsilon})) \\ &= -C \exp\left(-\frac{1}{\epsilon} \Delta(x_k, M_{k-1})\right) (1 + O(\epsilon^{1/2} |\log \epsilon|)), \end{aligned} \tag{3.12}$$

where $\rho_{x_k}(S_{k-1})$ is the hitting time of S_{k-1} starting at x_k and the prefactor is given by

$$C = \frac{|\eta|}{2\pi} \frac{\sqrt{|\det(\nabla^2 V(x_k))|}}{\sqrt{|\det(\nabla^2 V(z(x_k, M_{k-1})))|}},$$

where η denotes the unique negative eigenvalue of the Hessian $\nabla^2 V$ of V at the saddle point $z(x_k, M_{k-1})$.

Remark 3.9 (Kramer’s rule). For a two-well potential, equation (3.12) is identical to Kramer’s rule as given in (3.10) for $\sigma = \sqrt{2\epsilon}$.

Langevin dynamics. Let us now consider the Langevin molecular dynamics model (2.14) with potential energy landscape V , and choose $\sigma = \sqrt{2\gamma\epsilon}$,

$$\dot{x} = \xi, \quad \dot{\xi} = -\nabla_x V(x) - \gamma \xi + \sqrt{2\gamma\epsilon} \dot{W}_t, \tag{3.13}$$

with states (x, ξ) , and the L_μ^2 -generator

$$L_{Lan} = \epsilon\gamma \Delta_\xi - \xi \cdot \nabla_x + \nabla_x V \cdot \nabla_\xi - \gamma \xi \cdot \nabla_\xi.$$

A rigorous statement about the fact that the dominant eigenvalues of the Langevin process are real-valued for small noise is given in the following theorem for small $\epsilon > 0$ in (3.13); see Hérau, Hitrik and Sjöstrand (2010).

Theorem 3.10 (exponentially small eigenvalues of Langevin generator). Under the exact above conditions on the potential, its minima x_0, \dots, x_m , their enumeration and saddle points, there is an $\epsilon_* > 0$ such that for all $0 < \epsilon < \epsilon_*$ there are exactly $m + 1$ exponentially small eigenvalues $\lambda_0 = 0 > \lambda_1 > \dots > \lambda_m$ of L_{Lan} , and for all $k = 1, \dots, m$ we have

$$\lambda_k = -C \exp\left(\frac{1}{\epsilon} \Delta(x_k, M_{k-1})\right) (1 + O(\epsilon)),$$

where the prefactor is given by

$$C = \frac{|v_k|}{\pi} \frac{\sqrt{|\det(\nabla_x^2 V(x_k))|}}{\sqrt{|\det(\nabla_x^2 V(z(x_k, M_{k-1})))|}}$$

and v_k denotes the unique negative eigenvalue of the block matrix

$$\begin{pmatrix} 0 & \text{Id} \\ H_k & \gamma \text{Id} \end{pmatrix}.$$

Here H_k denotes the Hessian of V at the saddle point $z(x_k, M_{k-1})$. A statement similar to (3.12) regarding the relation between the λ_k and the hitting times for starting in one well and hitting another is also available, but technically more involved.

Limitations. While these results show the relation between dominant eigenvalues and deep wells, they have two main disadvantages in the context of MD.

- (1) The small noise case corresponds to the case that the average energy per atom is smaller than the main energy barriers, which is true for the main wells of the energy landscape but *not* for the very large number of small wells contained within each of the deep wells.
- (2) These results do not characterize the basin of attraction around a particular minimum in which the process remains for long periods of time. Specifically we are interested in the basin of attraction of ‘all’ minima within one of the main wells.

The next approach allows us to characterize these basins of attraction and does not depend on a small noise assumption.

Dominant eigenvalues and small exit rates

We consider the asymptotic decay of the conditional distribution of exit times $\tau_x(A) = \rho_x(A^c)$ from a set A ,

$$F_x(s, t) = \mathbb{P}(\tau_x(A) \geq s + t \mid \tau_x(A) \geq t)$$

for $s, t \geq 0$. $F_x(s, t)$ describes the tail of the distribution for which the exit time is larger than the so-called waiting time t . If asymptotically $F_x(s, t) \propto \exp(-\Gamma s)$ for $t \rightarrow \infty$ then the decay rate is said to be equal to Γ ; see [Donsker and Varadhan \(1975\)](#). When aiming at a definition of decay rates for sets, there is one main problem: we have to expect the decay rate to depend on the starting point, i.e. $\Gamma = \Gamma_x$. As is shown in [Huisinga et al. \(2004\)](#) and [Schütte, Huisinga and Meyn \(2003\)](#), there are specific sets for which this is the case.

To this end, for the diffusive molecular dynamics process with a potential satisfying appropriate growth conditions such that the Markov process (X_t) is \mathcal{V} -uniformly ergodic, we consider the non-probabilistic semigroup

$$T_A^t g(x) = \mathbb{E}_x[g(X_t)\mathbf{1}(\tau_x(A) > t)], \quad g \in L^\infty. \tag{3.14}$$

Then the ergodicity assumption implies that T_A^t has a unique largest eigenvalue $\eta(T_A^t)$ that asymptotically decays exponentially in t such that the limit

$$r = \lim_{s \rightarrow \infty} \eta(T_A^s)^{1/s}$$

exists. Then the *exit rate* $\Gamma(A)$ from the set A is defined as

$$\Gamma(A) = -\log r.$$

With these preparations, the following statement holds ([Huisinga et al. 2004](#), [Schütte et al. 2003](#)); see also [Pinsky \(1985\)](#) and [Bris, Lelièvre, Luskin and Perez \(2012\)](#).

Theorem 3.11 (metastable sets and associated exit rates for diffusive MD). Let us assume there exists a twice continuously differentiable eigenfunction $v: \mathbb{X} \rightarrow \mathbb{R}$ of the generator L_{Diff} with eigenvalue $\lambda < 0$ such that

$$L_{\text{Diff}} v = \lambda v, \tag{3.15}$$

and the set M is an open connected set such that

- (i) $v(x) > 0$ or $v(x) < 0$ for all $x \in M$,
- (ii) $v(x) = 0$ and $(\nabla v(x))^\top (\nabla v(x)) > 0$ for $x \in \partial M$.

Then the exit rate of this set M is $\Gamma(M) = -\lambda$, and the conditional distribution of exit times $F_x(s, t)$ from M decays asymptotically exponentially with decay rate $\Gamma = \Gamma(M)$, independent of $x \in M$. Moreover, for all open connected true subsets $A \subset M$ we have $\Gamma(A) > \Gamma(M)$.

That is, the positive and negative components of the dominant eigenfunctions determine the metastable sets with the smallest exit rates. Moreover, they allow for a hierarchical decomposition of the state space into such metastable sets: the positive and negative components of the first eigenfunction v_1 yield a decomposition into metastable sets with smaller exit rates than the second eigenfunction v_2 , and so on.

Limitations. This characterization of metastable decompositions of the state space does not require any small noise assumptions, and thus gives us relationships between the dominant eigenvalues and the metastable sets as basins of attraction of the main wells. However, sets are characterized by the zeros of the eigenfunctions, i.e. by objects that are notoriously ill-conditioned when it comes to numerical computations.

The next approach advances the concept of metastable decompositions by making it independent of the zeros of the eigenfunctions by reformulating it as an optimization problem.

Metastable decompositions and transition probabilities

The transition probability of the Markov process (X_t) between two sets A and B is given by

$$p(t, A, B) = \frac{1}{\mu(A)} \mathbb{P}_\mu[X_t \in B \mid X_0 \in A]. \quad (3.16)$$

It can be expressed directly in terms of the associated transfer operator P^t :

$$p(t, A, B) = \frac{\langle P^t \mathbf{1}_A, \mathbf{1}_B \rangle_\mu}{\langle \mathbf{1}_A, \mathbf{1}_A \rangle_\mu} = \frac{\langle \mathbf{1}_A, P_*^t \mathbf{1}_B \rangle_\mu}{\langle \mathbf{1}_A, \mathbf{1}_A \rangle_\mu}. \quad (3.17)$$

Every decomposition \mathcal{D} of the state space into $m + 1$ disjoint sets A_0, \dots, A_m induces an $(m + 1) \times (m + 1)$ transition matrix \mathcal{T} with entries $\mathcal{T}_{ij} = p(t, A_i, A_j)$ with non-negative entries and row-sums equal to 1, i.e. a stochastic matrix. Intuitively, a *metastable* decomposition, i.e. one where every set A_i is metastable, should be strongly diagonally dominant or, more precisely, the diagonal entries should be close to 1 and the off-diagonal ones very small. This should be true at least for metastability timescales t that are not too small and not too large.

Therefore we fix a timescale τ and assume that it satisfies the constraints

$$\tau_0 \ll \tau \ll \min_{i=1, \dots, m} T_i, \quad (3.18)$$

where τ_0 denotes the fastest timescale of the process (X_t) and T_i the implied timescales related to the dominant eigenvalues as defined in (3.3). Furthermore, let $P = P^\tau$ denote the associated transfer operator, and $p(A, B) = p(\tau, A, B)$ the transition probabilities.

Then each decomposition $\mathcal{D} = \{A_0, \dots, A_m\}$ has a metastability index

$$M(\mathcal{D}) = p(A_0, A_0) + \dots + p(A_m, A_m)$$

that should be close to $m + 1$ if the decomposition is metastable.

Equation (3.17) allows us to give a mathematical statement relating dominant eigenvalues, the corresponding eigenfunctions and a decomposition \mathcal{D} of the state space into subsets (Huisinga and Schmidt 2002).

Theorem 3.12 (metastable decompositions of state space). Assume that P is a self-adjoint transfer operator with $m + 1$ dominant eigenvalues, i.e. with spectrum of the form $\text{spec}(P) \subset [a, b] \cup \{\Lambda_m\} \cup \dots \cup \{\Lambda_1\} \cup \{1\}$ with $-1 < a \leq b < \Lambda_m \leq \dots \leq \Lambda_0 = 1$ and isolated, not necessarily simple eigenvalues Λ_i of finite multiplicity that are counted according to multiplicity, and denote by $u_m, \dots, u_1, \mathbf{1}$ the corresponding eigenfunctions, normalized to $\|u_k\|_2 = 1$. Consider the decomposition $\mathcal{D} = \{A_0, \dots, A_m\}$, and let Q be the orthogonal projection of $L^2(\mu)$ onto $\text{span}\{\mathbf{1}_{A_0}, \dots, \mathbf{1}_{A_m}\}$. Then the metastability index of \mathcal{D} can be bounded by

$$1 + \kappa_1 \Lambda_1 + \dots + \kappa_m \Lambda_m + c \leq M(\mathcal{D}) \leq 1 + \Lambda_1 + \dots + \Lambda_m, \quad (3.19)$$

where $\kappa_j = \|Qu_j\|_{L^2(\mu)}^2$ with Q denoting the orthogonal projection with respect to $\langle \cdot, \cdot \rangle_\mu$ onto $\text{span}\{\mathbf{1}_{A_0}, \dots, \mathbf{1}_{A_m}\}$, and $c = a(1 - \kappa_1) \dots (1 - \kappa_n)$.

This result highlights the strong relation between a decomposition of the state space into metastable subsets and dominant eigenvalues close to 1. Due to (3.18) the dominant eigenvalues Λ_i are all very close to 1, and the upper bound $1 + \Lambda_1 + \dots + \Lambda_m$ is very close to $m + 1$. Equation (3.19) states that the metastability of an arbitrary decomposition \mathcal{D} cannot be larger than $1 + \Lambda_1 + \dots + \Lambda_m$, while it is at least $1 + \kappa_1 \Lambda_1 + \dots + \kappa_m \Lambda_m + c$, which is close to the upper bound whenever the dominant eigenfunctions u_1, \dots, u_m are almost constant on the metastable subsets A_0, \dots, A_m , implying $\kappa_j \approx 1$ and $c \approx 0$. The term c can be interpreted as a correction that is small whenever $a \approx 0$ or $\kappa_j \approx 1$. Huisinga and Schmidt (2002) have demonstrated that the lower and upper bounds are sharp and asymptotically exact.

The above statement ‘the dominant eigenfunctions are almost constant on the metastable sets’ is in fact valid in many cases of importance in MD. We will reflect on the reasons for this in the next section. However, in the transition regions around the interface of one metastable set and the next, the eigenfunctions will not be constant but may exhibit sharp gradients. This does not affect $\kappa_j \approx 1$ because the invariant measure is very small in these transition regions. However, it tells us that the metastability index will not change much if the boundary between two metastable sets is changed somewhat *within* the associated transition region. That is, finding the optimal metastable decomposition, i.e. the one that maximizes $M(\mathcal{D})$, is again an ill-conditioned problem.

Limitations. The characterization of metastable decompositions by transition probabilities is different from the other characterizations because it is not based on a pathwise concept, that is, the process might exit from the set and return during time τ and this excursion would still count as metastability. Nevertheless, the approach via transition probabilities has been most influential regarding the utilization of metastability in molecular dynamics, as it was the basis of Markov state models

(MSMs), one of the most prominent computational approaches for understanding the long-term behaviour of biomolecular systems; see Section 4.2.3.

3.2. Committor functions and fuzzy metastable states

3.2.1. Committor functions

The *committor function* (Metzner, Schütte and Vanden-Eijnden 2006, 2009b), associated with a pair of disjoint sets A and B , is the probability of hitting B next if starting at x , i.e. the probability that the process (X_t) , started at x , hits B before A . In terms of the hitting times of the two sets, this can be expressed as

$$q_{AB}(x) = \mathbb{P}_x(\rho_x(B) < \rho_x(A)), \quad A \cap B = \emptyset. \quad (3.20)$$

Obviously, the committor function is mainly of interest if the two hitting times $\rho_x(A)$ and $\rho_x(B)$ are almost surely finite, for example in the case of an ergodic process with an invariant measure μ with $\mu(A), \mu(B) > 0$.

Committor functions can be considered as a special case of the potential ϕ introduced in (2.34). This can be seen if we assume $\mathbb{P}(\tau_x(A) < \infty) = 1$, set $D = \mathbb{X} \setminus (A \cup B)$ such that $D^c = A \cup B$, $c = 0$ and choose $f = \mathbf{1}_B$, the indicator function of the set B . Putting this into (2.34), we arrive at

$$\begin{aligned} \phi(x) &= \mathbb{E}[\mathbf{1}_B(X_{\tau_x(A \cup B)})] = \mathbb{P}(X_{\tau_x(A \cup B)} \in B) \\ &= \mathbb{P}(\tau_x(B) < \tau_x(A)) = q_{AB}(x). \end{aligned}$$

Since we know that the potential ϕ is given by the linear equation (2.35), we therefore also know that the committor function satisfies

$$\begin{aligned} Lq_{AB} &= 0 \quad \text{in } \mathbb{X} \setminus (A \cup B), \\ q_{AB} &= 0 \quad \text{in } A, \\ q_{AB} &= 1 \quad \text{in } B, \end{aligned} \quad (3.21)$$

where L denotes the generator of the process under consideration.

Forward and backward committor. The function q_{AB} defined above is often also called the *forward committor* $q_{AB}^+ = q_{AB}$. Analogously, we can ask for the probability $q_{AB}^-(x)$ that the process observed at x came last from A and not from B . That is, q_{AB}^+ refers to hitting set B next and q_{AB}^- refers to coming last from set A . Therefore we define the *last exit time* of a set $C = A, B$, $\tau_x^-(C) = \inf\{t > 0: X_{-t} \in C\}$, conditioned on $X_0 = x$, and get

$$q_{AB}^-(x) = \mathbb{P}_x(\tau_x^-(A) < \tau_x^-(B)).$$

The backward committor is defined in terms of the process (X_t) reversed, i.e. backwards in time. Whenever we assume reversibility of (X_t) , then the process forward in time is indistinguishable from the reversed process. Thus, under the assumption of reversibility, we find that

$$q_{AB}^-(x) = \mathbb{P}_x(\tau_x^-(A) < \tau_x^-(B)) = \mathbb{P}_x(\rho_x(A) < \rho_x(B)) = 1 - q_{AB}^+(x).$$

Definition via single states. From the definition it is not clear whether the committor function really depends on the whole sets A and B or whether it can be expressed in more general form with single states instead of sets. Lu and Vanden-Eijnden (2014) answered the question for the case of diffusive molecular dynamics where (3.21) contains the generator $L = L_{\text{Diff}}$, i.e. a drift-diffusion operator, and the process is reversible. In this case one may define a function $\theta_{a,b}: \mathbb{X} \rightarrow (-\infty, \infty)$, defined as the unique solution of (Lu and Vanden-Eijnden 2014)

$$L_{\text{Diff}} \theta_{a,b}(x) = \frac{\tau}{\mu(x)}(\delta(x - a) - \delta(x - b)), \tag{3.22}$$

where $a, b \in \mathbb{X}$, $a \neq b$, are two arbitrary states, and $\tau > 0$ an arbitrary time introduced solely for dimensional consistency. That is, $\theta_{a,b}$ satisfies $L_{\text{Diff}} \theta_{a,b} = 0$ in $\mathbb{X} \setminus \{a, b\}$. By choosing two arbitrary real numbers $\theta_A < \theta_B$ and setting $A = \{x \in \mathbb{X}, \theta_{a,b}(x) \leq \theta_A\}$ and $B = \{x \in \mathbb{X}, \theta_{a,b}(x) \geq \theta_B\}$, we find that the (forward) committor q_{AB} of A and B satisfies

$$q_{AB}(x) = \frac{\theta_{ab}(x) - \theta_A}{\theta_B - \theta_A}, \quad x \in \mathbb{X} \setminus (A \cup B). \tag{3.23}$$

This representation comes in handy in many cases where it is not clear which sets A and B to choose. Lu and Vanden-Eijnden (2014) also discussed how this construction can be generalized to Langevin molecular dynamics.

Partition of unity. Assume several disjoint sets $C_1, \dots, C_m \subset \mathbb{X}$ are given that do not form a decomposition of state space but represent, for example, the cores of the main wells in the energy landscape. Then we define the committors q_1, \dots, q_m by

$$q_i(x) = q_{D_i, C_i}(x) = \mathbb{P}_x(\rho_x(C_i) < \rho_x(D_i)), \quad D_i = \bigcup_{\substack{j=1, \dots, m \\ j \neq i}} C_j, \tag{3.24}$$

that is, the probability of starting at x and hitting C_i before any other C_j , $j \neq i$. By construction, this family of committors forms a non-negative partition of unity, that is, for all $x \in \mathbb{X}$ we have

$$\sum_{j=1}^m q_j(x) = 1, \quad q_i(x) \geq 0 \quad \text{for all } i = 1, \dots, m.$$

This partition is very useful for several constructions that we will discuss later.

Importance in theory. There are several additional characterizations of the committor function. Perhaps the most prominent theoretical example is that the committor satisfies a variational principle, since it is the unique minimizer of the so-called Dirichlet form of the process and is deeply related to its stochastic capacity in the potential-theoretic approach to metastability (Bovier and Den Hollander 2016, Bianchi and Gaudillièr 2016).

Importance for MD. The committor is the key quantity of what is called *transition path theory* (TPT) in the literature. As we will see below in more detail, it allows

us to give a characterization of the dominant transition pathways and a measure of how much they contribute to the overall transition rate between the sets A and B . The TPT framework was developed by Vanden-Eijnden and coworkers in the context of diffusions (E and Vanden-Eijnden 2004, 2006, Metzner *et al.* 2006, E and Vanden-Eijnden 2010), generalized to discrete state spaces in Metzner *et al.* (2009*b*) and Metzner (2007), and applied to molecular systems, e.g. in Noé *et al.* (2009). Subsequently it was studied as a suitable reaction coordinate, in particular for protein folding (Krivov 2018). In Lu and Vanden-Eijnden (2014), based on the representation (3.23), it is even presented as the *optimal* reaction coordinate; see Section 3.3.2 for more details.

Before we go into details of the TPT framework and reaction coordinates, let us first discuss the relation between committor functions between metastable sets, and dominant eigenfunctions of the transfer operator.

3.2.2. Committors and dominant eigenfunctions

Let us first consider the case of a process with two main metastable sets with disjoint cores C_1 and C_2 and a (non-empty) transition region $T = \mathbb{X} \setminus (C_1 \cup C_2)$ in between. Let us denote the invariant measure of the sets by $\mu_i = \mu(C_i)$. Then the two associated committors q_1 and q_2 as defined in (3.24) satisfy $q_1 + q_2 = \mathbf{1}$, and we denote $q_1 = q$ and $q_2 = 1 - q$. Furthermore, assume that the process is reversible. In this situation, the transfer operator typically has two dominant eigenvalues, $\Lambda_0 = 1$ and $\Lambda_1 < 1$, with eigenfunctions $\mathbf{1}$ and u_1 such that $\langle \mathbf{1}, u_1 \rangle_\mu = 0$ and $\langle u_1, u_1 \rangle_\mu = 1$. Now we assume that these two committors are contained in the dominant eigenspace of the transfer operator, that is, there are real-valued coefficients α, β such that $q = \alpha \mathbf{1} + \beta u_1$. These coefficients can be computed based on the properties of the committor ($= 1$ on C_1 , $= 0$ on C_2) (see Roux 2022 for details), which leads to

$$q \approx \mu_1 \mathbf{1} - \sqrt{\mu_1 \mu_2} u_1 \quad \text{and} \quad u_1 \approx -\sqrt{\frac{\mu_2}{\mu_1}} \mathbf{1}_{C_1} + \sqrt{\frac{\mu_1}{\mu_2}} \mathbf{1}_{C_2}, \quad (3.25)$$

where the approximation quality is given by Theorem 3.5.

This rough computation illustrates the deep relation between committor functions and dominant eigenfunctions in the presence of strong metastability, i.e. for energy landscapes dominated by deep wells. In order to make this more precise, let us again consider the case of some core sets C_1, \dots, C_m with smooth boundaries, and the committor functions $\{q_1, \dots, q_m\}$ defined by (3.24). To understand whether the dominant eigenfunctions can be written approximately as linear combination of the committor functions, we introduce the orthogonal projection $Q: L^2(\mu) \rightarrow \text{span}\{q_1, \dots, q_m\}$ with respect to $\langle \cdot, \cdot \rangle_\mu$, that is,

$$Qv = \sum_{k,j=1}^n (S^{-1})_{kj} \langle q_k, v \rangle_\mu q_j, \quad \text{with} \quad S_{kj} = \langle q_k, q_j \rangle_\mu.$$

Moreover, let L denote the generator associated with a reversible process (X_t) , e.g. the diffusive molecular dynamics process. Then the following holds (Schütte and Sarich 2014, Theorem 18).

Theorem 3.13 (dominant eigenfunctions in committor space). Let

$$\mathcal{T} = \mathbb{X} \setminus \bigcup_{i=1}^n C_i$$

denote the transition region between the core sets, let $\lambda < 0$ denote one of the dominant eigenvalues of L , let u be the associated normalized eigenfunction and let $T = 1/|\lambda|$ be the associated implied dominant timescale of the process. Moreover, let $\mathbb{E}_\mu[\tau_x(\mathcal{T})]$ denote the mean first exit time of the process from the transition region, let $u|_{\mathcal{T}}(x) = u(x)\mathbf{1}_{\mathcal{T}}(x)$ be the restriction of u to the transition region, and let

$$\delta u(x) = \begin{cases} 0 & x \in \mathcal{T}, \\ u(x) - \frac{1}{\mu(C_i)} \int_{C_i} u(x)\mu(x) dx & x \in C_i, i = 1, \dots, n, \end{cases}$$

be the deviation of the eigenfunction u from its average on each core set. Then the projection error of u with respect to Q satisfies

$$\|Q^\perp u\|_\mu \leq \|\delta u\|_\mu + 2\mu(\mathcal{T})\|\delta u\|_\infty + \frac{\mathbb{E}_\mu[\tau_x(\mathcal{T})]}{T} \|u|_{\mathcal{T}}\|_\mu. \tag{3.26}$$

The upper bound in (3.26) contains two parts. The first, depending on the deviation δu of the eigenfunction u from its average value on the core sets, will be very small if the core sets are chosen such that the eigenfunction u is almost constant on the core sets. The second part depends on (i) the μ -weighted average of u in the transition region, $\|u|_{\mathcal{T}}\|_\mu$, and (ii) the time needed to leave the transition region on average, $\mathbb{E}_\mu[\tau_x(\mathcal{T})]$, and the dominant timescale T associated with u . Part (i) will be very small if the eigenfunction or the invariant measure take only small values in the transition region. Part (ii) will be very small if the C_i are the cores of the main metastable sets of the system, because then $\mathbb{E}_\mu[\tau_x(\mathcal{T})]$ will be much smaller than the timescale T of transition between these metastable sets.

Example 3.14 (dominant eigenfunctions and committors). The link between dominant eigenfunctions and committors between metastable sets is nicely illustrated in Figure 3.3. There we consider diffusive molecular dynamics with $\gamma = 1$ and two $\beta = 1.67$ and $\beta = 6.67$ (larger and smaller noise) in a rugged two-dimensional energy landscape that exhibits two main wells and a smaller and less deep third well. We see that the second eigenfunction (the first one is given by $u_0 = \mathbf{1}$) is very similar to the committor function between the two main wells for both cases, larger and smaller noise.

The next result shows that this is the case for diffusive molecular dynamics with small noise.

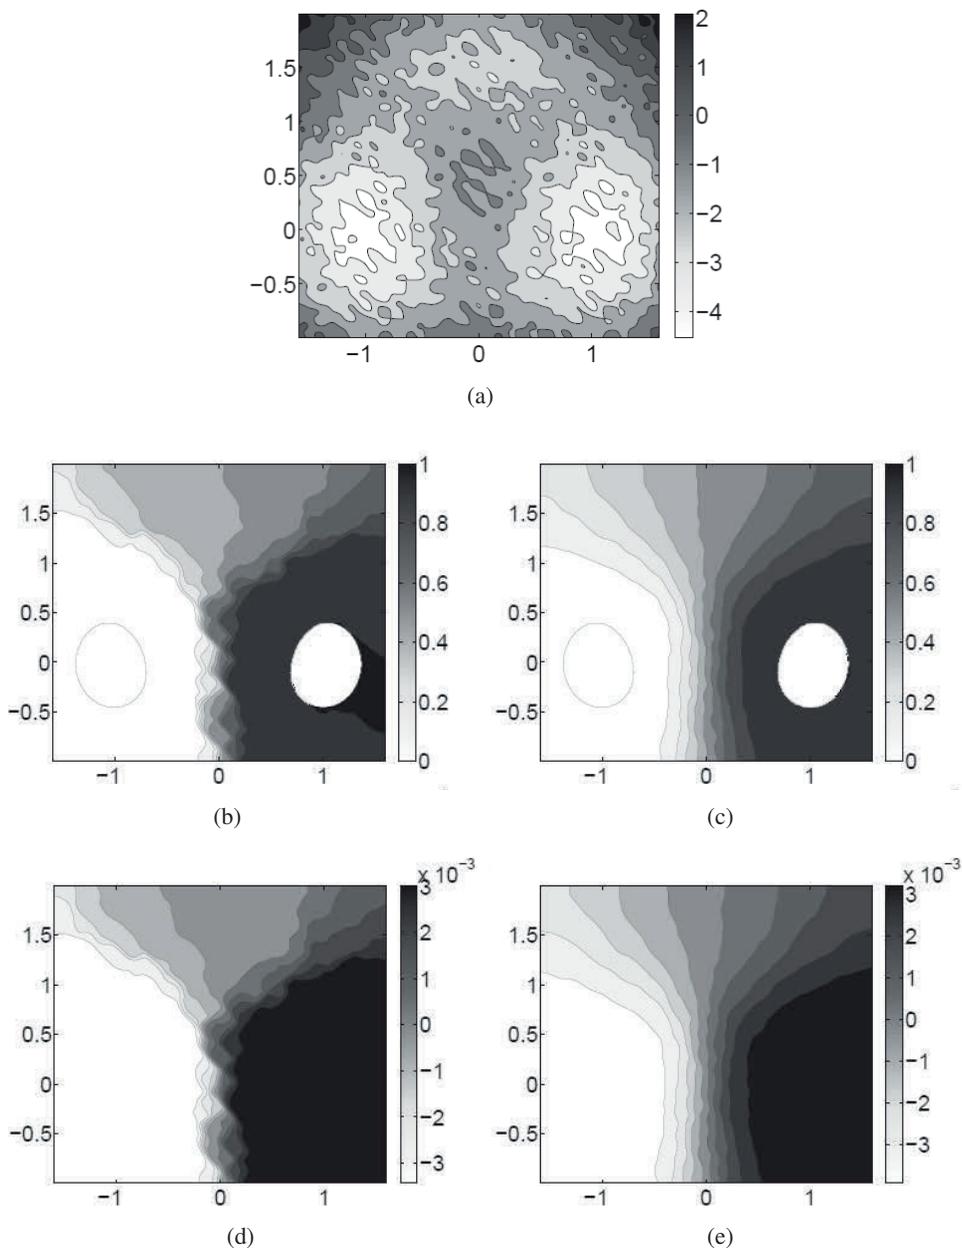


Figure 3.3. (a) Rugged three-well energy landscape V with left main wells around $(-1, 0)$ and right main well $(1, 0)$ and a less deep well around $(0, 1)$. (b–e) Committor functions q_{AB} (b,c) and associated second eigenvalues (d,e) for diffusion molecular dynamics with $\beta = 1.67$ (b,d) and $\beta = 6.67$ (c,e) and $\gamma = 1$ for the sets A (left main well) and B (right main well) for the rugged three-well energy landscape.

Small noise. Let us now consider the small noise case for the diffusive molecular dynamics model (2.17) with $\gamma = 1$ and $\sigma = \sqrt{2\epsilon}$,

$$\dot{x}_t = -\nabla V(x_t) + \sqrt{2\epsilon}\dot{W}_t,$$

and associated generator $L_{\text{Diff}} = \epsilon\Delta_x - \nabla_x V(q) \cdot \nabla_x$. We have the following theorem connecting the dominant eigenfunctions and committor functions (Bovier *et al.* 2002c).

Theorem 3.15 (eigenfunctions and committors for small noise diffusive MD).

Let the assumption on the potential (minima x_0, \dots, x_m , ordered according to decreasing well depth, saddles, energy barriers) of Theorem 3.8 be valid. Let B_i again denote a ball of radius ϵ around the minimum x_i , and $S_k = \bigcup_{i=0}^k B_i$. Let $q_k = q_{B_k, S_{k-1}}$ denote the committor function between the ball around the k th minimum and the union of balls around the minima x_0, \dots, x_{k-1} . Then there are $m + 1$ exponentially small eigenvalues $\lambda_1 = 0 > \lambda_1 > \dots > \lambda_n$ of L_{Diff} with normalized eigenfunctions $u_k \in L^2_\mu$. Furthermore, there exists a $\delta > 0$ such that, for $k = 1, \dots, m$, possibly after replacing u_k with $-u_k$,

$$u_k(y) = \frac{q_k(y)}{\|q_k\|_2} (1 + O(e^{-\delta/\epsilon})) + O(e^{-\delta/\epsilon}). \tag{3.27}$$

This result seems to contradict our previous remarks and results in so far as (3.27) seems to indicate that the leading eigenfunctions are approximately equal to a single committor function and therefore non-negative. In general, this is not the case. In the situation considered in Theorem 3.15, however, formally for $\epsilon \rightarrow 0$ and two minima, one well attracts almost all of the invariant measure such that (3.25) takes the following form: with $\mu_1 \approx 1$ and $\mu_2 \approx 0$, we get $q_{C_1, C_2} \approx \mathbf{1}_{C_2}$, $q_{C_2, C_1} \approx \mathbf{1}_{C_1}$ and $u_1 \approx (1 - \mu_1)^{-1/2} \mathbf{1}_{C_2}$, consistent with (3.27).

3.2.3. *Transition path theory*

The committor is the key quantity for understanding the reaction pathways between two disjoint sets A and B . In transition path theory, we first consider *reactive trajectories*, that is, individual realizations of the underlying process that start in A and end up in B without returning to A in between. For an ergodic process we can get an ensemble of such reactive trajectories by pruning a generic infinitely long trajectory (cutting out the reactive parts). When the invariant measure is denoted by μ , the probability of finding the process in $x \in T = \mathbb{X} \setminus (A \cup B) = (A \cup B)^c$ is $\mu(x)$, and conditioned on being at x , the probability that the process came last from A and will next go to B is given by $q(x)q^-(x)$, where $q = q_{AB}$ denotes the forward and $q^- = q^-_{AB}$ the backward committor. Therefore the probability of observing a reactive trajectory at $x \in (A \cup B)^c$ is

$$\pi_{AB}(x) = \frac{1}{Z_{AB}} \mu(x) q(x) q^-(x),$$

where $Z_{AB} = \int_T \mu(x)q(x)q^-(x)$. The probability of observing a reactive trajectory is the basic quantity based on which one can find the reactive flux from A to B and the main reaction pathways. To this end, let us first restrict our attention to diffusive molecular dynamics with $\gamma = 1$ and invariant measure $\mu(x) \propto \exp(-\beta V(x))$. Then we have reversibility and thus $q^- = 1 - q$. As shown in E and Vanden-Eijnden (2004, 2006) and Metzner *et al.* (2006), the reactive flux J_{AB} , defined by

$$J_{AB}(x) = \frac{1}{\beta} \mu(x) \nabla q(x),$$

is the divergence-free total probability flux induced by reactive trajectories. This means that, for any surface S that divides \mathbb{X} into two parts with one containing A and the other B , the reaction rate k_{AB} (i.e. the mean frequency of observing reactive trajectories) is

$$k_{AB} = \int_S J_{AB}(x)n(x) d\sigma(x),$$

where $n(x)$ denotes the normal (column) vector on $x \in S$ in the direction of B , and $d\sigma$ is the surface element on S . The (assumed) smooth boundary ∂A of A is such a dividing surface; taking a subset $R \subset \partial A$, we will find that

$$\int_R J_{AB}(x)n(x) d\sigma(x) = p_R k_{AB},$$

meaning that the percentage p_R of all reactive flux from A to B goes through R . Further, the reactive flow exhibits streamlines s , given by

$$\frac{ds(\tau)}{d\tau} = J_{AB}(s(\tau))$$

in artificial time τ . Each streamline starts at an $x \in \partial A$ and connects this state with a state $y \in \partial B$, that is, for $s(0) = x \in \partial A$ there is a $\tau_0 > 0$ such that $s(\tau_0) \in \partial B$. Taking all the streamlines that start in $R \subset \partial A$, we get a reaction tube or *reaction pathway* that connects A and B , and through which p_R per cent of the total reactive flux occurs.

Transition path theory (TPT) utilizes the quantities π_{AB} and J_{AB} , both given by the committor q , to find reaction pathways. The TPT framework was developed by Vanden-Eijnden and coworkers (see E and Vanden-Eijnden 2004, 2006) in the context of diffusive molecular dynamics, and has been generalized to discrete state spaces in Metzner *et al.* (2009b) and Metzner (2007) and applied to molecular systems, e.g. Noé *et al.* (2009).

For Langevin molecular dynamics, which is not reversible, the expressions become more involved but similar (see Metzner *et al.* 2006 for an overview): the sets A and B have to be defined in phase space (positions and momenta), and consequently π_{AB} and J_{AB} are functions on phase space too.

3.2.4. Fuzzy metastable states

As we have seen, the committor functions q_i induced by some cores $C_i, i = 1, \dots, m$ of metastable sets (see (3.24)) form a partition of unity. Its interpretation is clear: if at state x the value $q_i(x)$ tells us the probability of going next to C_i before we enter any other core $C_j, j \neq i$.

Now assume again that there are $m + 1$ dominant eigenvalues $1 = \Lambda_0 \geq \Lambda_1 \geq \dots \geq \Lambda_m$ and associated eigenfunctions u_i such that the dominant eigenspace of the transfer operator P^t is $E = \text{span}\{u_0, \dots, u_m\}$. Normally the committors are not fully contained in E . We will now proceed under the assumption that there is another basis $\{\chi_0, \dots, \chi_m\}$ of E , where the χ_i are all non-negative functions on \mathbb{X} that form a partition of unity, and χ_i is almost constant to 1 on C_i and almost 0 on each other core set.

The functions χ_i are called *fuzzy* or *soft metastable states*; the intuition behind this name is that the value $\chi_i(x)$ gives the probability with which the state x belongs to the metastable set whose core is C_i . Thus the χ_i are often also called *membership functions*. By this interpretation, χ_i represents a fuzzy set belonging to the metastable core C_i because the characterization of metastability by crisp sets is ill-conditioned in the transition region between the cores C_i .

The transition matrix \mathcal{T}^t between fuzzy metastable states $\{\chi_i, i = 0, \dots, m\}$ is meant to generalize equation (3.17). Since \mathcal{T} is the matrix representation of P^t in the basis $\{\chi_i\}$ where $\langle \chi_i, \chi_j \rangle_\mu \neq 0$ for $i \neq j$ in general, it takes the form

$$\mathcal{T}^t = \hat{P}^t M^{-1}, \quad \text{with} \quad \hat{P}^t_{ij} = \frac{\langle P^t \chi_i, \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}, \quad \text{and} \quad M_{ij} = \frac{\langle \chi_i, \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}, \quad (3.28)$$

with stochastic mass matrix M and stochastic \hat{P}^t . In Schütte and Sarich (2014, Theorem 13) we find a generalization of Theorem 3.12 to this case, where the trace of \mathcal{T}^t (i.e. the sum of its diagonal entries) is bounded from above and below by the dominant eigenvalues. The generator L associated with P^t has an analogous matrix representation

$$\mathcal{L} = \hat{L} M^{-1}, \quad \text{with} \quad \hat{L}_{ij} = \frac{\langle L \chi_i, \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}.$$

By introducing fuzzy metastable states, many things become computationally easier (see below). However, there are several theoretical problems. For example, one may ask how long the process will typically stay within χ_i . If χ_i were the indicator function of a metastable sets C_i , then the answer would be given by the expected exit time from C_i . If it is not, what can replace the exit time or exit rate? This problem is discussed by Weber and Ernst (2017), who consider the case with two metastable cores, that is, we have one fuzzy metastable state, χ , and another one, $1 - \chi$, and we have two dominant eigenvalues of the associated generator L , $0 = \lambda_0 > \lambda_1$ with eigenfunctions $\mathbf{1} = u_0, u_1$ with orthogonality $\langle u_k, u_l \rangle_\mu = \delta_{kl}$.

The fuzzy metastable state χ has the mass

$$\mu_\chi = \langle \chi, \mathbf{1} \rangle_\mu = \int \chi(x) \mu(x) dx.$$

It is spanned by the leading eigenfunctions, that is, there are real coefficients c_0, c_1 such that

$$\chi = c_0 \mathbf{1} + c_1 u_1, \quad \text{with} \quad c_0 = \langle \chi, \mathbf{1} \rangle_\mu = \mu_\chi,$$

such that

$$L\chi = \lambda_1 c_1 v_1 = \lambda_1 (\chi - \mu_\chi \mathbf{1}) = \lambda_1 (1 - \mu_\chi) \chi - \lambda_1 \mu_\chi (\mathbf{1} - \chi). \quad (3.29)$$

In this case we can compute the matrix representation of L in the basis $\{\chi, \mathbf{1} - \chi\}$ explicitly:

$$\mathcal{L} = \lambda_1 \begin{pmatrix} 1 - \mu_\chi & \mu_\chi - 1 \\ -\mu_\chi & \mu_\chi \end{pmatrix}.$$

That is, \mathcal{L} is a rate matrix (with positive off-diagonal entries and row-sum equal to 0) with eigenvalues 0 and $\lambda_1 < 0$. We observe that the transition rate from fuzzy state χ to $\mathbf{1} - \chi$ is $\mathcal{L}_{12} = |\lambda_1| (1 - \mu_\chi)$. The main idea in [Weber and Ernst \(2017\)](#) is to introduce the *holding probability*

$$p_\chi(x, t) = \chi(x) \exp(-\Gamma_\chi t), \quad \Gamma_\chi = |\lambda_1| (1 - \mu_\chi) > 0. \quad (3.30)$$

The intuition is that the exit rate Γ_χ is the same for each state x relative to its membership $\chi(x)$, in analogy to the exit rate from metastable sets as in [Theorem 3.11](#). Based on [\(3.29\)](#), we can show that p_χ satisfies

$$\partial_t p_\chi(x, t) = L p_\chi(x, t) - F_\chi(x, t) p_\chi(x, t)$$

with discount rate

$$F_\chi(x) = |\lambda_1| \mu_\chi \frac{1 - \chi(x)}{\chi(x)}.$$

By means of the Feynman–Kac theorem, this yields the pathwise representation

$$p_\chi(x, t) = \mathbb{E}_x \left[\chi(X_t) \exp \left(- \int_0^t F_\chi(X_s) ds \right) \right]. \quad (3.31)$$

That is, the discount F_χ along a realization/path at time t is 0 if $\chi(X_t) = 1$ and ∞ where $\chi(X_t) = 0$.

In comparison, this is a generalization of the pathwise functional [\(3.14\)](#) used in the approach to exit rates from a set A which takes the form

$$p_A(x, t) = \mathbb{E}_x \left[\mathbf{1}_A(X_t) \exp \left(- \int_0^t F_A(X_s) ds \right) \right],$$

with $F_A = \infty \mathbf{1}_{A^c}$ such that the discount at time t also takes value 0 if $X_t \in A$ and ∞ if $X_t \notin A$.

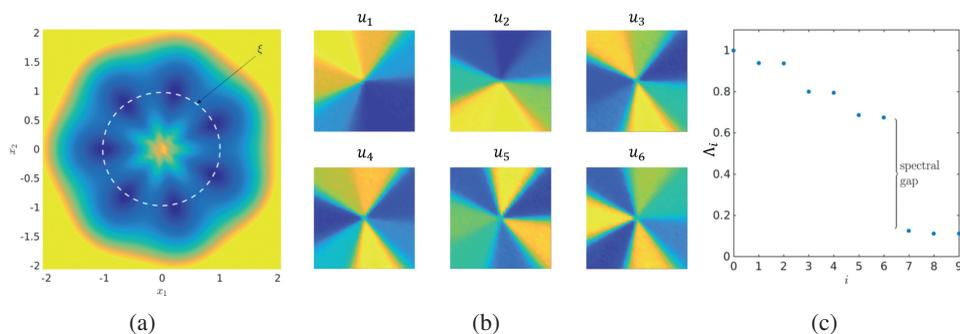


Figure 3.4. (a) Lemon-slice potential: dark blue indicates the seven deep wells in the energy landscape, light blue the transition region between them, and yellow/orange the high-energy regions. (b) Six dominant eigenfunctions u_1, \dots, u_6 of the transfer operator P^τ for diffusive molecular dynamics with $\gamma = \sigma = 1$ (eigenfunction $u_0 = \mathbf{1}$ not shown). (c) Leading eigenvalues Λ_i of P^τ exhibiting seven dominant eigenvalues (including $\Lambda_0 = 1$) and a significant gap to the remaining eigenvalues.

3.3. Reaction coordinates and transition manifolds

The notion of *reaction variables*, also called *collective variables*, is used in different versions in the literature. The basic concept is that the reaction coordinates are an abstract low-dimensional set of coordinates which represent progress along a transition pathway between an initial state and a target state that mostly belong to two different metastable sets. That is, a reaction variable is a nonlinear smooth map $\xi: \mathbb{X} \rightarrow \mathbb{R}^k$ that reduces the dimension $\mathcal{N} = \dim(\mathbb{X})$ to a significantly smaller dimension $k \ll \mathcal{N}$ with the main additional requirement that the ‘projection’ of full-dimensional molecular dynamics to the reaction coordinate allows for a ‘good’ reproduction of the long-term dynamical behaviour of the system under investigation. Figure 3.4 illustrates the basic ideas. The lemon-slice potential exhibits seven wells that are arranged around a circle. On long timescales, the dynamics will be characterized by the transitions between these wells (as can be seen from the dominant eigenvalues and eigenvectors of the associated transfer operator), and can thus be described in terms of the one-dimensional reaction coordinate $\xi(x_1, x_2) = \theta$, where θ is the polar angle. However, θ will generally not be a slow variable of the system: there is no ‘simple’ splitting into slow and fast coordinates, at least not on the short timescales. In order to get a dynamical description of the long timescales in terms of θ only, we have to understand how to project the transfer operator onto the reaction coordinate or how to design a dynamical system that describes the progress of the full dynamics in θ alone. Both options will be discussed below.

3.3.1. Reaction coordinates, marginal measure and free energy

Let $\xi: \mathbb{X} \rightarrow \mathbb{R}^k$ be a C^1 function, where $k \leq n$. Let $\mathbb{L}_z = \{x \in \mathbb{X} \mid \xi(x) = z\}$ be the z -level set of ξ . The so-called *coarea formula* (Federer 1969, Section 3.2) splits integrals over \mathbb{X} into consecutive integrals over level sets of ξ and then over the range of ξ . For $f \in L^2_\mu(\mathbb{X})$, with $L^2_\mu \subset L^1_\mu$, the coarea formula gives

$$\int_{\mathbb{X}} f(x) d\mu(x) = \int_{\xi(\mathbb{X})} \int_{\mathbb{L}_z} f(x') \varrho(x') \det(\nabla \xi(x')^\top \nabla \xi(x'))^{-1/2} d\sigma_z(x') dz, \tag{3.32}$$

where $z = \xi(x)$ and σ_z is the surface measure on \mathbb{L}_z . The *coordinate projection*, defined next, averages a given function along the level sets of a coordinate function ξ . For all $x \in \mathbb{L}_z$,

$$\begin{aligned} Q_\xi f(x) &= \int_{\mathbb{L}_z} f(x') d\mu_z(x') \\ &= \frac{1}{\Gamma(z)} \int_{\mathbb{L}_z} f(x') \varrho(x') \det(\nabla \xi(x')^\top \nabla \xi(x'))^{-1/2} d\sigma_z(x'), \end{aligned} \tag{3.33}$$

where μ_z is a probability measure on \mathbb{L}_z with density $(\varrho/\Gamma(z)) \det(\nabla \xi^\top \nabla \xi)^{-1/2}$ with respect to σ_z , and $\Gamma(z)$ is just the normalization constant so that μ_z becomes a probability measure. That is, $Q_\xi f$ is constant along the level sets \mathbb{L}_z to which x belongs. The residual projection is given by $Q_\xi^\perp = \text{Id} - Q_\xi$.

Another way to express Q_ξ uses that $Q_\xi f(x)$ is the expectation of $f(\mathbf{x}')$ with respect to μ conditional to $\xi(\mathbf{x}') = \xi(x)$, that is,

$$Q_\xi f(x) = \mathbb{E}_\mu [f(\mathbf{x}') \mid \xi(\mathbf{x}') = \xi(x)]. \tag{3.34}$$

Or, in other words, μ_z is the marginal of μ conditional to $\xi(x) = z$. Note in particular that $P_\xi f$ is itself a function on \mathbb{X} but it is constant on the level sets of ξ . The coordinate projection has the following properties: "conjugated momentum"

- (i) Q_ξ is a linear projection, i.e. $Q_\xi^2 = Q_\xi$,
- (ii) Q_ξ is self-adjoint with respect to $\langle \cdot, \cdot \rangle_\mu$,
- (iii) $Q_\xi : L^2_\mu(\mathbb{X}) \rightarrow L^2_\mu(\mathbb{X})$ is orthogonal, hence non-expansive, i.e.

$$\|Q_\xi f\|_{L^2_\mu} \leq \|f\|_{L^2_\mu}.$$

Thus $Q_\xi f$ is the average of f along the level sets of ξ , and constant on these level sets \mathbb{L}_z . Therefore $Q_\xi f$ can also be understood as a function on z space, $z \in \xi(\mathbb{X})$. For the canonical measure $\mu(dx) = Z^{-1} \exp(-\beta V(x)) dx$, a particular case is given by

$$Q_\xi \mathbf{1}(z) = \int_{\mathbb{L}_z} \mathbf{1}(x) \mu_\xi(dx)$$

and the associated *free energy* function

$$F_\xi(z) = -\frac{1}{\beta} \log Q_\xi \mathbf{1}(z). \tag{3.35}$$

Example 3.16 (free energy and potential of mean force). For the easiest case $\xi(x) = x_1$ (the first component of x), this reads (with $z = x_1$)

$$F_{\xi}(x_1) = -\frac{1}{\beta} \log \int Z^{-1} \exp(-\beta V(x)) dx_2 \dots z dx_n$$

such that

$$\frac{d}{dx_1} F_{\xi}(x_1) = \frac{\int \frac{dV}{dx_1}(x) \exp(-\beta V(x)) dx_2 \dots dx_n}{\int \exp(-\beta V(x)) dx_2 \dots dx_n},$$

which shows that F_{ξ} is identical to the potential of mean force in terms of $x_1 = \xi(x)$.

3.3.2. Committor functions as ‘optimal’ reaction coordinates

Let us again consider the diffusive molecular dynamics process

$$dX_t = -\nabla V(X_t) dt + \sigma dW_t,$$

with associated generator L_{Diff} and a scalar-valued reaction coordinate $\xi: \mathbb{X} \rightarrow \mathbb{R}$. The process (X_t) defines dynamics in a one-dimensional z -space $\xi(\mathbb{X})$ by

$$z_t = \xi(X_t),$$

and an application of Itô’s lemma yields

$$dz_t = (L_{\text{Diff}}\xi)(X_t) dt + \sigma \nabla \xi(X_t) \cdot dW_t. \tag{3.36}$$

Since the right-hand side of this equation depends on X_t and not only on z_t , the equation is not closed, which explains the key problem of dynamic coarse-graining.

However, as first pointed out in [Lu and Vanden-Eijnden \(2014\)](#), if we consider the special case of the scalar-valued committor reaction coordinate $\xi = \theta_{a,b}$ as defined in (3.22), we find that almost everywhere the process $z_t = \theta_{ab}(X_t)$ satisfies

$$dz_t = \sigma \nabla \theta_{ab}(X_t) \cdot dW_t = \sigma \sum_i \frac{\partial}{\partial x_i} \theta_{ab}(X_t) dW_{i,t}, \tag{3.37}$$

since $L_{\text{Diff}} \theta_{ab} = 0$ as long as we stay away from the two states a, b . The transformation of time t to the artificial (random) time s given by

$$s(t) = \frac{1}{2} \sigma^2 \int_0^t |\nabla \theta_{ab}(X_{t'})|^2 dt'$$

and transforms (3.37) into

$$dz_s = \sigma dW_s, \tag{3.38}$$

which is a closed equation in the committor reaction coordinate z_s !

This insight tells us that the committor RC θ_{ab} is the ‘optimal’ reaction coordinate in the sense that it allows us to reduce the dimension to 1 and reproduce the dynamical behaviour by means of a closed equation. As shown in [Lu and Vanden-Eijnden \(2014\)](#), this is even true for the more general form of diffusive

molecular dynamics (2.18) with position-dependent diffusion tensor (but only after a transformation to an artificial time), and also for Langevin molecular dynamics. However, there are several essential limitations.

Limitations. Two main bottlenecks are that there is no way to reconstruct the full state X_t from the reaction coordinate process $z_t = \theta_{ab}(X_t)$ and that the dynamics (3.38) is given in artificial time s and not in physical time. The dynamics (3.37) informs us only about the sequence in which the level sets \mathbb{L}_z are crossed; however, by averaging along the level sets of θ_{ab} ,

$$\tilde{\sigma}(z)^2 = \sigma^2 Q_\xi [|\nabla \theta_{ab}|^2](z) \quad \text{for } \xi = \theta_{ab},$$

which approximately transforms (3.37) into the equation

$$dz_t = \tilde{\sigma}(z_t) dW_t.$$

As shown in Lu and Vanden-Eijnden (2014), despite the fact that the last equation is not equivalent to (3.37), it allows us to compute the mean first passage time from one level set of θ_{ab} to another, as in other methods such as exact milestoning (Bello-Rivas and Elber 2015). Depending on the choice of a and b , there might not be a level set \mathbb{L}_z of θ_{ab} that separates a certain pair of main metastable sets; if so, the transition process between these sets will not be characterized at all. The choice of the singular states a and b is therefore essential for what emergent dynamical behaviour can be described. Moreover, the committor function is notoriously difficult to compute with sufficient precision. This fact renders many insights into properties of the committor RC theoretical and only hypothetically practical.

3.3.3. Transition manifolds

We will now utilize the transfer operator approach for characterizing good reaction coordinates. We restrict our consideration to reversible processes for the sake of simplicity. Therefore, let P^t denote a self-adjoint transfer operator in $L^2(\mu)$, and let $\xi: \mathbb{X} \rightarrow \mathbb{R}^k$ denote a reaction coordinate and Q_ξ the associated coordinate projection.

The effective transfer operator $P_\xi^t: L_\mu^2(\mathbb{X}) \rightarrow L_\mu^2(\mathbb{X})$ is defined by

$$P_\xi^t = Q_\xi P^t Q_\xi. \quad (3.39)$$

We obtain from the self-adjointness of P^t that P_ξ^t is a self-adjoint operator on $L_\mu^2(\mathbb{X})$. Moreover, $\|P_\xi^t\|_{L_\mu^2} \leq 1$. Thus the spectrum of the effective transfer operator also lies in the interval $[-1, 1]$.

Returning to the purpose of these constructions, we call ξ a *good reaction coordinate* if, for all dominant eigenvalues $\Lambda_i(t)$, $i = 0, \dots, m$, of P^t , we have an eigenvalue $\tilde{\Lambda}_i(t)$ of P_ξ^t such that

$$\Lambda_i(t) \approx \tilde{\Lambda}_i(t). \quad (3.40)$$

The reduction lies in the fact that P^t operates on functions over $\mathbb{X} \subseteq \mathbb{R}^n$, $n = \dim(\mathbb{X})$, but the effective transfer operator P^t_ξ operates *essentially* on functions over $\xi(\mathbb{X}) \subset \mathbb{R}^k$, although we embed those into \mathbb{X} through the level sets of ξ .

In [Bittracher et al. \(2018\)](#) we find the following general result on the closeness of eigenvalues of P^t and P^t_ξ .

Theorem 3.17 (approximation of dominant eigenvalues). Let P^t , P^t_ξ and Q_ξ be as above and $Q^\perp_\xi = \text{Id} - Q_\xi$. Let u with $\|u\| = 1$ be an eigenfunction of P^t , i.e. $P^t u = \Lambda(t)u$ for some $\Lambda(t) \in \mathbb{R}$. If $\|Q^\perp_\xi u\| < \varepsilon$, then P^t_ξ has an eigenvalue $\tilde{\Lambda}(t) \in \mathbb{R}$ with

$$|\Lambda(t) - \tilde{\Lambda}(t)| < \varepsilon / \sqrt{1 - \varepsilon^2}.$$

The interpretation is as follows: if there exists a reaction coordinate $\xi: \mathbb{X} \rightarrow \mathbb{R}^k$ so that the dominant eigenfunctions u are constant on the level sets of ξ , that is, if, for all dominant eigenfunctions u_i of P^t , there exist functions $\tilde{u}_i: \mathbb{R}^k \rightarrow \mathbb{R}$ such that $u_i = \tilde{u}_i \circ \xi$, then the projection error $\|Q^\perp_\xi u_i\|_{L^2_\mu}$ is zero. A generalization of this is that if the eigenfunctions u_i are *almost constant* on level sets of ξ , then the projection error is small, and the dominant eigenvalues of P^t and P^t_ξ are ε -close to each other.

This insight inspires us to look for reaction coordinates ξ for which the dominant eigenfunctions are almost constant on the level sets. The first obvious choice would be to take the dominant eigenfunctions of P^t , that is, $\xi = (\mathbf{1}, u_1, \dots, u_m): \mathbb{X} \rightarrow \mathbb{R}^{m+1}$ such that $Q^\perp_\xi u_i = 0$ for all $i = 0, \dots, m$, and perfect reproduction of dominant timescales. However, in many realistic cases, m can be too large in the sense that there is a reaction coordinate ξ with $\dim(\xi(\mathbb{X}))$ being considerably smaller than m but sufficient approximation quality regarding the dominant timescales.

This leads to the concept of *transition manifolds* as first introduced in [Bittracher et al. \(2018\)](#). To introduce the concept we first have to fix a timescale τ that is much larger than the fastest timescales and much faster than the dominant timescales T_i implied by the dominant eigenvalues of P^t :

$$t_{\text{fast}} \ll \tau \ll \min_i T_i.$$

Definition 3.18 (transition manifold). We call the process $X_t(\varepsilon, r)$ -*reducible* if there exists a smooth closed r -dimensional manifold $\mathbb{M} \subset L^2_{1/\mu} \subset L^1(\mathbb{X})$ such that for all $x \in \mathbb{X}$

$$\min_{f \in \mathbb{M}} \|f - p(\tau, x, \cdot)\|_{L^2_{1/\mu}} \leq \varepsilon. \tag{3.41}$$

We call \mathbb{M} the *transition manifold* and the map $\mathcal{Q}: \mathbb{X} \rightarrow \mathbb{M}$,

$$\mathcal{Q}(x) := \arg \min_{f \in \mathbb{M}} \|p(\tau, x, \cdot) - f\|_{L^2_{1/\mu}}, \tag{3.42}$$

the *mapping onto the transition manifold*.

Note that the $L^2_{1/\mu}$ -norm appears for technical reasons and that (since μ is a probability measure) we have

$$\|f\|_{L^1} \leq \|f\|_{L^2_{1/\mu}},$$

which means that if two transition functions are close in the $L^2_{1/\mu}$ -norm, they are also close in the L^1 -norm.

Next we see that (ε, r) -reducibility implies that dominant eigenfunctions are almost constant on the level sets of the reaction coordinate given by the transition manifold.

Theorem 3.19 (eigenvalues and eigenvectors projected to transition manifold).

Let X_t be (ε, r) -reducible and reversible, and let $\mathcal{E}: L^1(\mu) \rightarrow \mathbb{R}^{2r+1}$ be one-to-one on the transition manifold \mathbb{M} and its image. Moreover, let the reaction coordinate $\xi: \mathbb{X} \rightarrow \mathbb{R}^{2r+1}$ be given by $\xi(x) = \mathcal{E}(Q(x))$. Then, for an eigenfunction u_i of the transfer operator P^τ with associated eigenvalue Λ_i , there exists a function $\tilde{u}_i: \mathbb{M} \rightarrow \mathbb{R}$ such that

$$|\varphi_i(x) - \tilde{\varphi}_i(\xi(x))| \leq \frac{\varepsilon}{|\Lambda_i|},$$

which implies that $\|Q_\xi u_i\|_{L^2_\mu} \leq 2\varepsilon/|\Lambda_i|$. Then Theorem 3.17 implies that there is an eigenvalue $\tilde{\Lambda}_i$ of the effective transfer operator P^τ_ξ associated with $\xi = Q$ such that

$$|\Lambda_i - \tilde{\Lambda}_i| \leq 2 \frac{\varepsilon}{\sqrt{\Lambda_i^2 - 4\varepsilon^2}}.$$

The existence of such an embedding \mathcal{E} is a consequence of infinite-dimensional embedding theorems; see [Bittracher et al. \(2018, Corollary 4.11\)](#). According to these theorems, almost every bounded linear function does the job (see Theorem 4.10) and $\mathcal{E}(\mathbb{M})$ is an r -dimensional smooth manifold in \mathbb{R}^{2r+1} .

Theorem 3.19 states that (ε, r) -reducibility of the process gives us an at most $(2r+1)$ -dimensional reaction coordinate, for which the dominant eigenvalues Λ_i of the transfer operator P^t of the original dynamics (which are very close to 1 based on the choice of τ) are nicely approximated by the eigenvalues of the effective transfer operator P^τ_ξ . Thus the existence of a transition manifold implies the existence of a good reaction coordinate. This reaction coordinate allows better reproduction of the dominant eigenvalues as ε decreases; the interpretation of (3.41) is as follows: for all $x \in \mathbb{X}$, the transition function $p(\tau, x, \cdot)$ of the process started at x is close to the transition manifold \mathbb{M} with best approximation $Q(x) \in \mathbb{M}$. That is, the timescale τ is sufficiently large such that the fast equilibration from the transition region into one metastable set and the slower equilibration there towards the associated quasi-equilibrium have happened but the global equilibration between the main metastable sets has not. The transition manifold represents the backbone of the

transition events between the main metastable sets (i.e. the transition pathways), which is of considerably lower dimension than the full state space.

Bittracher and Schütte (2020) introduced a weaker concept of transition manifolds: for $\varepsilon > 0$, $r \leq n$, $\tau \in \mathbb{R}_0^+$, the process (X_t) is called *weakly* (ε, r, τ) -reducible if there exists an r -dimensionally parametrizable smooth manifold $\mathbb{M} \subset \{p(\tau, x, \cdot), x \in \mathbb{X}\}$ such that for all $x \in \mathbb{X}$

$$\int_{\mathbb{L}_z} \|\mathcal{Q}(x) - p(\tau, x', \cdot)\|_{L^2_{1/\mu}} d\mu_z(x') \leq \varepsilon, \tag{3.43}$$

where μ_z is the marginal invariant measure on \mathbb{L}_z . Any smooth manifold \mathbb{M} that fulfils (3.43) is called a *weak transition manifold*. Bittracher and Schütte (2020) showed that the above statements about approximation of the dominant eigenvalues generalize to weak transition manifolds.

3.3.4. Variational characterization of reaction coordinates

Following Bittracher *et al.* (2021), we will now introduce two seemingly different conditions for a system/reaction coordinate pair. Each condition may be taken individually as a definition of a good reaction coordinate. However, it will turn out that the two conditions are equivalent to each other for reversible systems, so a good reaction coordinate with respect to one condition is a good reaction coordinate with respect to the other.

The two new properties of the process (X_t) to be considered next are defined via conditions on its transition function $p(t, \cdot, \cdot): \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}^+$ relative to a smooth reaction coordinate $\xi: \mathbb{X} \rightarrow \mathbb{R}^r$ (ξ is C^1 and its level sets are smooth topological submanifolds of \mathbb{X}) with $\mathbb{Z} = \xi(\mathbb{X})$.

Lumpability. The process (X_t) is called ε -lumpable with respect to ξ if there is a reduced transition function $p_L(t, \cdot, \cdot): \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}^+$ and a lag time $\tau > 0$ such that

$$\frac{1}{|\mathbb{Z}|} \|p(t, \cdot, \cdot) - p_L^t(\xi(\cdot), \cdot)\|_{\mathbb{K}} \leq \varepsilon \tag{3.44}$$

for all $t \geq \tau$ and $\mathbb{K} = L^1(\mu \times \lambda)$ with Lebesgue measure λ . Here \cdot is a placeholder for the two state variables on which the respective transition functions depend. In words, lumpability means that for sufficiently large t , the transition densities $p(t, x, \cdot)$ essentially depend only on the value $\xi(x)$ of the reaction coordinate at x , and not on the precise location of x on the level set \mathbb{L}_z with $z = \xi(x)$.

Decomposability. Similarly, (X_t) is called ε -decomposable with respect to ξ if there is a reduced transition function $p_D(t, \cdot, \cdot): \mathbb{X} \times \mathbb{Z} \rightarrow \mathbb{R}^+$ and a lag time $\tau > 0$ such that

$$\frac{1}{|\mathbb{Z}|} \|p(t, \cdot, \cdot) - p_D(t, \cdot, \xi(\cdot))\mu(\cdot)\|_{\mathbb{K}} \leq \varepsilon \tag{3.45}$$

for $t \geq \tau$. That is, decomposability holds, for instance, for systems such that this transition function $p(t, x, y)$ can be decomposed into a slow transition from x to

anywhere on the level set \mathbb{L}_z , $z = \xi(y)$, for which the transition probability is given by $p_D(t, x, \xi(y))$, followed by an instantaneous equilibration on that level set with respect to the invariant density $\mu(y)$, which no longer depends on the starting point.

Whenever a system is ε -lumpable or ε -decomposable, there exists a reduced transition kernel $p_{LD}^t: \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{R}^+$ such that

$$\|p^t(\cdot, \cdot) - p_{LD}^t(\xi(\cdot), \xi(\cdot))\mu(\cdot)\|_{\mathbb{K}} \leq \varepsilon \quad (3.46)$$

for $t \geq \tau$. Under this condition, we know that ξ and p_{LD}^t allow us to approximately reconstruct the effective long-term dynamics of the full system.

Bittracher *et al.* (2021) showed that if the process is weakly (ε, r, τ) -reducible as defined in (3.43), then it is also ε -lumpable with respect to the transition manifold reaction coordinate $\xi = \mathcal{E}(\mathcal{Q}(\cdot))$, and that ε -lumpability and ε -decomposability are equivalent for reversible processes. Conclusively, we can find good reaction coordinates by looking at lumpability or deflatability. The main difficulty, however, is that the lumpability and deflatability depend on the unknown reduced transition functions p_L and p_D and we do not know how to construct them explicitly.

A variational formulation helps to overcome this obstacle. To this end, let us fix a timescale τ and set $p(x, \cdot) = p(\tau, x, \cdot)$, define the functional

$$\mathcal{L}(\vartheta) := \frac{1}{|\mathbb{Z}|} \min_{p_L: \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}^+} \|p(\cdot, \cdot) - p_L(\vartheta(\cdot), \cdot)\|_{\mathbb{K}} \quad (3.47)$$

and consider its minimizers as *optimal* reaction coordinates:

$$\xi \in \arg \min_{\vartheta \in \mathcal{C}(\mathbb{X}, \mathbb{Z})} \mathcal{L}(\vartheta), \quad (3.48)$$

This variational formulation still requires a minimization over transition functions that proves almost infeasible in practice. However, this can be avoided. To this end, we first define lumpability and decomposability loss functionals that act on the set $\mathcal{S}_r(\mathbb{Z})$ of all smooth reaction coordinates $\xi: \mathbb{X} \rightarrow \mathbb{R}^r$ such that the domain $\mathbb{Z} = \xi(\mathbb{X})$ is the same for all $\xi \in \mathcal{S}_r(\mathbb{Z})$. The *lumpability loss functional* $\mathcal{F}_L: \mathcal{S}(\mathbb{X}, \mathbb{Z}) \rightarrow \mathbb{R}^+$ and the *deflatability loss functional* $\mathcal{F}_D: \mathcal{S}(\mathbb{X}, \mathbb{Z}) \rightarrow \mathbb{R}^+$ are defined by

$$\mathcal{F}_L(\vartheta) = \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\mathbb{L}_z(\vartheta)} \int_{\mathbb{L}_z(\vartheta)} \|p(x^{(1)}, \cdot) - p(x^{(2)}, \cdot)\|_{L^1} d\mu_z(x^{(1)}) d\mu_z(x^{(2)}) dz, \quad (3.49)$$

$$\mathcal{F}_D(\vartheta) = \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\mathbb{L}_z(\vartheta)} \int_{\mathbb{L}_z(\vartheta)} \left\| \frac{p(\cdot, y^{(1)})}{\mu(y^{(1)})} - \frac{p(\cdot, y^{(2)})}{\mu(y^{(2)})} \right\|_{L^1_{\mu}} d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz \quad (3.50)$$

and, in addition to the reaction coordinate, depend solely on the transition function of the process and its invariant measure, not on any reduced transition function. They allow for the following direct variational formulation of good reaction coordinates.

Theorem 3.20 (variational principle for optimal reaction coordinates). Let the process be (ε, r, τ) -reducible as defined in (3.43). Then the process is ε -lumpable and ε -decomposable with $\mathcal{L}(\xi) = \varepsilon$ for an optimal reaction coordinate ξ satisfying (3.48), and reaction coordinates ξ_L and ξ_D that minimize the lumpability and deflatability loss functionals,

$$\xi_L \in \arg \min_{\vartheta \in \mathcal{S}_r(\mathbb{Z})} \mathcal{F}_L(\vartheta) \quad \text{and} \quad \xi_D \in \arg \min_{\vartheta \in \mathcal{S}_r(\mathbb{Z})} \mathcal{F}_D(\vartheta), \tag{3.51}$$

satisfy $\mathcal{F}_L(\xi_L) \leq 2\varepsilon$ and $\mathcal{F}_D(\xi_D) \leq 2\varepsilon$.

That is, the minimizers of the lumpability and deflatability functionals can be understood as *quasi-optimal* reaction coordinates. As we will see, this fact opens the door to efficient algorithms for computing such reaction coordinates.

3.3.5. Slow variables

In order to understand the relation between the concepts for reaction coordinates introduced above and slow collective variables, let us briefly discuss a diffusive molecular dynamics process $(X_t) = (Y_t, Z_t)$ with explicit and simple separation between fast scales (Y_t) and slow scales (Z_t) :

$$\begin{aligned} \varepsilon dY_t &= -\nabla_y V(Y_t, Z_t) dt + \sqrt{\varepsilon} \sigma dW_t^y, \\ dZ_t &= -\nabla_z V(Y_t, Z_t) dt + \sigma dW_t^z, \end{aligned} \tag{3.52}$$

with invariant measure $\mu(y, z) = \exp(-\beta V(y, z))/Z$, where $\beta = 2/\sigma^2$. The intuitive understanding is that $\xi(x) = \xi(y, z) = z$ is the (optimal) reaction coordinate if ε is small enough; more precisely, the equilibration of (Y_t) for an arbitrary fixed $Z_t = z$ (i.e. on every level set \mathbb{L}_z) must take place on an $o(1)$ timescale. The literature contains different approaches to the behaviour for small ε (stochastic averaging, homogenization, ...). In fact, with ε small enough, the process is $(\varepsilon, r = 1)$ -reducible and ε -lumpable and ε -deflatable with respect to $\xi(x) = z$; see [Bittracher et al. \(2021\)](#).

Moreover, the lumped and deflated transition function

$$p_{LD}(\xi(x), \xi(y)) = p_{LD}(z, z')$$

that satisfies (3.46) is given by the transition function of the process

$$dZ_t = -\nabla_z F(Z_t) dt + \sigma dW_t^z, \quad \text{with} \quad F(z) = -\frac{1}{\beta} \log \int \mu(y, z) dy, \tag{3.53}$$

where F is the free energy with respect to $\xi(y, z) = z$ as introduced in (3.35). This identity is the background for the widespread belief in MD that the free energy associated with a reaction coordinate allows for the derivation of dynamical information such as transition timescales via free energy barriers (which is not true in general).

Example 3.21 (fast–slow process with entropic barrier). Let us consider a two-dimensional energy landscape given by

$$V(y, z) = W(z) + \frac{1}{2}\omega(z)^2 y^2, \quad \omega(z) = 1 + 15 \exp(100(z - 0.8)^2), \quad (3.54)$$

with double well potential $W(z) = (z^2 - 1)/2$ in the z -direction with minima at $z = \pm 1$. The other direction is governed by a quadratic potential centred at $y = 0$, with a z -dependent stiffness ω^2 that assumes its maximum value at $z = 0.8$. From Figure 3.5 we see that the sharp peak at $z = 0.8$ generates a vertically tapered passage that gives rise to an entropic barrier in the horizontal direction. (Note that the potential is essentially flat in the middle of the passage, i.e. there is no energy barrier.) The resulting free energy can be computed explicitly:

$$F(z) = -\frac{1}{\beta} \log \int \mu(y, z) dy = W(z) + \frac{1}{\beta} \log \omega(z) + C,$$

with a constant C that is independent of z . If we define the conditional entropy, S_z , of the y -variable for given z as

$$S_z(y) = - \int \mu(y, z) \log \mu(y, z) dy = -\log \omega(z) - \beta C,$$

we observe that the second term in the free energy, $F = W - \beta^{-1} S_z$, that is induced by averaging out y indeed represents the *entropic contribution* to the free energy.

From Figure 3.5(b) it can be seen that F is a three-well potential with an *entropic barrier* in the free energy landscape located at $z \approx 0.8$. Note that in contrast to potential energy barriers that are easier to overcome at large noise, the opposite is true for an entropic barrier, since the entropic barrier height in the free energy is proportional to β^{-1} .

The leading eigenvalues of the generator L_{Diff} of the fast–slow dynamics (3.52) with $\epsilon = 0.1$, $\sigma = \sqrt{2/\beta}$ and $\beta = 5.5$ are

$$\lambda_0 = 0, \quad \lambda_1 = -0.0541, \quad \lambda_2 = -0.6646, \quad \lambda_3 = -2.4716,$$

showing that there are in fact three dominant eigenvalues, taking into account that the entropic contribution of the fast variable splits the right well of the double well potential into two parts, which results in three metastable sets.

Figure 3.5(c,d) shows the associated eigenfunctions u_1 and u_2 which are (i) almost constant on the three wells of the free energy, and (ii) almost constant on the level sets $\mathbb{L}_z = \{(y, z), y \in \mathbb{R}\}$ of the reaction coordinate $\xi(y, z) = z$.

3.4. Effective dynamics

Given a smooth reaction coordinate $\xi: \mathbb{X} \rightarrow \mathbb{R}^k$ for some kind of molecular dynamics process (X_t) , it defines a process $z_t = \xi(X_t)$ on reaction coordinate space $\mathbb{Z} = \xi(\mathbb{X}) \subset \mathbb{R}^k$. The process z_t is called the effective dynamics with respect to ξ .

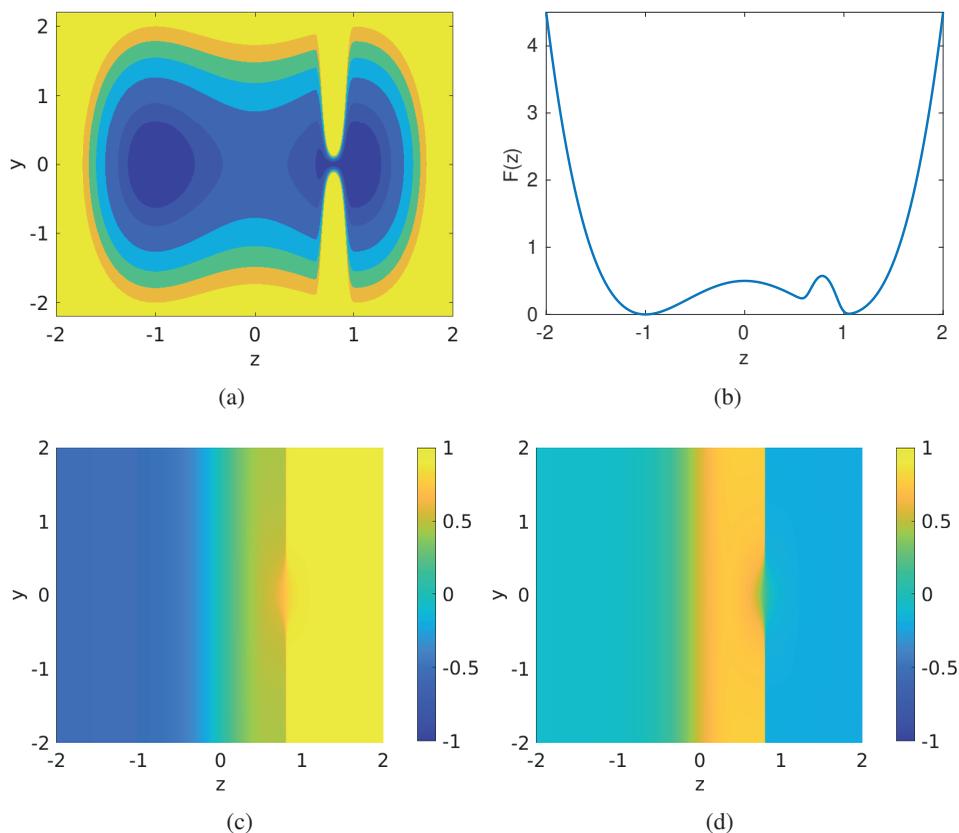


Figure 3.5. (a) Energy landscape V given in (3.54). (b) Free energy for $\beta = 5.5$. (c,d) Second and third eigenfunctions of the generator of the fast–slow process (3.52) with $\epsilon = 0.1$.

Typically there is no closed equation that describes z_t without reference to the original full process (X_t) . Therefore the phrase ‘effective dynamics’ is often used for a (closed) dynamical system whose solution process Z_t is an almost accurate approximation of z_t in particular on long timescales.

There is a huge variety of different approaches to finding such dynamical systems, for example via conditional expectations (Legoll and Lelièvre 2010), via equation-free approaches (Kevrekidis and Samaey 2009) or transfer operator theory (Froyland, Gottwald and Hammerlindl 2013), to name just a few examples. The most prominent example, however, may be the approach to effective dynamics using the Mori–Zwanzig formalism, which is an operator approach to studying the coarse-graining of dynamical systems (Mori 1965, Zwanzig 1973, Grabert 1982). To this end, let us assume that we are dealing with a process (X_t) with continuous

path, generator L and transfer operator $P^t = \exp(tL)$. Moreover, let Q_ξ denote the associated projection defined in (3.33). Following the Mori–Zwanzig approach, for an arbitrary given function $f \in D(\mathcal{L}) \subset L^2(\mu)$, we can directly verify the following identity for all $t \geq 0$:

$$\begin{aligned} \frac{d}{dt} e^{tL} f &= e^{tL} L f = e^{tL} Q_\xi L f + \int_0^t e^{(t-s)L} Q_\xi L r_s ds + r_t, \\ r_t &= e^{tQ_\xi^\perp L} Q_\xi^\perp L f \end{aligned} \quad (3.55)$$

where $Q_\xi^\perp = \text{Id} - Q_\xi$. Letting Q_ξ act on this equation, the third term vanishes, $Q_\xi r_t = 0$, and we get

$$\frac{d}{dt} Q_\xi P^t f = P_\xi^t L f + \int_0^t P_\xi^{t-s} L r_s ds, \quad (3.56)$$

where $P_\xi^t = Q_\xi P^t Q_\xi$ denotes the effective transfer operator considered above. If we follow reasoning similar to that of *optimal prediction* (Chorin, Hald and Kupferman 2000, Hijón, Español, Vanden-Eijnden and Delgado-Buscalioni 2010) to drop the second term in (3.56), we end up with an equation that expresses the evolution of a function f under the full original dynamics considered in terms of the reaction coordinate, $Q_\xi P^t f$, in a Markovian way in terms of the effective transfer operator P_ξ^t . In the next section we will see that this Markovian effective dynamics can be worked out in detail, and how good its approximation of the full dynamics might be. In Section 3.4.2, however, we will take the second term in (3.56) seriously. Memory enters into the description via the second term, and taking it into account therefore leads to non-Markovian descriptions of the effective dynamics.

3.4.1. Markovian effective dynamics

In order to see which form the Markovian description takes explicitly, let us again consider the diffusive molecular dynamics model (2.17) with potential energy landscape V , and choose $\gamma = 1$ and $\sigma = \sqrt{2/\beta}$,

$$\dot{x}_t = -\nabla V(x_t) + \sqrt{2\beta^{-1}} \dot{W}_t,$$

with invariant measure $\mu(x) \propto \exp(-\beta V(x))$ and generator L_{Diff} . Given a smooth reaction coordinate $\xi: \mathbb{X} \rightarrow \mathbb{R}^k$, the diffusive molecular dynamics process (X_t) defines a process $z_t = \xi(X_t)$ on reaction coordinate space $\mathbb{Z} = \xi(\mathbb{X}) \subset \mathbb{R}^k$ that is governed by equation (3.36), which generally does not have a closed form (i.e. it can be written down in terms of z_t alone). However, we can write down its projection onto the reaction coordinate space by means of the projection Q_ξ . The result is the SDE (Zhang, Hartmann and Schütte 2016)

$$dz_s = \tilde{b}(z_s) ds + \sqrt{2\beta^{-1}} \tilde{\sigma}(z_s) dw_s, \quad (3.57)$$

where $z_s \in \mathbb{R}^k$, w_s is a Brownian motion on \mathbb{R}^k and the coefficients $\tilde{b}: \mathbb{R}^k \rightarrow \mathbb{R}^k$, $\tilde{\sigma}: \mathbb{R}^k \rightarrow \mathbb{R}^{k \times k}$ are given by (Zhang *et al.* 2016)

$$\begin{aligned} \tilde{b}_l(z) &= Q_\xi(L_{\text{Diff}}\xi_l)(z) = Q_\xi\left(-\nabla V \cdot \nabla \xi_l + \frac{1}{\beta} \Delta \xi_l\right)(z), \\ \tilde{a}_{ll'}(z) &= (\tilde{\sigma} \tilde{\sigma}^\top)_{ll'}(z) = Q_\xi\left(\sum_{i=1}^n \frac{\partial \xi_l}{\partial x_i} \frac{\partial \xi_{l'}}{\partial x_i}\right)(z), \end{aligned} \tag{3.58}$$

for all $z \in \mathbb{Z} \subset \mathbb{R}^k$, $1 \leq l, l' \leq k$. The infinitesimal generator of the reduced process governed by (3.57) is given by

$$\tilde{\mathcal{L}} = \sum_{l=1}^k \tilde{b}_l \frac{\partial}{\partial z_l} + \frac{1}{\beta} \sum_{l,l'=1}^k \tilde{a}_{ll'} \frac{\partial^2}{\partial z_l \partial z_{l'}}, \tag{3.59}$$

which is a self-adjoint operator on the reduced Hilbert space with discrete spectrum under appropriate conditions on ξ and V . Moreover, for all $f = \tilde{f} \circ \xi$, we have

$$Q_\xi L_{\text{Diff}} f = (\tilde{\mathcal{L}} \tilde{f}) \circ \xi, \tag{3.60}$$

and the following result (Zhang and Schütte 2017).

Theorem 3.22 (long-term error of Markovian effective dynamics). Let u be a normalized eigenfunction of the operator L_{Diff} corresponding to the eigenvalue λ . Define constants

$$\delta_1 = \|L_{\text{Diff}} Q_\xi^\perp u\|_\mu, \quad \delta_2 = \|Q_\xi^\perp u\|_\mu \leq 1,$$

and suppose that $0 < \delta_2 < 1$. Then there is an eigenvalue κ of the operator $Q_\xi L_{\text{Diff}}$, such that

$$|\kappa - \lambda| \leq \frac{\delta_1}{(1 - \delta_2^2)^{1/2}}. \tag{3.61}$$

Moreover, there is an eigenvalue $\tilde{\lambda}$ of $\tilde{\mathcal{L}}$ with eigenfunction \tilde{u} such that

$$\lambda \leq \tilde{\lambda} \leq \lambda + \frac{1}{\beta} \int_{\mathbb{X}} |\nabla(u - \tilde{u} \circ \xi)(x)|^2 \mu(x) \, dx.$$

That is, when considering the reduced process (3.57) as a model for the effective dynamics, then the dominant timescales/eigenvalues of the original process are approximated only imperfectly with an error given in (3.61). For the perfect reaction coordinate, i.e. the one given by all dominant eigenfunctions of L_{Diff} , the error is zero and the approximation exact.

Remark 3.23 (fast–slow process). For reversible fast–slow processes, such as (3.52), the Markovian effective dynamics (3.57) is identical to the averaged process given by (3.53). See Hartmann, Neureither and Sharma (2020) for a discussion

of non-reversible dynamics, for which the Markovian effective dynamics and the averaged equation turn out to be different in general.

Conclusion. The Markovian effective dynamics (3.57) offers a good approximation of the full dynamics on the dominant timescales only for ‘good’ reaction coordinates for which $\delta_1 = \|L_{\text{Diff}} Q_{\xi}^{\perp} u\|_{\mu}$ is small for all dominant eigenfunctions.

3.4.2. Non-Markovian effective dynamics

In the physics-oriented literature, the Zwanzig–Mori identity (3.55) is taken as the starting point for deriving so-called *generalized Langevin equations* (GLE) as non-Markovian approximations of the effective dynamics $z_t = \xi(X_t)$. In the standard approach we start with Hamiltonian dynamics (2.12) for the molecular system with Hamiltonian $H(x, p) = p^{\top} M^{-1} p / 2 + V(x)$, written in the second-order form $M\ddot{x}_t = -\nabla V(x_t)$, and invariant measure $\mu(x, p) \propto \exp(-\beta H(x, p))$. For a scalar reaction coordinate $\xi = \xi(x)$ that depends on the positions x only, the standard approach yields, by means of several *ad hoc* assumptions and approximations (Lange and Grubmüller 2006, Ayaz *et al.* 2021),

$$m\ddot{z}_t = -\nabla_z U(z_t) - \int_0^t \Gamma(t-s)\dot{z}_s ds + F_R(t), \quad (3.62)$$

where m denotes the so-called effective mass, $U(z) = F_{\xi}(z)$ the associated free energy or potential of mean force with respect to the invariant measure μ , and F_R the so-called ‘random force’ that is connected to the memory kernel Γ via the dissipation–fluctuation principle, that is,

$$\mathbb{E}[F_R(t)F_R(s)] = \frac{1}{\beta}\Gamma(t-s).$$

This form of the GLE has often been used in practical applications; in many cases documented in the literature, it allows for a much better approximation of the ‘true’ effective dynamics $z_t = \xi(X_t)$ than available Markovian approaches such as (3.57), at least regarding collective processes on long timescales such as transition between metastable sets (barrier crossing) (Ayaz *et al.* 2021, Kappler *et al.* 2018). However, there is also a considerable number of articles in which improvements to the GLE (3.62) are discussed; see Ayaz, Scalfi, Dalton and Netz (2022) for a recent contribution that compares several approaches. The observation is that for ‘good’ reaction coordinates the memory kernel decays very fast as a function of time, while for ‘bad’ reaction coordinates it does not.

Moreover, in Hijón *et al.* (2010), the derivation starting from the original Zwanzig–Mori identity (3.55) is performed on a sound mathematical basis, yielding an equation that is a formally exact reformulation of (3.55) (Hijón *et al.* 2010, eq. (15)), from which a more general form of GLE with random forcing as in (3.62) is derived with a memory kernel that does not depend only on time as in (3.62).

4. Numerical analysis of transfer operators

In this section we will consider different methods for discretizing transfer operators and generators, that is, methods for finding finite-dimensional approximations of these operators for which we can get some form of error estimate telling us how closely the dominant timescales/eigenvalues of the finite-dimensional approximation are to the exact ones. To this end, we will focus our considerations on spatial discretization and see that different methods based on spatial discretization allow for deriving matrix representations of the respective operators with some form of *discretization error*. However, at first we will *not* discuss how to compute the entries of these matrix representations; in Section 5 we will see that the entries can be computed by means of stochastic approximations based on MD simulation *data* but that these stochastic methods introduce another form of *statistical error* that will be considered independently from the discretization error.

4.1. Spatial discretization

Due to the huge dimension of the state space \mathbb{X} , spatial discretizations of transfer operators or generators associated with realistic molecular system have to face the *curse of dimensionality*, that is, standard numerical discretization schemes lead to exponential growth of the size of the respective discretization matrix with the dimension of \mathbb{X} , rendering it practically infeasible in high dimensions. Therefore the discretization schemes discussed herein will for the most part utilize non-standard approaches, for which we will subsequently discuss whether they allow circumvention of the curse of dimensionality.

4.1.1. Spatial discretization of the Fokker–Planck equation

Spatial discretization of the Fokker–Planck equation in high dimensions requires careful consideration, especially if we wish the discretized version to inherit the main structural properties (e.g. invariant measure, reversibility, stability, stochastic interpretation). One prominent example of a spatial discretization scheme meeting these requirements is the so-called square-root approximation (SQRA) (Donati, Heida, Keller and Weber 2018, Donati, Weber and Keller 2021, Latorre, Metzner, Hartmann and Schütte 2011). In order to see how it works, we consider diffusive molecular dynamics, i.e. a reversible process. In its compact form (2.27), the Fokker–Planck equation for diffusive molecular dynamics reads

$$\partial_t u(x, t) = \frac{1}{\beta} \frac{1}{\mu(x)} \nabla \cdot (\mu(x) \nabla_x u(x, t)) = L_{\text{Diff}} u(x, t), \quad (4.1)$$

with $\mu(x) = \exp(-\beta V(x))/Z$.

Square-root approximation. SQRA is based on a partition of state space into disjoint sets A_i , $i = 1, \dots, n$, of which we assume that the A_i are convex polytopes resulting from a Voronoi discretization defined by grid point $x_i \in \mathbb{X}$ such that $x_i \in A_i$. Following Sikorski (2023), we let ∂_i denote the boundary of A_i , and

let ∂_{ij} be the boundary that is shared by neighbouring polytopes A_i and A_j , such that $\partial_i = \bigcup_{j \in N(i)} \partial_{ij}$, where $j \in N(i)$ means that A_j shares a boundary with A_i . Integration of (4.1) over A_i , using Gauss's theorem, yields

$$\begin{aligned} \beta \int_{A_i} \partial_t u(x, t) \mu(x) dx &= \int_{A_i} \nabla \cdot (\mu(x) \nabla u(x, t)) dx \\ &= \sum_{j \in N(i)} \int_{\partial_{ij}} \mu(x) \nabla u(x, t) \cdot n_{ij}(x) dS_{ij}(x), \end{aligned}$$

where n_{ij} denotes the outer unit normal of ∂_{ij} and dS_{ij} the corresponding surface element. Based on this, standard finite volume discretization is used based on the approximation of the flux across the boundary and of the volume integral,

$$\nabla u(x, t) n_{ij}(x) \Big|_{\partial_{ij}} \approx \frac{u_j(t) - u_i(t)}{h_{ij}}, \quad \int_{A_i} \partial_t u(x, t) \mu(x) dx \approx \dot{u}_i(t) \mu_i(A_i)$$

with $u_i(t) = u(x_i, t)$ and $h_{ij} = \|x_j - x_i\|$. In contrast to standard volume discretization, SQRA uses

$$\mu(x) \Big|_{\partial_{ij}} \approx \frac{1}{Z} \exp(-\beta(V(x_i) + V(x_j))/2) = \sqrt{\mu_i \mu_j}$$

as the approximation to the value of the measure at the surface ∂_{ij} , with $\mu_i = \mu(x_i)$. This yields the discrete scheme

$$\beta \dot{u}_i(t) \mu(A_i) = \sum_{j \in N(i)} (u_j(t) - u_i(t)) \frac{a_{ij}}{h_{ij}} \sqrt{\mu_i \mu_j},$$

where $a_{ij} = \int_{\partial_{ij}} dS_{ij}(x)$. With $\mu(A_i) \approx \mu(x_i) \text{Vol}(A_i)$, we get the master equation

$$\dot{u}_i(t) = \sum_{j \in N(i)} R_{ij} u_j(t) - R_{ii} u_i(t), \quad (4.2)$$

where the matrix $R = (R_{ij})$ contains the entries

$$R_{ij} = \begin{cases} C_{ij} \sqrt{\frac{\mu_j}{\mu_i}} & \text{if } j \in N(i), \\ - \sum_{j \in N(i)} R_{ij} & \text{if } i = j, \\ 0 & \text{otherwise,} \end{cases}$$

with

$$C_{ij} = \frac{1}{\beta} \frac{a_{ij}}{h_{ij} \text{Vol}(A_i)}.$$

For regular grids, $C_{ij} = C$ is the same for all index pairs i, j .

The master equation (4.2) inherits the main structural properties of the Fokker-Planck equation (4.1): R is a rate matrix (non-negative off-diagonal entries and

row-sums all equal to 0) with leading eigenvalues 0 and it satisfies the detailed balance condition, that is,

$$\hat{\mu}_i R_{ij} = \hat{\mu}_j R_{ji},$$

where $\hat{\mu}_i = \mu_i \text{Vol}(A_i)$, which is thus also the invariant measure of the master equation (4.2). Donati *et al.* (2021) and Latorre *et al.* (2011) have given several alternative derivations of SQRA and its variants (leading to different C_{ij}).

SQRA is a method to calculate transition rates as a ratio of the invariant measure of neighbouring discretization sets times a flux. Donati *et al.* (2018) used MD simulations to determine the flux across the set boundaries. Donati *et al.* (2021) used several alternative methods to calculate the exact or approximate flux for various grid types, and thus estimate the rate matrix without a simulation, showing very high accuracy and efficiency for regular grids. Latorre *et al.* (2011) demonstrated, based on numerical experiments on Lennard-Jones clusters, that the dominant eigenmodes of R are highly accurate approximations of those of the generator L_{Diff} already for rather coarse grids.

Furthermore, a deeper mathematical analysis can be found in Heida (2018), including a convergence result for SQRA for smooth potentials V .

High dimensions. Finite volume discretizations based on regular grids are standard numerical schemes and thus suffer from the curse of dimensionality. We will see later that SQRA can be used in contexts where this problem is avoided to some extent.

4.2. Projected transfer operators

4.2.1. Galerkin projection

A large class of methods uses the Galerkin projection of transfer operators and generators onto a finite-dimensional subspace $D \subset L^2(\mu)$ of dimension n . In what follows we will assume that we are given a basis $\{\phi_1, \dots, \phi_n\} \subset L^2_\mu$ of D , that is, our finite-dimensional ansatz space is spanned by the linearly independent functions ϕ_i ,

$$D = \text{span}\{\phi_1, \dots, \phi_n\}.$$

Definition 4.1. We call the basis $\{\phi_1, \dots, \phi_n\}$ a partition of unity if the following conditions are satisfied.

- (i) The ϕ_k are *non-negative*, linearly independent functions.
- (ii) The ϕ_k sum up to unity almost everywhere, i.e. $\sum_{k=1}^n \phi_k = \mathbf{1}_X$.

Now let $Q: L^2(\mu) \rightarrow D$ denote the *Galerkin projection* onto the associated finite-dimensional ansatz space D . Then Q has the form

$$Qv = \sum_{k,j=1}^n (S^{-1})_{kj} \langle \phi_k, v \rangle_\mu \phi_j, \quad (4.3)$$

where the non-negative, symmetric matrix $S \in \mathbb{R}^{n \times n}$ has entries

$$S_{kj} = \langle \phi_k, \phi_j \rangle_\mu, \quad (4.4)$$

which is invertible since it is a Gramian matrix of a set of linearly independent functions.

The special case results from a full partition of state space $\{A_k\}_{k=1, \dots, n}$, formed of disjoint, measurable, non-null sets A_k with $\mathbb{X} = \bigcup_{k=1}^n A_k$. Then the projection Q associated with the partition of unity $\{\phi_1, \dots, \phi_n\}$ with ansatz functions being the indicator functions $\phi_k = \phi_{A_k}$ of the partition sets, is called *full partition projection*. In this case we have that S is a diagonal matrix with entries

$$S_{kk} = \langle \phi_k, \phi_k \rangle_\mu = \mu(A_k) \quad \text{such that} \quad Qv = \sum_{k=1}^n \frac{1}{\mu(A_k)} \langle \phi_k, v \rangle_\mu \phi_k.$$

Next we consider a transfer operator $P^t: L^2(\mu) \rightarrow L^2(\mu)$ for some time t and write $P = P^t$ for simplicity.

Definition 4.2. Under Galerkin projection with Q , the transfer operator

$$P: L^2(\mu) \rightarrow L^2(\mu)$$

yields the projected operator $QP|_D: D \rightarrow D$. In addition we consider the projected operator $QPQ: L^2(\mu) \rightarrow D$. Both operators will be called *projected transfer operators* in what follows and we will use the common notation QPQ for simplicity. Whether $QP|_D$ or QPQ is meant will become clear from the context.

It is easy to show that projected transfer operator QPQ has the matrix representation

$$P_Q = \mathcal{P}M^{-1},$$

with the invertible mass matrix $M \in \mathbb{R}^{n \times n}$ and $\mathcal{P} \in \mathbb{R}^{n \times n}$ with entries

$$M_{kj} = \frac{S_{kj}}{\langle \phi_k, \mathbf{1} \rangle_\mu} = \frac{\langle \phi_k, \phi_j \rangle_\mu}{\langle \phi_k, \mathbf{1} \rangle_\mu} \quad \text{and} \quad \mathcal{P}_{kj} = \frac{\langle \phi_k, P\phi_j \rangle_\mu}{\langle \phi_k, \mathbf{1} \rangle_\mu}. \quad (4.5)$$

Let us now assume that the ϕ_i form a partition of unity. Then both M and \mathcal{P} are stochastic matrices. Let $e \in \mathbb{R}^n$ denote the vector with all entries being identical to 1, i.e. the vector representation of the function $\mathbf{1} \in D$. Then $P_Q e = e$. That is, P_Q also has an eigenvalue $\lambda = 1$ with the constant vector as the associated left eigenvector. In search of the associated right eigenvector, we define the vector $\hat{\mu} \in \mathbb{R}^n$ with entries

$$\hat{\mu}_k = \langle \mathbf{1}, \phi_k \rangle_\mu = \int \phi_k(x) \mu(dx) > 0, \quad (4.6)$$

and observe that $\hat{\mu}$ is also the left eigenvector associated with the eigenvalue $\lambda = 1$ of P_Q . This leads to the following result (Schütte and Sarich 2014).

Theorem 4.3 (basic properties of projected transfer operators). Let P be a transfer operator with unique, positive invariant measure μ , and let $\{\phi_1, \dots, \phi_n\}$ be a partition of unity with associated projection Q onto the subspace spanned by ϕ . Then the projected transfer operator QPQ has the matrix representation $P_Q = \mathcal{P}M^{-1}$ with the two stochastic matrices M and \mathcal{P} as defined in (4.5) and associated generalized eigenvalue problem $\mathcal{P}u = \Lambda Mu$, with a pair of stochastic matrices such that all the eigenvalues Λ of P_Q satisfy $|\Lambda| \leq 1$. Moreover, we have the following properties.

- (i) The row vector $\hat{\mu} \in \mathbb{R}^n$ defined in (4.6) is a left eigenvector corresponding to the eigenvalue $\Lambda = 1$ of P_Q , M and \mathcal{P} . The associated right eigenvector of P_Q , M and \mathcal{P} is the constant vector $e = (1, \dots, 1)^T$.
- (ii) Whenever P is self-adjoint in $L^2(\mu)$, then so is QPQ . Moreover, \mathcal{P} and M satisfy the *detailed balance condition* with respect to $\hat{\mu}$, i.e. $\hat{\mu}_k \mathcal{P}_{kl} = \hat{\mu}_l \mathcal{P}_{lk}$ for every pair $k, l \in \{1, \dots, n\}$. Hence all eigenvalues of \mathcal{P} and M are real-valued and contained in the interval $[-1, 1]$.
- (iii) Whenever Q is a full partition projection, then P_Q is a stochastic matrix too, and the statements in (ii) hold for P_Q . Then P_Q induces a Markov chain on state space $\{1, \dots, n\}$ with transition probabilities

$$P_{Q,kj} = \mathbb{P}_\mu[X_t \in A_k \mid X_0 \in A_j] = \frac{1}{\mu(A_k)} \int_{A_j} p(t, x, A_k) \mu(x) dx$$

given by the transition probabilities of the Markov process (X_t) underlying $P = P^t$ between the full partition sets A_k , $k = 1, \dots, n$.

Obviously we can also consider the projection of the generator L of (X_t) onto the space D . Analogously, the projected generator QLQ has the matrix representation $L_Q = \tilde{L}M^{-1}$ with the same mass matrix and

$$\tilde{L} = \frac{\langle \phi_k, L\phi_j \rangle_\mu}{\langle \phi_k, \mathbf{1} \rangle_\mu}. \tag{4.7}$$

Results concerning the properties of \tilde{L} similar to those of Theorem 4.3 can be found in Schütte and Sarich (2014). Moreover, the results in Donati *et al.* (2021) show that, for diffusive molecular dynamics and $\phi_i = \mathbf{1}_{A_i}$ with Voronoi cells A_i , the square-root approximation of the generator yields the same discretization matrix as (4.7).

Projection onto reaction coordinates. When given a reaction coordinate $\xi: \mathbb{X} \rightarrow \mathbb{R}^k$, a set of linearly independent ansatz functions $\{\tilde{\phi}_1, \dots, \tilde{\phi}_n\}$, $\tilde{\phi}_i: \mathbb{Z} \rightarrow \mathbb{R}$ on the reaction coordinate space $\mathbb{Z} = \xi(\mathbb{X})$ induces a set of ansatz functions $\phi_i: \mathbb{X} \rightarrow \mathbb{R}$, $i = 1, \dots, n$ by means of

$$\phi_i = \tilde{\phi}_i \circ \xi.$$

The matrix representation $P_Q = \mathcal{P}M^{-1}$ of the associated projected transfer operator QPQ is then given by

$$M_{kj} = \frac{\langle \tilde{\phi}_k, \tilde{\phi}_j \rangle_\nu}{\langle \tilde{\phi}_k, \mathbf{1} \rangle_\nu} \quad \text{and} \quad \mathcal{P}_{kj} = \frac{\langle \tilde{\phi}_k, P\tilde{\phi}_j \rangle_\nu}{\langle \tilde{\phi}_k, \mathbf{1} \rangle_\nu}, \quad (4.8)$$

where ν denotes the invariant measure of the effective dynamics $z_t = \xi(X_t)$ on the much lower-dimensional reaction coordinate space \mathbb{Z} . Consequently, a full set partition $\{B_1, \dots, B_n\}$ of the reaction coordinate space $\mathbb{Z} = \xi(\mathbb{X})$ induces a full partition $\{A_1, \dots, A_n\}$ of \mathbb{X} with indicator sets $\phi_{A_i} = \phi_{B_i} \circ \xi$, and the associated transition matrix has entries

$$\mathbb{P}_\mu[X_\tau \in A_j \mid X_0 \in A_k] = \mathbb{P}_\nu[z_\tau \in B_j \mid z_0 \in B_i]. \quad (4.9)$$

When turning to diffusive molecular dynamics and considering the matrix representation $L_Q = \tilde{L}M^{-1}$ of the projected generator QLQ induced by the basis $\{\tilde{\phi}_1, \dots, \tilde{\phi}_n\}$, we find the following interesting consequence:

$$\begin{aligned} \tilde{L} &= \frac{\langle \phi_k, L\phi_j \rangle_\mu}{\langle \phi_k, \mathbf{1} \rangle_\mu} = \frac{\langle Q_\xi \phi_k, L\phi_j \rangle_\mu}{\langle \phi_k, \mathbf{1} \rangle_\mu} = \frac{\langle \phi_k, Q_\xi L\phi_j \rangle_\mu}{\langle \phi_k, \mathbf{1} \rangle_\mu} \\ &= \frac{\langle \phi_k, (\tilde{L}\tilde{\phi}_j) \circ \xi \rangle_\mu}{\langle \phi_k, \mathbf{1} \rangle_\mu} = \frac{\langle \tilde{\phi}_k, \tilde{L}\tilde{\phi}_j \rangle_\nu}{\langle \tilde{\phi}_k, \mathbf{1} \rangle_\nu}, \end{aligned}$$

where Q_ξ denotes the projection induced by ξ and \tilde{L} is the generator of the Markovian approximation (3.57) to the effective dynamics, given in (3.59). That is, the Markovian approximation (3.57) to the effective dynamics is sufficient to compute the matrix representation of QLQ , even if it is insufficient as an approximation of $z_t = \xi(X_t)$.

Dominant eigenvalues. The following result from Schütte and Sarich (2014) shows that the dominant eigenvalues of the full transfer operator $P = P^t$ are well approximated by the dominant eigenvalues of the projected transfer operator QPQ if the dominant eigenfunction of P can be well represented in the ansatz space D (independent of whether the basis $\{\phi_1, \dots, \phi_n\}$ of D forms a partition of unity).

Theorem 4.4 (approximation of dominant eigenvalues). Let the transfer operator P be self-adjoint with spectrum $\text{sp}(P)$ such that there is a non-negative $r < 1$ so that $\text{spec}(P) \cap [-r, r]$ contains only isolated eigenvalues, among them the dominant eigenvalues $1 = \Lambda_0 > \Lambda_1 > \dots > \Lambda_m$. Let u_0, u_1, \dots, u_m be the corresponding normalized eigenvectors, and $D \subset L^2(\mu)$ a subspace with

$$\mathbf{1} \in D \quad \text{and} \quad \dim(D) =: n > m. \quad (4.10)$$

Let Q denote the orthogonal projection onto D , and let $1 = \hat{\Lambda}_0 > \hat{\Lambda}_1 > \dots > \hat{\Lambda}_m$ be the dominant eigenvalues of the projected operator QPQ . Then

$$\max_{i=1, \dots, m} |\Lambda_i - \hat{\Lambda}_i| \leq m\Lambda_1 \varepsilon^2, \quad \text{with} \quad \varepsilon = \max_{i=1, \dots, m} \|Q^\perp u_i\|_\mu. \quad (4.11)$$

Furthermore, for every isolated eigenvalue Λ of P with corresponding normalized eigenvector u and $\delta = \|Q^\perp u\|_\mu$, there is an eigenvalue $\hat{\Lambda}$ of QPQ such that

$$|\Lambda - \hat{\Lambda}| \leq \lambda_1 \delta (1 - \delta^2)^{-1/2}.$$

Remark 4.5 (projected generators). Conrad, Sarich and Schütte (2012) have given a similar result for the dominant eigenvalues of the associated generator.

Convergence. By increasing the number of ansatz functions and thus the dimension of D , a convergence of the Galerkin approximation QPQ to the real operator P in the strong operator topology (i.e. pointwise on $L^2(\mu)$) as $n \rightarrow \infty$ can generally be obtained (Korda and Mezić 2018b). However, desirable spectral convergence results typically require a convergence in operator norm (Kato 1995). The spectral convergence is therefore ultimately limited by the pointwise convergence of numerical projection methods. As an alternative, there exist RKHS-based versions where the basis functions are adapted to the data (Klus et al. 2019b). These methods allow for stronger modes of convergence than the classical projection methods (Mollenhauer and Koltai 2020). They will be discussed in more detail in Section 5 on data-driven methods.

Projection error. General estimates on the projection error ε in Theorem 4.4 are given in terms of the distance between P and QPQ . The theorem of Davis and Kahan (1970) and its generalization in Yu, Wang and Samworth (2015) show that the error will generally depend on the spectral gap: if P and E are compact and positive self-adjoint in $L^2(\mu)$, and $\Pi_m(P)$ and $\Pi_m(P + E)$ denote the orthonormal projections on the first m eigenfunctions of P and $P + E$, respectively, then

$$\|\Pi_m(P) - \Pi_m(P + E)\|_{HS} \leq \frac{2^{2/3} \sqrt{m}}{\Lambda_m - \Lambda_{m+1}} \|E\|_{L^2(\mu)}, \tag{4.12}$$

where $\|\cdot\|_{HS}$ denotes the associated Hilbert–Schmidt norm and $\Lambda_m > \Lambda_{m+1}$ the m th and $(m + 1)$ th eigenvalue of P . For a generalization to non-self-adjoint transfer operators in terms of singular values, see Mollenhauer (2022, Theorem 3.5.6). However, despite its theoretical value, this estimate is of limited use in practice since one would be required to estimate $\|E\|_{L^2(\mu)} = \|P - QPQ\|_{L^2(\mu)}$. It will, however, become handy in controlling statistical errors; see Section 5, for example.

Propagation error. In Theorem 4.4 the time t in $P = P^t$ is not explicitly specified. In general we have $(QPQ)^k \neq QP^kQ$ such that the long term evolution of functions under $(QPQ)^k$ will differ from the projected evolution under QP^kQ such that we should know about the propagation error

$$E(k) = \|(QPQ)^k - QP^kQ\|_\mu.$$

Therefore let us fix a specific lag time τ and set $P = P^\tau$. Under the same assumptions on $P = P^\tau$ and D as in Theorem 4.4, and with ε as defined in (4.11),

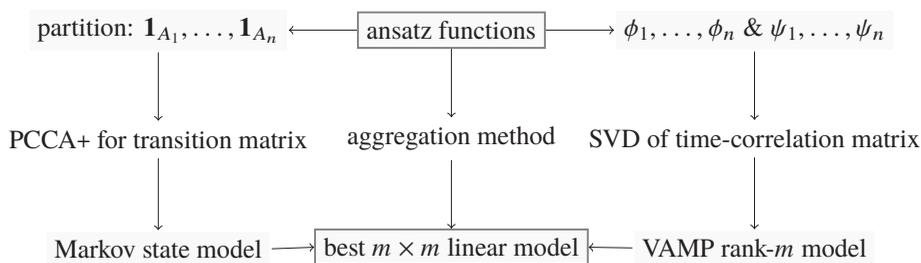


Figure 4.1. There are several approaches to aggregation of the ansatz space for finding the best low-dimensional linear model. Here two of these approaches are illustrated: on the left, Markov state model (MSM) building, where the aggregation of an $n \times n$ transition matrix using membership functions is done by PCCA+ (see Section 4.2.3); on the right, the variational approach to Markov processes (VAMP), where the aggregation is performed by singular value decomposition (SVD) of the time-correlation matrix (see Section 4.3.1).

it is shown in Schütte and Sarich (2014) and Sarich, Noé and Schütte (2010) that

$$E(k) \leq \min[2; C(k)] \cdot \Lambda_1^k, \quad (4.13)$$

with $\eta = r/\Lambda_1 < 1$ and

$$C(k) = (m\varepsilon + \eta) \left[m^{1/2}(k-1)\varepsilon + \frac{\eta}{1-\eta}(1-\eta^{k-1}) \right]. \quad (4.14)$$

Sarich *et al.* (2010) have shown that by choosing an appropriate ansatz space (with small ε) and large enough lag time τ , we can control the maximal propagation error.

4.2.2. Best linear model

We saw that given an ansatz space $D = \text{span}\{\phi_1, \dots, \phi_n\}$, the projected transfer operator $QP^\tau Q$, respectively its $n \times n$ matrix representation, is the best approximation in the sense of $\|\cdot\|_\mu$. If P^τ has just $m \ll n$ dominant eigenvalues, the natural objective is to look for ways to find the best $m \times m$ linear model. Figure 4.1 illustrates two prominent strategies: *Markov state models* (MSMs) (see Section 4.2.3) and the *variational approach to Markov processes* (VAMP) (see Section 4.3.1).

Several other approaches will be explained in detail below, e.g. the *milestoning* approach using committor functions as ansatz functions (see Section 4.2.4), the *time-lagged independent component analysis* (TICA), which maximizes the auto-correlation time of aggregated coordinates (see Section 5), and *extended dynamic mode decomposition* (EDMD), a data-based method where a least-squares approach is used for aggregation (see Section 5.1.1).

4.2.3. Markov state models

The special case of a full partition projection based on a complete partition of the state space into disjoint subsets $\{A_k\}_{k=1, \dots, n}$, $\mathbb{X} = \bigcup_{k=1}^n A_k$, associated with the

partition of unity $\{\phi_1, \dots, \phi_n\}$ spanned by the indicator functions $\phi_k = \phi_{A_k}$ of the partition sets, is the basis of the most prominent approach to discretizing transfer operators, namely *Markov state models* (MSMs). The projected transfer operator $QP^\tau Q$ is represented by the stochastic matrix P_Q^τ with entries

$$P_{Q,kj}^\tau = \mathbb{P}_\mu[X_\tau \in A_k \mid X_0 \in A_j], \tag{4.15}$$

that induces a Markov chain on the state space $\{1, \dots, n\}$ with transition probabilities $P_{Q,kj}^\tau$. The phrase ‘Markov state models’ originates from this fact.

In the previous section we mentioned that in general

$$(P_Q^\tau)^k \neq (P^{k\tau})_Q,$$

which is due to non-Markovian behaviour. Let us illustrate this by means of the discrete *index process* $(\tilde{X}_k)_{k \in \mathbb{N}}$ on the finite index space $\{1, \dots, n\}$, by setting

$$\tilde{X}_k = i \iff X_{k\tau} \in A_i. \tag{4.16}$$

Here (\tilde{X}_k) describes the snapshot dynamics of the continuous process (X_t) with lag time τ between the sets A_1, \dots, A_n . The entries of the transition matrix P_Q^τ are then given by $P_{Q,ij}^\tau = \mathbb{P}_\mu[\tilde{X}_{k+1} = j \mid \tilde{X}_k = i]$. However, in general, the index process (\tilde{X}_k) is not Markovian, that is,

$$\mathbb{P}_\mu[\tilde{X}_{k+1} = j \mid \tilde{X}_k = i_k, \tilde{X}_{k-1} = i_{k-1}, \dots, \tilde{X}_0 = i_0] \neq \mathbb{P}_\mu[\tilde{X}_{k+1} = j \mid \tilde{X}_k = i_k].$$

Example 4.6 (two deep wells). This can be understood on the basis of an energy landscape with just two deep wells and two sets A and B around these wells that form a full partitioning of state space. When considering a lag time τ that is small compared to the expected transition time from A to B , the knowledge of $X_{(k-1)\tau} \in A$ implies that at time $k\tau$ most trajectories that arrived in set B are still close to set A because the lag time τ is not large enough. That is, they are still inside the transition region and not close enough to the minimum in set B , as is to be expected with overwhelming probability if we just know $X_{k\tau} \in B$. Therefore

$$\mathbb{P}_\mu[X_{(k+1)\tau} \in A \mid X_{k\tau} \in B, X_{(k-1)\tau} \in A] > \mathbb{P}_\mu[X_{(k+1)\tau} \in A \mid X_{k\tau} \in B].$$

In other words, if τ is too small, the effect of *re-crossing* causes memory, thus making the index process non-Markovian.

This re-crossing problem has several other undesired consequences. For example, let us consider the implied dominant timescales T_i of the process (X_t) as defined in (3.3). The T_i are given by means of the dominant eigenvalues $\Lambda_i(\tau)$ of the associated transfer operator P^τ but do *not* depend on the chosen timescale τ . Naturally we aim to approximate the T_i via the implied dominant timescales $T_{Q,i}(\tau)$ associated with $QP^\tau Q$. To this end, we define

$$T_{Q,i}(\tau) = -\frac{1}{\tau} \log(\hat{\Lambda}_i(\tau)), \tag{4.17}$$

where $\hat{\Lambda}_i(\tau)$ are the dominant eigenvalues of $QP^\tau Q$, respectively, of its matrix representation P_Q . When computing the $T_{Q,i}(\tau)$, we observe that they increase strongly with τ for small values of τ and then slowly (but monotonically) reach a plateau for large values of τ . The plateau value $T_{Q,i}$ is an approximation of the dominant timescales T_i of the full process with an accuracy depending on the accuracy of the dominant eigenvalues and therefore on the projection error of the chosen subspace D ; see Theorem 4.4. However, the increase of $T_{Q,i}(\tau)$ with τ is an effect of the re-crossing phenomenon and thus of the non-Markovianity of the index process (\tilde{X}_k) ; the plateau $T_{Q,i}$ is reached for values of τ for which this problem becomes less and less of an issue, that is, as soon as (\tilde{X}_k) becomes closer and closer to Markovianity.

Perron-cluster cluster analysis (PCCA). PCCA (Deuffhard and Weber 2005) exploits the structure of the eigenvectors in order to define the metastable (long-lived) states of a Markov chain. PCCA+ and its variants (Röblitz and Weber 2013, Frank, Sikorski and Röblitz 2022) are advanced methods that find an optimal linear transformation of the eigenvector into membership functions. In particular, PCCA+ ensures that all membership functions are non-negative, and partitions of unity. Using PCCA+ or PCCA++ with a Markov model is a common approach to finding long-lived states in molecular dynamics trajectories.

PCCA starts from a Markov state model with transition matrix P_Q of the form (4.15), originating from a set partition A_1, \dots, A_n of the state space. If P_Q results from a self-adjoint transfer operator, that is, if it is reversible in the sense of Theorem 4.3(ii), then the dominant subspace of P_Q is spanned by its first m eigenvectors, that is, if the matrix $U \in \mathbb{R}^{n \times m}$ contains the dominant eigenvectors in its columns and $\Lambda \in \mathbb{R}^{m \times m}$ denotes the diagonal matrix of its m dominant eigenvalues, then $U^\top D_\mu U = \text{Id}$ with $D_\mu = \text{diag}(\mu(A_i))$ in accordance with orthonormality with respect to $\langle \cdot, \cdot \rangle_\mu$, and

$$P_Q U = U \Lambda. \quad (4.18)$$

If P_Q is not reversible, then we do not have orthonormal eigenvectors and real-valued eigenvalues. However, the decomposition (4.18) is still possible based on the Schur decomposition, that is, U is an orthonormal matrix ($U^\top D_\mu U = \text{Id}$) whose columns are called the Schur vectors, and Λ is an upper quasi-triangular (1×1 and 2×2 blocks on its diagonal) matrix, called the Schur form. PCCA uses the Schur reordering algorithm (Kressner 2006, Röblitz and Weber 2013). The columns of U are called the Schur vectors of P_Q . The eigenvalues of P_Q appear on the diagonal of Λ , where complex conjugate eigenvalues correspond to the 2×2 blocks. The first eigenvalue 1 always has the eigenvector $e = (1, \dots, 1)^\top$, thus $U_{i,1} = 1$ for all $i = 1, \dots, n$.

PCCA aims to find a linear transformation matrix $A \in \mathbb{R}^{m \times m}$ such that the column of U (the eigenvectors or Schur vectors) is transformed into m membership

vectors. To this end, we set

$$\chi = UA,$$

such that the m column vectors of $\chi = (\chi_1, \dots, \chi_m) \in \mathbb{R}^{n \times m}$, $\chi_i = (\chi_{ji})_{j=1, \dots, n} \in \mathbb{R}^n$, are membership vectors, that is, they are non-negative and form a partition of unity:

$$\chi_{ij} \geq 0, \tag{4.19}$$

$$\sum_{i=1}^n \chi_{ji} = 1 \quad \text{for all } j = 1, \dots, m. \tag{4.20}$$

The set of feasible matrices $\mathcal{A} \subset \mathbb{R}^{n \times n}$ forms a convex polytope with constraints originating from (4.19) and (4.20) such that $A \in \mathcal{A}$ if and only if

$$A_{1j} \geq - \sum_{k=2}^m U_{ik} A_{kj}, \quad i = 1, \dots, n, \quad j = 1, \dots, m, \tag{4.21}$$

$$A_{i1} = \delta_{i1} - \sum_{j=2}^m A_{ij}, \quad i = 1, \dots, m. \tag{4.22}$$

Röblitz and Weber (2013) argue that the freedom in choosing $A \in \mathcal{A}$ should be used to find membership vectors that are as crisp as possible, i.e. as close as possible to characteristic functions. This leads to the maximization problem

$$\max_{A \in \mathcal{A}} I(A), \quad \text{with} \quad I(A) = \sum_{i,j=1}^m \frac{A_{ij}^2}{A_{1j}}, \tag{4.23}$$

which is then solved using an efficient optimization algorithm based on the Schur reordering algorithm that takes the side constraints (4.21) and (4.22) into account by projection techniques; see Röblitz and Weber (2013) and Sikorski (2015) for details. Moreover, there are several extensions and variants, for example to take into account specifics of non-reversible cases; see Frank *et al.* (2022).

Resulting MSM. What did we achieve? We computed the transition matrix P_Q , given by (4.15), associated with the transfer operator P^τ based on the set partition $\{A_1, \dots, A_n\}$ of state space. If τ is large enough, the number of dominant eigenvalues of P_Q , denoted by m , is taken as a good approximation of the m leading eigenvalues of P^τ . Next we use PCCA+ to find the aggregation of P_Q into an $m \times m$ transition matrix with the same m eigenvalues. To this end, we compute the optimal matrix A (solution of (4.23)) and the associated collection of membership vectors, the columns of $\chi = UA$. These vectors allow us to define *membership functions* in state space,

$$\tilde{\chi}_i(x) = \sum_{j=1}^m \chi_{ji} \mathbf{1}_{A_j}(x),$$

that form a partition of unity in \mathbb{X} . Taking these functions as ansatz functions of a Galerkin projection, we get the following matrix representation of the associated projected transfer operator:

$$P_\chi = \mathcal{P}M^{-1}, \quad \mathcal{P}_{ij} = \frac{\langle \tilde{\chi}_i, P^\tau \tilde{\chi}_j \rangle_\mu}{\langle \tilde{\chi}_i, \mathbf{1} \rangle_\mu}, \quad M = \frac{\langle \tilde{\chi}_i, \tilde{\chi}_j \rangle_\mu}{\langle \tilde{\chi}_i, \mathbf{1} \rangle_\mu}.$$

With $D_\mu = \text{diag}(\mu(A_i))$, we find that P_χ can be computed from P_Q directly,

$$P_\chi = (\chi^\top D_\mu P_Q \chi) (\chi^\top D_\mu \chi)^{-1}, \quad (4.24)$$

and by inserting $\chi = AU$ and $P_Q U = U\Lambda$ immediately

$$P_\chi = (A^\top \Lambda A) (A^\top A)^{-1},$$

which shows that P_χ and P_Q have the same m dominant eigenvalues, the diagonal entries of Λ . Moreover, the projection error related to P_χ is the same as that with respect to P_Q . Therefore P_χ in (4.24) is called the resulting Markov state model (MSM) with m metastable states, the best $m \times m$ approximation of P_Q in the sense of (4.23).

4.2.4. Projection on committor functions: milestoneing

There are several ways to circumvent the re-crossing problem of standard MSMs. The most natural one uses a partition of unity that connects to the dynamics of the underlying process with respect to its metastable sets and avoids transition regions between the metastable sets having to be intersected by the boundaries between the partition sets A_i of the MSM. To this end, assume several disjoint sets $C_1, \dots, C_n \subset \mathbb{X}$ are given that do not form a decomposition of state space but represent, for example, the cores of the main wells in the energy landscape. Then we define the committors q_1, \dots, q_n for these core sets according to (3.24). Thus, for every core set C_i , the associated committor function $q_i(x)$ is the probability that the process will hit the core set C_i next among the other core sets when the process is started in state x . Obviously, the committor functions form a partition of unity under the condition that the process is sufficiently ergodic and will hit almost surely one core set in finite time.

Schütte *et al.* (2011) introduced a projection onto the committor space

$$D = \text{span}\{q_1, \dots, q_n\}$$

and analysed the associated projected transfer operator P . From Section 4.2.1 it is clear that the matrix representation of the projected transfer operator QPQ has the form $P_Q = \mathcal{P}M^{-1}$ with

$$M_{kj} = \frac{\langle q_k, q_j \rangle_\mu}{\langle q_k, \mathbf{1} \rangle_\mu} \quad \text{and} \quad \mathcal{P}_{kj} = \frac{\langle q_k, Pq_j \rangle_\mu}{\langle q_k, \mathbf{1} \rangle_\mu}. \quad (4.25)$$

These expressions seem to indicate that we first have to compute all m committor functions before being able to compute the representation matrices M and \mathcal{P} . This is not the case, as stated by the following theorem (Sarich 2011, Schütte *et al.* 2011).

Theorem 4.7 (milestoning process). Let $P = P^\tau$ be the transfer operator of the Markov process (X_t) for given lag time τ . Let Q be the orthogonal projection onto the space $D = \{q_1, \dots, q_n\}$ spanned by the committor functions with respect to some core sets C_1, \dots, C_n , and let $P_Q = \mathcal{P}M^{-1}$ denote the matrix representation of QPQ as given in (4.25). Moreover, let (\tilde{X}_k^-) and (\tilde{X}_k^+) be the backward and forward discrete-time milestoning processes defined by

$$\tilde{X}_k^- = i \iff X_{\sigma(k\tau)} \in C_i, \quad \text{with} \quad \sigma(t) = \sup_{s \leq t} \left\{ X_s \in \bigcup_{k=1}^n C_k \right\} \tag{4.26}$$

and

$$\tilde{X}_k^+ = i \iff X_{\sigma^+(k\tau)} \in C_i, \quad \text{with} \quad \sigma^+(t) = \inf_{s \geq t} \left\{ X_s \in \bigcup_{k=1}^n C_k \right\}. \tag{4.27}$$

Then

$$M_{ij} = \mathbb{P}[\tilde{X}_k^+ = j \mid \tilde{X}_k^- = i] \tag{4.28}$$

and

$$\mathcal{P}_{ij} = \mathbb{P}[\tilde{X}_{k+1}^+ = j \mid \tilde{X}_k^- = i]. \tag{4.29}$$

The milestoning processes are index processes that simply keep track of the last and next core set hits: $\tilde{X}_k^- = i$ means that the index of the core set where the process last came from before time $k\tau$ was i . So (\tilde{X}_k^-) always stays in a state i until the original process hits another core set $C_j, j \neq i$. Similarly, $\tilde{X}_k^+ = j$ holds if the index of the core set that is hit next after time $k\tau$ is j .

Thus Theorem 4.7 shows that when projecting onto a committor space, one can construct a matrix representation P_Q for the projected transfer operator QPQ using transition probabilities between the core sets as given by the milestoning processes. This leads to a strong computational advantage because the committor functions *never* need to be computed, and nor does any inner product.

According to Theorem 4.4, the deviation of the eigenvalues $\hat{\Lambda}_i$ of the projected transfer operator QPQ from the true eigenvalues Λ_i can be estimated according to $\max_i |\Lambda_i - \hat{\Lambda}_i| \leq n\Lambda_1 \max_i \|Q^\perp u_i\|_\mu$, with the projection error $\|Q^\perp u_i\|_\mu$ of the dominant eigenfunctions u_i of P . For the subspace spanned by committor functions, we have already seen in Theorem 3.13 that for reversible processes this projection error can be understood more precisely. The bound given there means that

$$\max_i \|Q^\perp u_i\|_\mu \ll 1$$

if the core sets C_i form the cores of the main metastable sets, i.e. the centres of the basins of attraction of the deepest wells in the energy landscape, such that

- the dominant eigenfunctions u_i are almost constant on the core sets,
- the mean first exit time from the transition region $\mathbb{X} \setminus (\cup_i C_i)$ is much smaller than the dominant implied timescales T_i associated with the dominant eigenvalues, and
- the invariant measure of the transition region is small.

Remark 4.8 (milestoning). Milestoning was introduced as a computational approach in a different context, without reference to committor functions or transfer operators, in Faradjian and Elber (2004) and Bello-Rivas and Elber (2015). It starts by choosing a set of non-intersecting, sequential surfaces, called milestones, and initiating short trajectories from each milestone, which are terminated when they reach an adjacent milestone for the first time. From the average duration of these trajectories and the probabilities of where they terminate, a rate matrix can be constructed and then used to calculate the mean first-passage time (MFPT) between any two milestones (Berezchkovskii and Szabo 2019).

4.3. Variational approaches

The min-max formula or Rayleigh–Ritz principle allows us to compute the largest eigenvalues of a self-adjoint, positive definite operator by means of a *variational principle*.

For example, let L_{Diff} again denote the generator of the diffusive molecular dynamics process, let the growth conditions (3.7) on the potential energy function V be satisfied and let $\lambda_k < 0$ denote the k th largest eigenvalue after $\lambda_0 = 1$. Then

$$-\lambda_k = \min_{H_{k+1}} \max_{\substack{f \in H_{k+1} \\ \|f\|_\mu = 1}} \langle f, -L_{\text{Diff}} f \rangle_\mu,$$

where the minimum runs over all $(k+1)$ -dimensional subspaces H_{k+1} of the Sobolev space \mathcal{H}^1 (Zhang *et al.* 2022). This statement shows that the projected generator of (4.7), since it belongs to a specific finite-dimensional subspace, always has eigenvalues that approximate the exact eigenvalues from below.

Furthermore, the Galerkin approach, based on fixing a specific finite-dimensional subspace D , is fundamentally insufficient in the sense of *parametric models* in statistics, the parameter choice being the fixed ansatz functions spanning the ansatz space. Instead one should look for *non-parametric* alternatives, where the optimal subspace itself is becoming part of a variational formulation. In the context of molecular dynamics, this idea was first introduced in Noé and Nüske (2013). However, most of the subsequent approaches use variational principles not for eigenvalues but for singular values instead, thus allowing for non-reversible cases. They aim to find the best finite-dimensional linear representation of the given transfer operator without specifying a subspace in advance.

We will follow the so-called *variational approach for Markov processes* (VAMP) (Mardt *et al.* 2018, Wu and Noé 2020). For the sake of simplicity we restrict our attention to transfer operators $P = P^t$ in $L^2(\mu)$ that are Hilbert–Schmidt operators, and how to best approximate them by low-rank operators. To this end, we start in a slightly more general setting in which we consider linear operators $A : H \rightarrow F$ between two arbitrary separable Hilbert spaces.

A bounded operator $A : H \rightarrow F$ is said to be r -dimensional if $\text{rank}(A) = r$. If $r < \infty$, we say that A is finite-rank. The given operator $A : H \rightarrow F$ is finite-rank with $\text{rank}(A) = r$ if and only if there exist linearly independent sets $\{f_i, i = 1, \dots, r\} \subset H$ and $\{g_i, i = 1, \dots, r\} \subset F$ such that

$$A = \sum_{i=1}^r \langle f_i, \cdot \rangle_H g_i,$$

where $\langle \cdot, \cdot \rangle_H$ is the inner product in H . Furthermore, then the adjoint of A is given by

$$A_* = \sum_{i=1}^r \langle g_i, \cdot \rangle_F f_i.$$

The class of finite-rank operators is a dense subset of the class of compact operators with respect to the operator norm.

Let $\{f_i, i \in I\} \subset H$ be a complete orthonormal system. An operator $A : H \rightarrow F$ is called a *Hilbert–Schmidt operator* if

$$\|A\|_{HS}^2 = \sum_{i \in I} \|A f_i\|_F^2 < \infty, \tag{4.30}$$

and then the Hilbert–Schmidt norm $\|A\|_{HS}$ of A is independent of the orthonormal system. The space of finite-rank operators is a dense subset of the space of Hilbert–Schmidt operators with respect to the Hilbert–Schmidt norm.

The transfer operator $P^t : L^2(\mu) \rightarrow L^2(\mu)$ with

$$P^t u(y)\mu(y) = \int p(t, x, y)u(x)\mu(x) dx$$

is Hilbert–Schmidt if

$$\|P^t\|_{HS}^2 = \int_{\mathbb{X}} \int_{\mathbb{X}} p(t, x, y)^2 \frac{\mu(x)}{\mu(y)} dx dy < \infty.$$

Best rank- r approximation. Every Hilbert–Schmidt operator is compact and therefore admits an SVD as in (3.5). This allows the following best low-rank approximation statement, here stated as in Wu and Noé (2020) and Froyland *et al.* (2013).

Theorem 4.9 (optimal rank- r approximation of transfer operators). Let the transfer operator $P = P^t : H \rightarrow F$ be Hilbert–Schmidt with singular value decomposition

$$P = \sum_{i \in I} s_i \langle u_i, \cdot \rangle_H v_i, \tag{4.31}$$

with orthonormal systems $\{u_i, i \in I\} \subset H$ and $\{v_i, i \in I\} \subset F$, and singular values $\{s_i, i \in I\}$, ordered such that $s_i \geq s_{i+1}$. Then, for given $r \in \mathbb{N}$, the unique solution of the variational problem

$$\min_{\substack{A: H \rightarrow F, \\ \text{rank}(A)=r}} \|P - A\|_{HS}$$

is given by

$$P_r = \sum_{i=1}^r s_i \langle u_i, \cdot \rangle_H v_i.$$

Moreover, for sets of functions $f = \{f_1, \dots, f_r\} \subset H$ and $g = \{g_1, \dots, g_r\} \subset F$, and any positive integer $k \in \mathbb{N}$,

$$\sum_{i=1}^r s_i^k = \max_{f, g} \sum_{i=1}^r \langle P f_i, g_i \rangle_F^k, \quad \text{with } f, g \text{ such that } \langle g_i, g_j \rangle_F = \delta_{ij} = \langle f_i, f_j \rangle_H. \tag{4.32}$$

4.3.1. Variational approach to Markov processes (VAMP)

In VAMP (Mardt *et al.* 2018, Wu and Noé 2020), for given $r \in \mathbb{N}$, the VAMP- k -score of the transfer operator P^τ in $H = L^2(\mu)$ with respect to orthonormal sets $f = \{f_1, \dots, f_r\} \subset H$ and $g = \{g_1, \dots, g_r\} \subset F$ is defined as

$$\mathcal{R}_k(f, g, \tau) = \sum_{i=1}^r \langle P^\tau f_i, g_i \rangle_F^k = \sum_{i=1}^r \langle f_i, P_*^\tau g_i \rangle_\mu^k, \tag{4.33}$$

which can also be expressed as

$$\mathcal{R}_k(f, g, \tau) = \sum_{i=1}^r \mathbb{E}_\mu(f_i(X_0)g_i(X_\tau))^k = \sum_{i=1}^r C_{f_i, g_i}(\tau)^k, \tag{4.34}$$

in terms of the time correlations of the functions f_i and g_i , as introduced in (2.5).

Now let $\phi = (\phi_1, \dots, \phi_n)^\top \in H^n$ and $\psi = (\psi_1, \dots, \psi_n) \in F^n$ denote linearly independent ansatz functions. The representation of $f = \{f_i, i = 1, \dots, r\}$ and $g = \{g_i, i = 1, \dots, n\}$ as linear combinations of these ansatz functions, that is,

$$\begin{aligned} f_i &= U_i \phi, & U_i &\in \mathbb{R}^{1 \times n}, \\ g_i &= \psi V_j, & V_j &\in \mathbb{R}^{n \times 1}, \end{aligned}$$

transforms (4.33) into

$$\mathcal{R}_k(f, g, \tau) = \sum_{i=1}^r \left(\sum_{j, l=1}^n U_{i,j} C_{\phi_j, \psi_l}(\tau) V_{i,l} \right)^k = \sum_{i=1}^r (U_i C_{\phi, \psi}(\tau) V_i)^k,$$

where $C_{\phi,\psi}(\tau) \in \mathbb{R}^{n \times n}$ denotes the time-correlation matrix between the ansatz functions, with entries $C_{\phi_i,\psi_j}(\tau)$. The orthogonality constraints on f and g transform to

$$\begin{aligned} \langle f_i, f_j \rangle_H &= U_i S_\phi U_j^\top, & S_\phi &= (S_{\phi,kl}), & S_{\phi,kl} &= \langle \phi_k, \phi_l \rangle_H, \\ \langle g_i, g_j \rangle_F &= V_i^\top S_\psi V_j, & S_\psi &= (S_{\psi,kl}), & S_{\psi,kl} &= \langle \psi_k, \psi_l \rangle_F. \end{aligned}$$

Therefore the variational problem (4.32) transforms into the variational problem

$$\max_{U,V} \sum_{i=1}^r (U_i C_{\phi,\psi}(\tau) V_i)^k \quad \text{such that } U_i S_\phi U_j^\top = \delta_{ij} = V_i^\top S_\psi V_j. \quad (4.35)$$

This immediately leads to the following.

Theorem 4.10 (optimal rank- r VAMP model). Let the ansatz functions be given by (ϕ_1, \dots, ϕ_n) and (ψ_1, \dots, ψ_n) , let $C_{\phi,\psi}(\tau)$ be the associated time-correlation matrix, and let S_ϕ and S_ψ be the associated mass matrices. Then the matrix representation (4.35) of the variational problem (4.32) with respect to ansatz functions has the form $U_i = \tilde{U}_i^\top S_\phi^{-1/2}$ and $V_i = S_\psi^{-1/2} \tilde{V}_i$, with $\tilde{U}_i, \tilde{V}_i \in \mathbb{R}^{n \times 1}$ being the left and right singular vectors of the matrix

$$S_\phi^{-1/2} C_{\phi,\psi}(\tau) S_\psi^{-1/2} = \sum_{i=1}^r \tilde{s}_i \tilde{U}_i \tilde{V}_i^\top, \quad (4.36)$$

whose singular values \tilde{s}_i are the approximations of the singular values s_i of the transfer operator P from (4.31) with respect to the ansatz spaces.

Wu and Noé (2020) then used this result to construct data-based algorithms for finding the best rank- r approximation of the transfer operator; see Section 5. There we will also see that we can go beyond this via a data-based computational of the optimal ansatz functions.

The VAMP approach has proved to be quite general and powerful; in particular, the relative freedom we still have to choose Hilbert spaces H and F also allows us to extend the transfer operator approach to non-equilibrium processes and the identification of coherent sets (Koltai, Wu, Noé and Schütte 2018); see also Section 5.

5. Data-driven methods

In this section we will provide an overview of different data-driven methods for the approximation of transfer operators and their eigenvalues and eigenfunctions. Similar approaches can also be used to estimate the generators of transfer operators. All the algorithms except for the methods to approximate the generators are implemented in the *deeptime* toolbox (Hoffmann et al. 2021).

Mollenhauer, Mücke and Sullivan (2022) have considered the problem of learning a linear operator between two Hilbert spaces from empirical observations in general. Their framework allows for the elegant derivation of *dimension-free* rates for generic learning algorithms, that is, learning a transfer operator from data is not fundamentally more difficult than learning a matrix or tensor.

5.1. Data-driven methods for time-homogeneous systems

Let us consider a discrete dynamical system $S: \mathbb{X} \rightarrow \mathbb{X}$, where $\mathbb{X} \subset \mathbb{R}^{\mathcal{N}}$, so that $x \in \mathbb{X}$ evolves as $\{x, S(x), S^2(x), \dots\}$. In our setting, S is typically given by the flow map Φ^τ for a fixed lag time $\tau > 0$. That is, we have trajectories of the form $\{x(0), x(\tau), x(2\tau), \dots\}$. In what follows, we assume that the process is time-homogeneous, i.e. the transition probabilities from time t_1 to time t_2 depend only on the difference $t_2 - t_1$, and that we have training data

$$\{(x^{(l)}, y^{(l)})\}_{l=1}^m,$$

with $y^{(l)} = S(x^{(l)})$. Such a training data set can, for instance, be extracted from one long trajectory $\{x^{(1)}, x^{(2)}, x^{(3)}, \dots, x^{(m+1)}\}$, where $y^{(l)} = x^{(l+1)}$. Another way to generate training data is to sample m initial conditions $x^{(l)}$ from a given distribution, e.g. the uniform distribution or the invariant density, and to apply the dynamical system S to each initial condition to obtain the corresponding $y^{(l)}$. Note that the estimated operators generally depend on the underlying density of the training data.

5.1.1. Extended dynamic mode decomposition

One of the simplest and most popular methods to approximate transfer operators from data is *extended dynamic mode decomposition* (EDMD) (Williams *et al.* 2015a). Although originally developed for the Koopman operator, it can also be used to estimate the Perron–Frobenius operator or the Perron–Frobenius operator reweighted with respect to the invariant density (Klus, Koltai and Schütte 2016). EDMD is a nonlinear extension of dynamic mode decomposition (DMD) (Schmid 2010), which was initially intended for the analysis of fluid dynamics problems, but has in the meantime been used for a host of other applications, including molecular dynamics, stock market data, traffic data, EEG data and quantum systems (Klus *et al.* 2018a, Hua *et al.* 2017, Marrouch, Slawinska, Giannakis and Read 2020, Avila and Mezić 2020, Klus, Nüske and Peitz 2022), to name just a few. The main difference between DMD and EDMD is that EDMD maps the data to a higher-dimensional feature space before solving the regression problem. This will be described in more detail below.

Estimation of the Koopman operator. Given training data $\{(x^{(l)}, y^{(l)})\}_{l=1}^m$, we first define data matrices $X \in \mathbb{R}^{\mathcal{N} \times m}$ and $Y \in \mathbb{R}^{\mathcal{N} \times m}$ by

$$X = [x^{(1)}, x^{(2)}, \dots, x^{(m)}] \quad \text{and} \quad Y = [y^{(1)}, y^{(2)}, \dots, y^{(m)}].$$

Further, we have to select a set of basis functions $\{\phi_i\}_{i=1}^n$, with $\phi_i: \mathbb{X} \rightarrow \mathbb{R}$, also called a *dictionary*. The optimal choice of basis functions depends strongly on the dynamical system, for which we aim to estimate associated transfer operators. Ideally, the basis functions are chosen in such a way that eigenfunctions can be well approximated by linear combinations of basis functions. In practice, however, the properties of the eigenfunctions are not known in advance. Typical choices of basis functions include monomials up to a certain order, indicator functions, radial basis functions or combinations thereof. Note that we do not assume here that the basis functions form a partition of unity. Now let $\phi: \mathbb{X} \rightarrow \mathbb{R}^n$ be the vector-valued function

$$\phi(x) = [\phi_1(x), \phi_2(x), \dots, \phi_n(x)]^\top.$$

We then define the transformed data matrices $\Phi_x, \Phi_y \in \mathbb{R}^{n \times m}$ by

$$\Phi_x = [\phi(x^{(1)}), \phi(x^{(2)}), \dots, \phi(x^{(m)})]$$

and

$$\Phi_y = [\phi(y^{(1)}), \phi(y^{(2)}), \dots, \phi(y^{(m)})],$$

and solve the minimization problem

$$\min_{\widehat{K} \in \mathbb{R}^{n \times n}} \|\Phi_y - \widehat{K}^\top \Phi_x\|_F,$$

whose solution is given by

$$\widehat{K}^\top = \Phi_y \Phi_x^+ = (\Phi_y \Phi_x^\top) (\Phi_x \Phi_x^\top)^+ = C_{yx} C_{xx}^+,$$

where $\|\cdot\|_F$ denotes the Frobenius norm and $^+$ the pseudoinverse. Furthermore, C_{xx} and C_{yx} are the (uncentred) covariance and cross-covariance matrices associated with the transformed data matrices, that is,

$$C_{xx} = \frac{1}{m} \Phi_x \Phi_x^\top = \frac{1}{m} \sum_{l=1}^m \phi(x^{(l)}) \otimes \phi(x^{(l)}),$$

$$C_{yx} = \frac{1}{m} \Phi_y \Phi_x^\top = \frac{1}{m} \sum_{l=1}^m \phi(y^{(l)}) \otimes \phi(x^{(l)}).$$

In some applications the mean-subtracted data are used, which essentially eliminates the trivial constant eigenfunction of the Koopman operator corresponding to the eigenvalue $\Lambda = 1$. The matrix $\widehat{K} = C_{xx}^+ C_{xy}$ is a representation of the projected Koopman operator with respect to the chosen dictionary. Any function contained in the function space spanned by the dictionary can be written as

$$f(x) = c^\top \phi(x) = \sum_{i=1}^n c_i \phi_i(x),$$

where $c \in \mathbb{R}^n$ contains the coefficients for the basis expansion. By solving the above minimization problem, \widehat{K} is chosen such that $\phi(y) = (K\phi)(x) \approx \widehat{K}^\top \phi(x)$,

that is,

$$(K\phi_i)(x) \approx \sum_{j=1}^n \widehat{K}_{ji} \phi_j(x).$$

It follows that

$$\begin{aligned} (Kf)(x) &= \sum_{i=1}^n c_i (K\phi_i)(x) \approx \sum_{i=1}^n c_i \sum_{j=1}^n \widehat{K}_{ji} \phi_j(x) \\ &= \sum_{j=1}^n \left[\sum_{i=1}^n \widehat{K}_{ji} c_i \right] \phi_j(x) = \sum_{j=1}^n [\widehat{K}c]_j \phi_j(x) \\ &= [\widehat{K}c]^\top \phi(x). \end{aligned}$$

Assuming that the data are i.i.d. and sampled from the probability distribution μ , EDMD then converges to the projected Koopman operator in the infinite-data limit.

Theorem 5.1. For $m \rightarrow \infty$, EDMD converges to a Galerkin projection of the Koopman operator onto the space $\text{span}\{\phi_i\}_{i=1}^n$.

Proof. It holds that

$$\begin{aligned} [C_{yx}]_{ij} &= \frac{1}{m} \sum_{l=1}^m \phi_i(y^{(l)}) \phi_j(x^{(l)}) \xrightarrow{m \rightarrow \infty} \int (K\phi_i)(x) \phi_j(x) d\mu(x) = \langle K\phi_i, \phi_j \rangle_\mu, \\ [C_{xx}]_{ij} &= \frac{1}{m} \sum_{l=1}^m \phi_i(x^{(l)}) \phi_j(x^{(l)}) \xrightarrow{m \rightarrow \infty} \int \phi_i(x) \phi_j(x) d\mu(x) = \langle \phi_i, \phi_j \rangle_\mu, \end{aligned}$$

where $x^{(l)} \sim \mu$; see also Williams *et al.* (2015a) and Klus *et al.* (2016). \square

Note that μ does not necessarily have to be the invariant density, but can be any distribution. The matrices defined in (4.5) are normalized in a slightly different fashion, but lead to the same representation of the operator, where C_{xx} corresponds to M and C_{yx} to \mathcal{P} . In addition to the number of data points m , we can also let the number of basis functions n go to infinity. A detailed convergence analysis is outlined in Korda and Mezić (2018b). Statements regarding the convergence of C_{xx} and C_{xy} and the concentration of measure can also be found in Mollenhauer (2022).

Estimation of the Perron–Frobenius operator. The adjoint operator can also be easily computed using EDMD. In this case the matrix representation \widehat{P} of the Perron–Frobenius operator P is

$$\widehat{P} = C_{xx}^+ C_{yx}.$$

Note, however, that the estimated operator depends on the underlying density of the training data. If $x^{(l)}$ is drawn from the uniform distribution, we obtain the Perron–Frobenius operator with respect to the Lebesgue measure, and if the data

were extracted from one long equilibrated trajectory, then we would obtain an estimate of the Perron–Frobenius operator with respect to the invariant density. For a comprehensive derivation, we refer the reader to [Klus *et al.* \(2016\)](#).

Spectral decomposition. We are particularly interested in the dominant eigenvalues of transfer operators, which represent the slowest timescales, and the associated eigenfunctions, which contain information about metastable sets. Let v be an eigenvector of \widehat{K} corresponding to the eigenvalue Λ , i.e. $\widehat{K}v = \Lambda v$, and define $u(x) = v^\top \phi(x)$. Then

$$(Ku)(x) \approx (\widehat{K}v)^\top \phi(x) = \Lambda v^\top \phi(x) = \Lambda u(x).$$

That is, approximations of eigenfunctions are determined by the eigenvectors of \widehat{K} . Provided that the covariance matrix C_{xx} is non-singular, which implies that $C_{xx}^+ = C_{xx}^{-1}$, the eigenvalue problem $\widehat{K}v = C_{xx}^{-1}C_{xy}v = \Lambda v$ can be rewritten as a generalized eigenvalue problem of the form

$$C_{xy}v = \Lambda C_{xx}v$$

and eliminates the pseudoinverse computation. A necessary condition¹ for the existence of the inverse of C_{xx} is that the basis functions are linearly independent. Similarly, in order to compute eigenfunctions of the Perron–Frobenius operator, we have to solve the generalized eigenvalue problem

$$C_{yx}v = \Lambda C_{xx}v.$$

Example 5.2. In order to illustrate a typical workflow, let us consider the stochastic differential equation

$$dX_t = -\nabla V(X_t) dt + \sqrt{2\beta^{-1}} dW_t,$$

where the potential $V(x)$ is given by the Himmellblau function

$$V(x) = (x_1^2 + x_2 - 11)^2 + (x_1 + x_2^2 - 7)^2$$

and $\beta = \frac{1}{10}$. We generate $m = 10000$ uniformly distributed initial conditions $x^{(l)}$ in $\mathbb{X} = [-6, 6]^2$ and apply the Euler–Maruyama integrator with a step size $h = 10^{-3}$ to compute the corresponding vectors $y^{(l)}$ using the lag time $\tau = 10$. Selecting a basis comprising 900 Gaussians with bandwidth $\varsigma = 1.2$ centred at the midpoints of a uniform grid, we then estimate the Koopman operator and the Perron–Frobenius operator with the aid of EDMD and compute the dominant eigenvalues and associated eigenfunctions. The results are shown in [Figure 5.1](#). By clustering the state space into metastable sets, we obtain a coarse-grained model.

¹ This is not a sufficient condition since the integrals are computed via Monte Carlo integration and depend on the data.

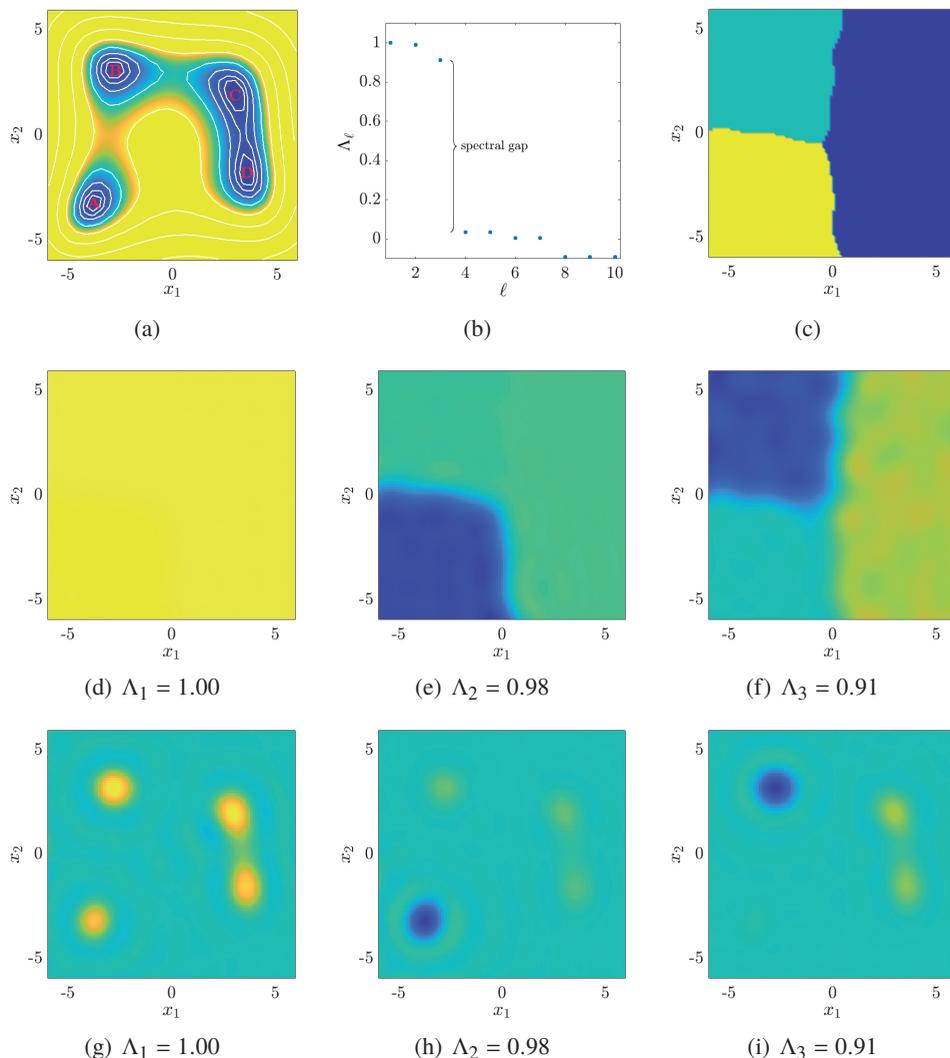


Figure 5.1. (a) Himmelblau potential, where blue corresponds to small values and yellow to large values. (b) Spectrum of the Koopman operator. There are three dominant eigenvalues close to 1, implying three metastable sets. Due to the proximity of the minima C and D and the large noise, these two wells constitute one metastable set. (c) Clustering of the dominant Koopman eigenfunctions (shown below) into three metastable sets. (d–f) Dominant eigenfunctions of the Koopman operator. The eigenfunction corresponding to $\Lambda_1 = 1$ is the constant function. The second eigenfunction separates well A from the other wells, which forms the highest energy barrier. The third eigenfunction separates B and C \cup D. (g–i) Dominant eigenfunctions of the Perron–Frobenius operator. These eigenfunctions again contain information about the locations of the metastable sets. The eigenfunction corresponding to $\Lambda_1 = 1$ is the invariant density $\mu \sim \exp(-\beta V)$.

Koopman modes and system identification. In addition to computing spectral properties of transfer operators, EDMD can also be used for prediction. Let $g: \mathbb{X} \rightarrow \mathbb{X}$ be the vector-valued observable defined by $g(x) = x$. The function g is also sometimes referred to as the *full-state observable*. Furthermore, let $B \in \mathbb{R}^{n \times \mathcal{N}}$ be the matrix such that

$$g(x) = B^\top \phi(x). \tag{5.1}$$

The easiest way to accomplish this is to add the observables $\{x_i\}_{i=1}^{\mathcal{N}}$ to the set of basis functions if these functions are not already contained in the dictionary. We define

$$V = [v_1, v_2, \dots, v_n],$$

where v_i denotes the i th eigenvector of \widehat{K} . Then

$$u(x) := [u_1(x), u_2(x), \dots, u_n(x)]^\top = V^\top \phi(x).$$

It follows that

$$g(x) = B^\top \phi(x) = \underbrace{B^\top V^{-\top}}_{:=\Xi} u(x) = \sum_{i=1}^n \xi_i u_i(x).$$

The columns of the matrix Ξ , denoted by ξ_i , are called *Koopman modes*. Applying the Koopman operator to the full-state observable yields

$$(Kg)(x) = g(S(x)) = S(x) \approx \sum_{i=1}^n \Lambda_i \xi_i u_i(x).$$

That is, the dynamical system can be written in terms of the Koopman eigenvalues, eigenfunctions and modes. For instance, we can now use the Koopman-based approximation of the potentially unknown dynamical system S for prediction and control (Korda and Mezić 2018a, Peitz and Klus 2019). For stochastic systems we will only be able to predict the expected value of the next state.

The above derivation of EDMD is based on the solution of a least-squares problem. Alternatively, given a reversible dynamical system, we can formulate a variational problem (Noé and Nüske 2013); see also Section 4. It was shown in Klus *et al.* (2018b) that the resulting eigenvalue problems are identical.

Curse of dimensionality. The above derivations show that we can estimate transfer operators and their spectra from data. In practice, however, accurate approximations of eigenfunctions often require a prohibitively large number of basis functions. Assume that we want to use n_i basis functions for each variable x_i ; then the product basis contains $n = \prod_{i=1}^d n_i$ functions. Storing the covariance and cross-covariance

matrices and solving the resulting eigenvalue problem thus often becomes infeasible for high-dimensional problems. This is a consequence of the *curse of dimensionality*. Different approaches have been developed over the last few years to mitigate this effect, for example kernel-based and tensor-based methods as well as deep-learning techniques.

5.1.2. Kernel-based extensions

There are two different scenarios: the number of data points m is larger than (or equal to) the number of basis functions n , or vice versa. Conventional EDMD requires computing (cross-) covariance matrices, which in turn requires computing outer products of the form $\phi(x^{(l)}) \otimes \phi(x^{(l)})$ and $\phi(y^{(l)}) \otimes \phi(x^{(l)})$, and solving an n -dimensional eigenvalue problem. If, on the other hand, $n > m$, which is, for instance, often the case for fluid dynamics problems, it is possible to construct a dual m -dimensional problem. The idea is to rewrite the algorithm in such a way that only inner products need to be computed, i.e. $\langle \phi(x^{(l)}), \phi(x^{(l)}) \rangle$ and $\langle \phi(y^{(l)}), \phi(x^{(l)}) \rangle$. This additionally allows us to replace the inner product in feature space by a kernel function $k: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ such that $k(x, x') = \langle \phi(x), \phi(x') \rangle$, which can be evaluated efficiently without explicitly constructing the feature space mapping ϕ by exploiting the so-called *kernel trick*. Instead of covariance and cross-covariance matrices, we then have to compute (generalized) Gram matrices, i.e. G_{xx} and G_{yx} with entries

$$[G_{xx}]_{ij} = k(x^{(i)}, x^{(j)}) \quad \text{and} \quad [G_{yx}]_{ij} = k(y^{(i)}, x^{(j)}). \quad (5.2)$$

Furthermore, kernels allow us to map the data into potentially infinite-dimensional feature spaces. An example of a kernel that spans an infinite-dimensional feature space is the standard Gaussian function

$$k(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\zeta^2}\right),$$

where the bandwidth ζ is a hyperparameter. Introducing the required reproducing kernel Hilbert spaces (Schölkopf and Smola 2001, Steinwart and Christmann 2008) in detail, however, is beyond the scope of this review. We will briefly summarize the main results. For a detailed derivation and high-dimensional molecular dynamics examples, we refer the reader to Williams *et al.* (2015b) and Klus *et al.* (2018a, 2019b).

Eigenfunctions of the Koopman operator can be approximated by solving the generalized eigenvalue problem

$$G_{yx} w = \Lambda G_{xx} w$$

and defining

$$u(x) = \sum_{i=1}^m w_i k(x^{(i)}, x).$$

Analogously, for the Perron–Frobenius operator, we solve

$$G_{xy} w = \Lambda G_{xx} w$$

and set

$$u(x) = \sum_{i=1}^m \tilde{w}_i k(x^{(i)}, x),$$

where $\tilde{w} = G_{xx}^{-1} w$. Note that computing these eigenfunctions does not require explicit evaluation of the feature map ϕ : the algorithms rely solely on kernel evaluations.

The advantage of kernel EDMD is that it also works well for high-dimensional systems. However, the number of data points m that can be taken into account is limited since the complexity of solving the typically dense generalized eigenvalue problem is $O(m^3)$. Another benefit is that we do not have to explicitly construct a set of basis functions: the feature space mapping ϕ is implicitly defined by the kernel. Choosing the right kernel for a given system – or tuning its hyperparameters such as the bandwidth – is still an open problem. In practice, the Gaussian kernel is often a good choice. A kernel-based variant of TICA that is based on the variational principle has been proposed in Schwantes and Pande (2015).

5.1.3. Tensor-based variants

EDMD maps the data to a generally higher-dimensional feature space. The size of the resulting data matrices $\Phi_x, \Phi_y \in \mathbb{R}^{n \times m}$ depends on the number of basis functions n and the number of snapshots m . The idea behind tensor-based methods is to construct a tensor product basis of simple (typically univariate) functions

$$\phi(x) = \phi_1(x) \otimes \cdots \otimes \phi_p(x) = \begin{bmatrix} \phi_{1,1}(x) \\ \vdots \\ \phi_{1,n_1}(x) \end{bmatrix} \otimes \cdots \otimes \begin{bmatrix} \phi_{p,1}(x) \\ \vdots \\ \phi_{p,n_p}(x) \end{bmatrix} \in \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_p}$$

so that $\Phi_x, \Phi_y \in \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_p \times m}$. Choosing, for example, $p = \mathcal{N}$ and $\phi_i(x) = [1, x_i, x_i^2, \dots, x_i^{n_i-1}]^\top$ results in a basis comprising all monomials of the form $\prod_{i=1}^{\mathcal{N}} x_i^{q_i}$, with $0 \leq q_i \leq n_i - 1$. Since the size of a tensor grows exponentially with the dimension \mathcal{N} , these high-dimensional arrays are decomposed into a network of lower-dimensional tensors using higher-order singular value decompositions or other types of low-rank approximations. There are many different tensor formats (Tucker 1964, Carroll and Chang 1970, Hackbusch 2014), one of the most powerful and efficient being the tensor-train format (Oseledets 2011, Gelb 2017). Tensor-based variants of the variational approach and EDMD, which exploit the low-rank structure to compute approximations of the eigenvalues and eigenfunctions of transfer operators without constructing the full tensor, have been proposed in Nüske, Schneider, Vitalini and Noé (2016) and Nüske, Gelb, Klus and Clementi (2021).

5.1.4. Deep learning approaches

Selecting a suitable dictionary often requires domain knowledge. A deep learning approach that automatically generates dictionaries is described in [Yeung, Kundu and Hodas \(2019\)](#). The basis functions are defined to be the output of a deep neural network and the dictionary and the Koopman operator are thus learned simultaneously. Since the basis functions are then tailored to the given system, it is possible to reduce the size of the dictionary significantly. A similar method to learn suitable observables using autoencoders can be found in [Alford-Lago, Curtis, Ihler and Issan \(2022\)](#).

5.2. Data-driven methods for time-inhomogeneous systems

So far we have assumed that the system is time-homogeneous. However, in general this is not necessarily the case. The energy potential and hence also the metastable sets could, for instance, change over time or the system might depend on external forcing terms. For such problems, the notion of metastability needs to be generalized. Instead of computing *metastable sets* we then compute *coherent sets*, which can be regarded as time-dependent metastable sets and play an important role in the analysis of transport and mixing processes in fluid dynamics. The computation of such coherent sets is based on a generalized operator corresponding to the forward–backward dynamics of the system ([Froyland 2013](#), [Banisch and Koltai 2017](#), [Froyland and Junge 2018](#)). Various other approaches for detecting coherent sets have been developed in the past; an overview of Lagrangian-based methods is presented in [Allshouse and Peacock \(2015\)](#).

5.2.1. Canonical correlation analysis

Let $\mu(x)$ be a reference density (not necessarily the invariant density, which might not exist) at time $t = 0$ and let

$$\nu(x) = \int p_\tau(x, y) \mu(x) dx \quad (5.3)$$

be the density mapped forward by the dynamics, i.e. the image density at time $t = \tau$. The *forward–backward operator* F is then defined by

$$Ff(x) = \int p^\tau(x, y) \frac{1}{\nu(y)} \int p^\tau(z, y) f(z) \mu(z) dz dy \quad (5.4)$$

and can be regarded as the composition of the Koopman operator and a reweighted Perron–Frobenius operator; see [Banisch and Koltai \(2017\)](#) for a detailed derivation. We are again interested in spectral properties of this operator. By applying clustering techniques to the dominant eigenfunctions, we obtain sets that are almost invariant under the forward–backward dynamics. These sets are now time-dependent. Initial conditions starting in such a coherent set will typically stay close together over long timescales. [Koltai et al. \(2018\)](#), [Klus, Husic, Mollenhauer and Noé \(2019a\)](#) and [Wu and Noé \(2020\)](#) have shown that the forward–backward

operator can also be approximated using covariance and cross-covariance matrices, and that identifying dominant eigenvalues and eigenfunctions is related to applying nonlinear variants of *canonical correlation analysis* (CCA) (Hotelling 1936, Melzer, Reiter and Bischof 2001) to Lagrangian data.

Estimation of the forward–backward operator. Given two multi-dimensional random variables X and Y , CCA identifies basis vectors such that the correlation between the projections of X and Y onto these basis vectors is maximized. We apply CCA to trajectory data, where Y is a time-lagged version of X , but in order to find nonlinear transformations, we first map the data to a typically higher-dimensional feature space. The derivation of CCA requires the data to have zero mean, that is, we have to centre the data in feature space. We will show that CCA is related to eigenfunctions of the forward–backward operator. We again use training data of the form $\{(x^{(l)}, y^{(l)})\}_{l=1}^m$ and select a set of basis functions $\{\phi_i\}_{i=1}^n$, with $\phi_i: \mathbb{X} \rightarrow \mathbb{R}$. The goal is to find two nonlinear mappings $f(x) = c_x^\top \phi(x)$ and $g(y) = c_y^\top \phi(y)$ such that the correlation

$$\rho = \frac{\mathbb{E}[f(x)g(y)]}{\sqrt{\mathbb{E}[f(x)^2]}\sqrt{\mathbb{E}[g(y)^2]}}$$

is maximized. We have

$$\mathbb{E}[f(x)g(y)] = \frac{1}{m} \sum_{l=1}^m f(x^{(l)})g(y^{(l)}) = \frac{1}{m} \sum_{l=1}^m c_x^\top \phi(x^{(l)})c_y^\top \phi(y^{(l)}) = c_x^\top C_{xy} c_y$$

and, analogously, $\mathbb{E}[f(x)^2] = c_x^\top C_{xx} c_x$ and $\mathbb{E}[g(y)^2] = c_y^\top C_{yy} c_y$. That is, we want to maximize

$$\rho = \frac{c_x^\top C_{xy} c_y}{\sqrt{c_x^\top C_{xx} c_x} \sqrt{c_y^\top C_{yy} c_y}}.$$

Since the vectors c_x and c_y are only determined up to multiplicative constants (see Shawe-Taylor and Cristianini 2004), we write this as a constrained optimization problem:

$$\max_{c_x, c_y} c_x^\top C_{xy} c_y \quad \text{subject to} \quad \begin{cases} c_x^\top C_{xx} c_x = 1, \\ c_y^\top C_{yy} c_y = 1. \end{cases}$$

It can be shown – using Lagrange multipliers – that the solution is given by the eigenvectors corresponding to the largest eigenvalue of the problem

$$\begin{cases} C_{xy} c_y = \rho C_{xx} c_x, \\ C_{yx} c_x = \rho C_{yy} c_y. \end{cases} \tag{5.5}$$

Solving the second equation for c_y and inserting the result into the first equation yields

$$C_{xx}^+ C_{xy} C_{yy}^+ C_{yx} c_x = \rho^2 c_x, \tag{5.6}$$

where we replaced the inverses, which might not exist, with pseudoinverses. Alternatively, Tikhonov regularization could be used. The obtained expression closely resembles the matrix representations of the transfer operators derived in Section 5.1.1, where $C_{xx}^+ C_{xy}$ is the Koopman operator and $C_{yy}^+ C_{yx}$ is a reweighted Perron–Frobenius operator.

Theorem 5.3. For $m \rightarrow \infty$, CCA converges to an approximation of the forward–backward operator (5.4) in the space spanned by the basis functions.

Proof. The first part of (5.6), $C_{xx}^+ C_{xy}$, converges to the projected Koopman operator as shown above. Furthermore, it holds that

$$[C_{xy}]_{ij} = \frac{1}{m} \sum_{l=1}^m \phi_i(x^{(l)}) \phi_j(y^{(l)}) \xrightarrow{m \rightarrow \infty} \int \phi_i(x) (K \phi_j)(x) d\mu(x) = \langle \phi_i, K \phi_j \rangle_\mu,$$

$$[C_{yy}]_{ij} = \frac{1}{m} \sum_{l=1}^m \phi_i(y^{(l)}) \phi_j(y^{(l)}) \xrightarrow{m \rightarrow \infty} \int \phi_i(y) \phi_j(y) d\nu(y) = \langle \phi_i, \phi_j \rangle_\nu,$$

and $\langle \phi_i, K \phi_j \rangle_\mu = \langle A \phi_i, \phi_j \rangle_\nu$, where

$$A f(y) = \frac{1}{\nu(y)} \int p^\tau(x, y) f(x) \mu(x) dx.$$

The composition of these two operators approximates the forward–backward operator; see Banisch and Koltai (2017) and Klus *et al.* (2019a) for a detailed derivation. \square

Spectral decomposition. The dominant eigenfunctions of the forward–backward operator encode information about the coherent sets. Approximations of these eigenfunctions can be obtained by computing the dominant eigenvalues and associated eigenvectors of the matrix

$$\widehat{F} = C_{xx}^+ C_{xy} C_{yy}^+ C_{yx},$$

i.e. $\widehat{F} v = \Lambda v$, and defining $u(x) = v^\top \phi(x)$.

Example 5.4. We construct a time-dependent potential $V(x, t)$ by modifying the Himmelblau potential introduced in Example 5.2 and define

$$V(x, t) = (y_1^2 + y_2 - 11 + 2 \sin(3s))^2 + (y_1 + y_2^2 - 7 + 2 \sin(3s))^2,$$

where $y = R(s)x$ and $R(s)$ is the matrix that rotates vectors in clockwise direction by the angle $s = (2\pi/20)t$. The distances between the wells now change in time and the wells are also rotating around the origin; see Figure 5.2(a–c). Consequently, the resulting stochastic differential equation and the associated transfer operators explicitly depend on the time t and we would not be able to detect metastable sets using standard EDMD. As in Example 5.2, we define $\mathbb{X} = [-6, 6]^2$ and choose the lag time $\tau = 10$ and a basis comprising 900 Gaussians with bandwidth

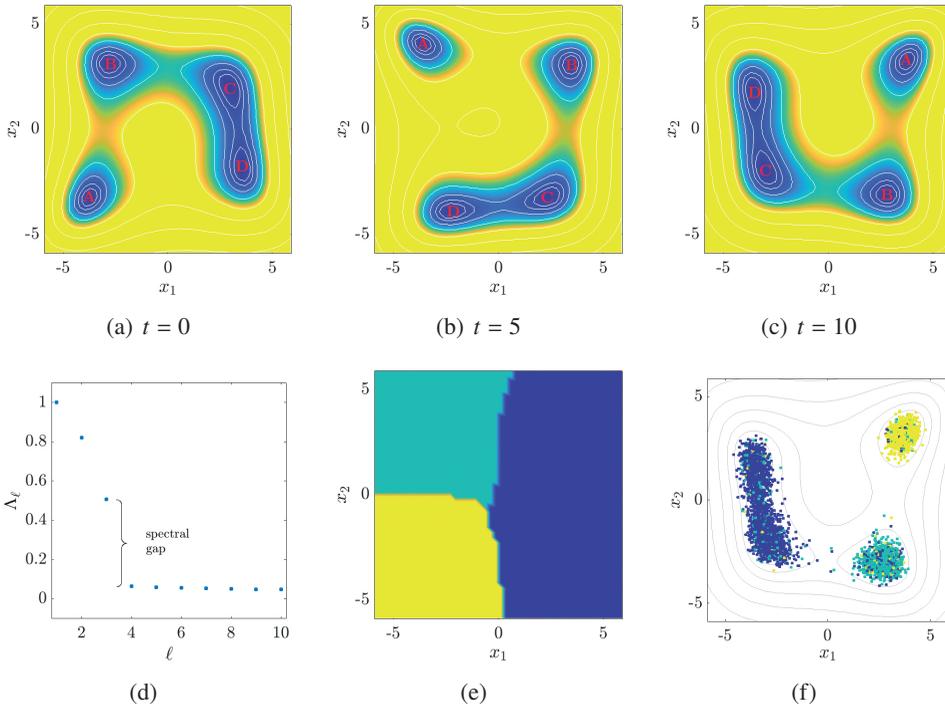


Figure 5.2. (a–c) Himmelblau potential at different times t . (d) Spectrum of the forward–backward operator. There are three dominant eigenvalues close to 1, implying the existence of three coherent sets. The corresponding eigenfunctions are similar to the eigenfunctions of the Koopman operator shown in Figure 5.1. (e) Clustering of the dominant eigenfunctions into three coherent sets at the initial time $t = 0$. (f) End points of the trajectories at time $t = 10$ coloured according to the obtained clustering. Only a few trajectories escape the coherent sets, i.e. there is little mixing.

$\zeta = 1.2$. We then generate $m = 5000$ initial conditions $x^{(l)}$ sampled from the uniform distribution and again use the Euler–Maruyama method to compute $y^{(l)}$. By applying CCA, we obtain three coherent sets corresponding to the wells A, B and $C \cup D$. Trajectories starting in these coherent sets are typically trapped within the sets for a long time. The numerical results are shown in Figure 5.2(d–f).

Curse of dimensionality. Just like EDMD, standard CCA (in feature space) requires a set of basis functions. Approximating eigenfunctions associated with high-dimensional systems might thus be infeasible. There are again many different variants and extensions of CCA using reproducing kernels, tensors or deep learning techniques.

5.2.2. Kernel-based extensions

The kernel trick can again be used to replace the covariance and cross-covariance matrices in (5.6) by Gram matrices, resulting in kernel CCA (Melzer *et al.* 2001, Shawe-Taylor and Cristianini 2004). An interpretation of kernel CCA applied to Lagrangian data in terms of transfer operators is given in Klus *et al.* (2019a). In order to compute eigenfunctions of the forward–backward operator, we have to solve the eigenvalue problem

$$G_{xx}^+ G_{yy}^+ G_{yy} G_{xx} w = \Lambda w$$

and set

$$u(x) = \sum_{i=1}^m w_i k(x^{(i)}, x).$$

The eigenfunctions corresponding to the largest eigenvalues encode information about the coherent sets. To extract these sets, we apply clustering techniques such as k -means to the dominant eigenfunctions. Note that again only kernel evaluations are required: the feature map ϕ itself is not needed. There exist many different variants; see Klus *et al.* (2019a) for a detailed derivation and examples. A kernel embedding-based variational approach called KVAD is described in Tian and Wu (2021).

5.2.3. Tensor-based variants

A tensor-based formulation of CCA can be found in Nüske *et al.* (2021), along with molecular dynamics and fluid dynamics examples. The idea is again to mitigate the curse of dimensionality by using low-rank tensor approximations. In addition to the multilinear singular value decomposition, an alternative tensor decomposition based on a higher-order CUR decomposition (Oseledets and Tyrtysnikov 2010) is presented.

5.2.4. Deep learning approaches

VAMP (Wu and Noé 2020) is formulated as an optimization problem. The goal is to find nonlinear functions that maximize the so-called VAMP- r score, which can be regarded as a generalization of the CCA cost function. A deep-learning formulation of the VAMP approach, called VAMPnets, is described in Mardt *et al.* (2018). The training data points $x^{(l)}$ and $y^{(l)}$ are fed into two deep networks, whose outputs are then merged and used to compute the VAMP- r score. The main advantage is that the neural network not only optimizes the loss function but also automatically selects basis functions that encode the slow dynamics of the system. Deep CCA was also derived in Andrew, Arora, Bilmes and Livescu (2013). A recent survey of machine learning techniques for enhanced sampling and discovering collective variables in molecular dynamics can be found in Sidky, Chen and Ferguson (2020).

5.3. Data-driven approximation of generators

In the previous sections we presented data-driven techniques to learn the Koopman operator or the Perron–Frobenius operator. Our goal now is to estimate their generators, which can again be posed as a regression problem.

5.3.1. Generator EDMD

Conventional EDMD requires time-series data and a dictionary containing basis functions to estimate associated transfer operators. A method to estimate the Koopman generator and its adjoint from data, called *generator EDMD* or *gEDMD*, was proposed in Klus *et al.* (2020b).

Estimation of the Koopman generator. Given a stochastic differential equation of the form (2.28), assume we have training data $\{(x^{(l)}, b(x^{(l)}), \sigma(x^{(l)}))\}_{l=1}^m$, where $x^{(l)}$ is again sampled from an arbitrary distribution μ , and $b(x^{(l)})$ and $\sigma(x^{(l)})$ are the drift and diffusion terms evaluated at the training data points. That is, in order to generate the training data, we need a black-box model that allows us either to compute the drift and diffusion terms directly or to generate trajectory data from which we can then estimate these terms using Kramers–Moyal formulae, that is,

$$b(x) = \lim_{t \rightarrow 0} \mathbb{E} \left[\frac{1}{t} (X_t - x) \mid X_0 = x \right],$$

$$a(x) = \lim_{t \rightarrow 0} \mathbb{E} \left[\frac{1}{t} (X_t - x)(X_t - x)^\top \mid X_0 = x \right].$$

Further, gEDMD requires a set of basis functions $\{\phi_i(x)\}_{i=1}^n$ and, since the Koopman generator is a second-order differential operator, their first and second derivatives.² These derivatives can be computed either analytically or via automatic differentiation capabilities. In what follows, we will denote the generator of the Perron–Frobenius operator (2.29) by L_P and the generator of the Koopman operator (2.30) by L_K . We again start with the Koopman case and introduce the notation

$$\bar{d}\phi_k(x) = (L_K \phi_k)(x) = \sum_{i=1}^N b_i(x) \frac{\partial \phi_k}{\partial x_i}(x) + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N a_{ij}(x) \frac{\partial^2 \phi_k}{\partial x_i \partial x_j}(x)$$

and

$$\bar{d}\phi(x) = [\bar{d}\phi_1(x), \bar{d}\phi_2(x), \dots, \bar{d}\phi_n(x)]^\top.$$

That is, $\bar{d}\phi_k(x)$ represents the Koopman generator applied to the basis function ϕ_k evaluated in a data point x . Our goal is to derive a regression problem that allows us to obtain a *global* matrix representation of the Koopman generator based on

² Provided that the system is reversible, only first-order derivatives of the basis functions are required; see Klus *et al.* (2020b) for details.

training data. To this end, we construct a new data matrix $\bar{\mathbf{d}}\Phi_x \in \mathbb{R}^{n \times m}$, defined by

$$\bar{\mathbf{d}}\Phi_x = [\bar{\mathbf{d}}\phi(x^{(1)}), \bar{\mathbf{d}}\phi(x^{(2)}), \dots, \bar{\mathbf{d}}\phi(x^{(m)})],$$

and solve the minimization problem

$$\min_{\widehat{L}_K \in \mathbb{R}^{n \times n}} \|\bar{\mathbf{d}}\Phi_x - \widehat{L}_K^\top \Phi_x\|_F,$$

whose solution is

$$\widehat{L}_K^\top = \bar{\mathbf{d}}\Phi_x \Phi_x^+ = (\bar{\mathbf{d}}\Phi_x \Phi_x^\top) (\Phi_x \Phi_x^\top)^+ = C_{\dot{x}x}^+ C_{xx}.$$

That is, $\widehat{L}_K = C_{xx}^+ C_{\dot{x}x}$ is the matrix representation of the Koopman generator L_K and we have $\bar{\mathbf{d}}\phi(x) = (L_K \phi)(x) \approx \widehat{L}_K^\top \phi(x)$, where the generator is applied componentwise to the basis functions.

Theorem 5.5. For $m \rightarrow \infty$, gEDMD converges to a Galerkin projection of the Koopman generator onto the space span $\{\phi_j\}_{j=1}^n$.

Proof. The proof is equivalent to the counterpart for standard EDMD. We again assume that the data are i.i.d. sampled from μ . In the infinite-data limit, it holds that

$$[C_{\dot{x}x}]_{ij} = \frac{1}{m} \sum_{l=1}^m \bar{\mathbf{d}}\phi_i(x^{(l)}) \phi_j(x^{(l)}) \xrightarrow{m \rightarrow \infty} \int (L_K \phi_i)(x) \phi_j(x) \, \mathrm{d}\mu(x) = \langle L_K \phi_i, \phi_j \rangle_\mu,$$

$$[C_{xx}]_{ij} = \frac{1}{m} \sum_{l=1}^m \phi_i(x^{(l)}) \phi_j(x^{(l)}) \xrightarrow{m \rightarrow \infty} \int \phi_i(x) \phi_j(x) \, \mathrm{d}\mu(x) = \langle \phi_i, \phi_j \rangle_\mu,$$

where $x^{(l)} \sim \mu$; see [Klus et al. \(2020b\)](#) for details. \square

Estimation of the Perron–Frobenius generator. Using the duality between the Koopman generator and the Perron–Frobenius generator, we obtain a matrix representation \widehat{L}_P of the adjoint operator L_P , that is,

$$\widehat{L}_P = C_{xx}^+ C_{\dot{x}x}.$$

Additional details and examples can be found in [Klus et al. \(2020b\)](#).

Spectral decomposition. In order to approximate eigenfunctions of the Koopman generator, we thus have to solve the eigenvalue problem $\widehat{L}_K v = \lambda v$. The derivation is analogous to the EDMD counterpart. We then obtain

$$(L_K u)(x) \approx (\widehat{L}_K v)^\top \phi(x) = \lambda v^\top \phi(x) = \lambda u(x).$$

If the covariance matrix is invertible, the generalized eigenvalue problem

$$C_{\dot{x}x} v = \lambda C_{xx} v$$

can be solved instead. Equivalently, for the Perron–Frobenius generator, we have

$$C_{\dot{x}x} v = \lambda C_{xx} v.$$

The eigenfunctions are then again given by $u(x) = v^\top \phi(x)$.

Koopman modes and system identification. Using standard EDMD, we were able to learn the underlying discrete dynamical system, which in our case is typically the flow map associated with the stochastic system for a fixed lag time τ . With the aid of gEDMD, on the other hand, we can identify the governing equations of dynamical systems described by ordinary or stochastic differential equations. This again requires the full-state observable $g(x) = x$, written in terms of the basis functions ϕ ; see (5.1). We define

$$V = [v_1, v_2, \dots, v_n]$$

to be the matrix containing all the eigenvectors of \widehat{L}_K , so that we can write the eigenfunctions as

$$u(x) := [u_1(x), u_2(x), \dots, u_n(x)]^\top = V^\top \phi(x).$$

It follows that

$$g(x) = B^\top \phi(x) = \underbrace{B^\top V^{-\top}}_{:=\Xi} u(x) = \sum_{i=1}^n \xi_i u_i(x),$$

where the matrix Ξ contains the *Koopman modes*. Applying the Koopman generator to the full-state observable yields

$$(L_K g)(x) = b(x) \approx \sum_{i=1}^n \lambda_i \xi_i u_i(x).$$

That is, we can express the drift term using the eigenvalues, eigenfunctions and modes of the Koopman generator and decompose it into different timescales. Alternatively, we can directly write

$$(L_K g)(x) = b(x) \approx (\widehat{L}_K B)^\top \phi(x).$$

In order to estimate the diffusion term as well, note that applying the Koopman generator to a function of the form $\phi_k(x) = x_i x_j$, which we assume is also contained in the dictionary, yields

$$(L_K \phi_k)(x) = b_i(x)x_j + b_j(x)x_i + a_{ij}(x). \tag{5.7}$$

Provided that the functions $b_i(x)$, estimated in the previous step, can be accurately approximated and also that the functions $b_i(x)x_j$ and $b_j(x)x_i$ can be expressed by the set of basis functions, this allows us to identify $a_{ij}(x)$ and hence the diffusion matrix. Applying a Cholesky decomposition to the matrix $a(x)$ yields an estimate of $\sigma(x)$.

Example 5.6. Let us again consider the Himmelblau system introduced in Example 5.2. We could also use gEDMD to compute the eigenvalues and eigenfunctions of the Koopman generator, but our goal now is to learn the governing

equations. We sample $m = 100$ initial conditions from the uniform distribution, choose a dictionary comprising monomials of order up to four, and then apply gEDMD. Here we use the exact values for $b(x^{(l)})$ and $\sigma(x^{(l)})$, but we could also estimate these terms using the aforementioned Kramers–Moyal formulae. The first six columns of the matrix representation of the generator with respect to this basis are given by

$$\begin{array}{l}
 1 \\
 x_1 \\
 x_2 \\
 x_1^2 \\
 x_1 x_2 \\
 x_2^2 \\
 x_1^3 \\
 x_1^2 x_2 \\
 x_1 x_2^2 \\
 x_2^3 \\
 x_1^4 \\
 x_1^3 x_2 \\
 x_1^2 x_2^2 \\
 x_1 x_2^3 \\
 x_2^4
 \end{array}
 \begin{bmatrix}
 1 & x_1 & x_2 & x_1^2 & x_1 x_2 & x_2^2 \\
 0 & 14 & 22 & 20 & 0 & 20 \\
 0 & 42 & 0 & 28 & 22 & 0 \\
 0 & 0 & 26 & 0 & 14 & 44 \\
 0 & 0 & -2 & 84 & 0 & 0 \\
 0 & -4 & -4 & 0 & 68 & 0 \\
 0 & -2 & 0 & 0 & 0 & 52 \\
 0 & -4 & 0 & 0 & -2 & 0 \\
 0 & 0 & 0 & -8 & -4 & -4 \\
 0 & 0 & 0 & -4 & -4 & -8 \\
 0 & 0 & -4 & 0 & -2 & 0 \\
 0 & 0 & 0 & -8 & 0 & 0 \\
 0 & 0 & 0 & 0 & -4 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & -4 & 0 \\
 0 & 0 & 0 & 0 & 0 & -8
 \end{bmatrix}.$$

Using the second and third columns we extract the drift term

$$b(x) = \begin{bmatrix} 14 + 42x_1 - 4x_1 x_2 - 2x_2^2 - 4x_1^3 \\ 22 + 26x_2 - 2x_1^2 - 4x_1 x_2 - 4x_2^3 \end{bmatrix},$$

and using the third, fourth and fifth columns and exploiting (5.7) we extract the diffusion matrix

$$a(x) = \begin{bmatrix} 20 & 0 \\ 0 & 20 \end{bmatrix}.$$

The obtained function $b(x)$ is indeed the negative gradient of the potential $V(x)$, and since $\sigma(x) = \sqrt{2\beta^{-1}I}$ it follows that $a(x) = \sigma(x)\sigma(x)^T = 20I$. That is, we have successfully identified the stochastic differential equation.

This approach also works for non-isotropic diffusion terms; see [Klus et al. \(2020b\)](#). The example illustrates the close relationships between gEDMD and system identification. The Koopman generator can be used in the same way to identify the governing equations of deterministic systems described by ordinary differential equations. We then obtain, as a special case, the *sparse identification of nonlinear dynamics* (SINDy) approach proposed in [Brunton, Proctor and Kutz \(2016\)](#).

Example 5.7. In Example 5.6 we identified the drift and diffusion terms of the stochastic differential equation. Further, since $b(x) = -\nabla V(x)$, we obtain

$$\begin{aligned} V(x) &= - \int b_1(x) dx_1 + c_1(x_2) \\ &= -14x_1 - 21x_1^2 + 2x_1^2x_2 + 2x_1x_2^2 + x_1^4 + c_1(x_2), \\ V(x) &= - \int b_2(x) dx_2 + c_2(x_1) \\ &= -22x_2 - 13x_2^2 + 2x_1^2x_2 + 2x_1x_2^2 + x_2^4 + c_2(x_1), \end{aligned}$$

where $c_1(x_2)$ and $c_2(x_1)$ are unknown functions. However, by combining these two representations, it is possible to reconstruct the potential $V(x)$ up to additive constants, that is,

$$c_1(x_2) = -22x_2 - 13x_2^2 + x_2^4.$$

The antiderivatives of the basis functions, required for computing the above integrals, can again be computed analytically for suitable dictionaries. In order for this approach to work, we have to assume that the potential can be written in terms of the basis functions. The Himmelblau potential comprises monomials of order up to four and can thus be represented as a linear combination of the basis functions. While this will generally not be the case, it is still possible to obtain approximations of the potential.

The advantage of gEDMD is that we obtain interpretable equations for the drift and diffusion terms as well as the eigenfunctions and the potential. If, on the other hand, we are only interested in approximating the potential from data, it would also be possible to use kernel density estimation to compute the invariant density μ and to define $V = -\beta^{-1} \log(\mu)$ (assuming the process is reversible) or to learn a neural network representation.

Curse of dimensionality. For high-dimensional systems we again have to face the curse of dimensionality, but we can use the same tricks introduced in the EDMD and CCA sections. This will be described in more detail below.

5.3.2. Kernel-based extensions

We can derive a kernel-based formulation of gEDMD, which, in addition to the kernel itself, also requires partial derivatives of the kernel (Klus, Nüske and Hamzi 2020a). The derivatives of the feature space functions are then computed implicitly. For the Koopman generator, this leads to an eigenvalue problem of the form

$$G_{10} w = \lambda G_{00} w,$$

where

$$[G_{10}]_{ij} = \sum_{r=1}^d b_r(x^{(i)}) \frac{\partial}{\partial x_r} k(x^{(i)}, x^{(j)}) + \frac{1}{2} \sum_{r=1}^d \sum_{s=1}^d a_{rs}(x^{(i)}) \frac{\partial^2}{\partial x_r \partial x_s} k(x^{(i)}, x^{(j)})$$

and the derivatives are computed with respect to the first argument. The derivation requires derivative-reproducing properties (Zhou 2008). This approach can also be used to approximate related differential operators such as the Schrödinger operator.

5.3.3. Tensor-based variants

A tensor-based formulation of gEDMD, which can be regarded as a combination of the methods presented in Klus *et al.* (2020b) and Nüske *et al.* (2021), for the approximation of the Koopman generator is described in Lücke and Nüske (2022). The tensor-train format is again used to generate extremely high-dimensional feature spaces.

5.3.4. Deep learning approaches

We are not aware of deep learning-based approaches that directly approximate the Koopman generator using variants of EDMD, although similar methods have been developed for related differential operators (Pfau, Spencer, Matthews and Foulkes 2020, Han, Lu and Zhou 2020).

5.4. Related methods

Several different methods to compute projected transfer operators and their eigenvalues and eigenfunctions have been developed independently by the fluid dynamics, molecular dynamics and dynamical systems communities. Some of these methods are closely related or, under certain conditions, equivalent.

5.4.1. DMD and TICA

As already mentioned above, EDMD is a nonlinear variant of DMD. By choosing $\phi(x) = x$, we obtain the optimization problem

$$\min_{A \in \mathbb{R}^{d \times d}} \|Y - AX\|_F,$$

and thus DMD as a special case. Note, however, that DMD computes the eigenvectors of A (not A^\top) and thus the Koopman modes, whereas EDMD computes the Koopman eigenfunctions. *Time-lagged independent component analysis* (TICA), on the other hand, computes the Koopman eigenfunctions. (Of course, DMD can be rewritten in such a way that it also computes the Koopman eigenfunctions and TICA so that it computes the Koopman modes.) TICA was originally introduced in the signal processing literature (Molgedey and Schuster 1994) and later used in molecular dynamics as a method for extracting slow order parameters (Perez-Hernandez *et al.* 2013, Schwantes and Pande 2013). It is often used to project high-dimensional trajectory data onto the dynamically relevant slowest timescales. This dimensionality reduction step is typically required for high-dimensional problems to be able to then, for instance, derive Markov state models in the reduced space (Perez-Hernandez *et al.* 2013). TICA maximizes the auto-correlation of the transformed coordinates and can be regarded as a special case of EDMD and VAMP.

5.4.2. *Data-driven estimation of projected transfer operators*

Let us return to the projection of a transfer operator $P = P^\tau$ onto a finite-dimensional ansatz function space spanned by the ansatz functions $\{\phi_i\}_{i=1}^n$, and let Q be the projection onto the ansatz space. Then the matrix representation P_Q of the associated projected transfer operator QPQ , as defined in Section 4, has the form $P_Q = \mathcal{P}M^{-1}$ with

$$\begin{aligned} \mathcal{P}_{ij} &= \langle \phi_j, P^\tau \phi_i \rangle_\mu = \mathbb{E}_\mu(\phi_i(X_0)\phi_j(X_\tau)), \\ M_{ij} &= \langle \phi_i, \phi_j \rangle_\mu = \mathbb{E}_\mu(\phi_i(X_0)\phi_j(X_0)), \end{aligned}$$

which is again slightly different from (4.5) but leads to the same P_Q . Given a long time-series $\{x^{(l)}\}_{l=1}^{m+1}$ that results from sampling a long realization of (X_t) with lag time τ , let $\widehat{\mathcal{P}}$ and \widehat{M} denote the statistical estimates

$$\widehat{\mathcal{P}}_{ij}^{(m)} = \frac{1}{m} \sum_{l=1}^m \phi_i(x^{(l)})\phi_j(x^{(l+1)}), \quad \widehat{M}_{ij}^{(m)} = \frac{1}{m} \sum_{l=1}^m \phi_i(x^{(l)})\phi_j(x^{(l)}).$$

Note that these matrices are equivalent to the matrices C_{xy} and C_{xx} defined above, a fact that again illustrates the basic equivalence of the underlying strategies. Then $\mathcal{P}^{(m)}$ approximates \mathcal{P} in the following sense (Mollenhauer 2022).

Theorem 5.8. Let the underlying Markov process be ergodic. Then $\mathcal{P}^{(m)} \rightarrow \mathcal{P}$ for $m \rightarrow \infty$ \mathbb{P} -a.e. in Hilbert–Schmidt norm $\|\cdot\|_{HS}$. If the process is furthermore geometrically ergodic and there is a constant $B > 0$ such that $\phi_i(X_t) \leq B$ \mathbb{P} -a.e. for all $i = 1, \dots, n$, then \mathbb{P} -a.e. we have the convergence speed estimate

$$\|\mathcal{P} - \widehat{\mathcal{P}}^{(m)}\|_{HS} = O\left(\frac{(\log m)^{3/2}}{m^{1/2}}\right)$$

and the concentration of measure estimate

$$\mathbb{P}\left[\|\mathcal{P} - \widehat{\mathcal{P}}^{(m)}\|_{HS} \geq \epsilon\right] \leq 4 \exp\left(-\frac{m\epsilon^2}{K}\right) + C \exp(-cm),$$

with constants C and c that do not depend on n, m, ϵ, B and $K = 8\sigma^2 + \frac{4}{3}nB^2m\epsilon$, where σ^2 is a variance proxy that does not depend on ϵ and satisfies $\sigma^2 \leq 4n^2B^4C \exp(-cm)$.

There is an identical statement for the convergence and convergence speed for the mass matrix M and a similar one for the concentration of measure effect; see Mollenhauer (2022, Theorem 3.4.5) for details.

5.4.3. *Data-driven Markov state models*

In Section 4.2.3 we discussed how to construct the best MSM for the process (X_t) and the associated transfer operator P^τ for some given fixed lag time $\tau > 0$, based on an ansatz space spanned by indicator functions $\phi_i = \mathbf{1}_{A_i}$ defined by an initial set

partition $\{A_1, \dots, A_n\}$ of state space, and then aggregating it further. The related $n \times n$ transition matrix is

$$P_{ij} = \mathbb{P}_\mu[X_\tau \in A_j \mid X_0 \in A_i].$$

Given the long time-series $\{x^{(l)}\}_{l=1}^{m+1}$ sampled from (X_t) with lag time τ , we know from the discussion above that the estimator $\widehat{P}^{(m)}$ with entries

$$\widehat{P}_{ij}^{(m)} = \frac{\sum_{k=1}^m \mathbf{1}_{A_i}(x_k) \mathbf{1}_{A_j}(x_{k+1})}{\sum_{k=1}^m \mathbf{1}_{A_i}(x_k)}$$

converges to P for $m \rightarrow \infty$ with convergence speed given by Theorem 5.8. However, the statement in Theorem 5.8 is mostly useful as an asymptotic statement for large m . Therefore we are also interested in understanding the statistical uncertainty of quantities computed from $\widehat{P}^{(m)}$ for given m and potentially not asymptotically large m .

To this end, we define the transition counts between the partition sets,

$$C_{ij} = \sum_{l=1}^m \mathbf{1}_{A_i}(x^{(l)}) \mathbf{1}_{A_j}(x^{(l+1)}).$$

Now, under the condition that the hitting and transition counts C resulted from a realization of the Markov chain with states $\{A_1, \dots, A_n\}$ and transition matrix T with initial distribution π , then the probability of observing this time series would be

$$\mathbb{P}(C \mid T) = \pi(x^{(1)}) \prod_{i,j=1}^n T_{ij}^{C_{ij}}.$$

We are interested in the opposite question: What is the probability $\mathbb{P}(T \mid C)$ of the transition matrix T given the observed data? By virtue of Bayes' theorem it follows that the posterior probability $\mathbb{P}(T \mid C)$ is given by

$$\mathbb{P}(T \mid C) = \frac{\mathbb{P}(C \mid T) \mathbb{P}(T)}{\mathbb{P}(Y)},$$

where $\mathbb{P}(C \mid T)$ is the likelihood function, $\mathbb{P}(T)$ is the prior probability of transition matrices, and the normalization factor $\mathbb{P}(Y) = \int \mathbb{P}(C \mid T) \mathbb{P}(T) dT$ is called the evidence, where we do not distinguish between probabilities and probability densities for simplicity. Based on the above considerations, the likelihood takes the form

$$\mathbb{P}(C \mid T) = \prod_{i,j=1}^n T_{ij}^{C_{ij}},$$

whereas, if no other information is available, the prior $\mathbb{P}(T)$ is a uniform distribution on the space of $n \times n$ stochastic matrices. In this case the *optimal transition matrix*

\widehat{P} is the maximum of the posterior $\widehat{P}^{(m)} = \arg \max_T \mathbb{P}(C | T) \mathbb{P}(T)$, which can be computed explicitly by likelihood maximization with stochasticity side constraints,

$$\widehat{P}_{ij}^{(m)} = \frac{C_{ij}}{\sum_{j=1}^n C_{ij}},$$

which is just the estimator from above.

In MSM building the typical procedure is (i) to compute a long time series and the associated maximal posterior transition matrix $\widehat{T} \approx P$, and (ii) to use PCCA+ to aggregate \widehat{T} into the low-dimensional MSM P_ϕ , as described in Section 4.2.3. The MSM P_ϕ then optimally encodes the dominant eigenvalues of the transfer operator P^τ with errors resulting from discretization and the finiteness of the time series only.

The above Bayesian approach allows us to do uncertainty quantification for all the quantities that we may want to extract from the MSM, for example the dominant eigenvalues or timescales. To this end, we have to sample the posterior $\mathbb{P}(T | C)$ on the set of stochastic matrices. From the resulting distribution of transition matrices we can then extract uncertainties of the desired quantities. Several different approaches (different sampling strategies that allow for including additional side constraints) were introduced in Singhal and Pande (2005), Noé (2008), Metzner, Noé and Schütte (2009a) and Metzner, Weber and Schütte (2010). See therein for examples of uncertainty quantification regarding the dominant eigenvalues, for instance.

Related literature. There is a large collection of articles in which MSM building procedures have been applied to a wide variety of molecular systems, including protein folding (Noé *et al.* 2009), kinetic fingerprinting (Keller *et al.* 2011), spectroscopic observables (Prinz *et al.* 2011), RNA (Huang *et al.* 2010), protein–peptide association (Paul *et al.* 2017), ligand binding, rebinding and multivalency (Weber and Fackeldey 2014, Ge and Voelz 2021), non-equilibrium kinetics (Nüske *et al.* 2017) and numerical recipes (Pande *et al.* 2010), overviews of theory and algorithms (Bowman *et al.* 2014, Husic and Pande 2018, Schütte and Sarich 2014) and software (Senne *et al.* 2012, Scherer *et al.* 2015, Beauchamp *et al.* 2011, Hoffmann *et al.* 2021), to name just a few examples. A recent critical appraisal can be found in Suarez *et al.* (2021).

5.4.4. Data-driven methods for computing coherent sets

Various other methods have been proposed to detect coherent sets using space–time diffusion maps (Banisch and Koltai 2017), generalized Markov state models (Koltai *et al.* 2018) or FEM-based discretizations (Froyland and Junge 2018). The analogy between metastable sets and clusters in undirected graphs can also be exploited to develop clustering techniques for directed and time-evolving graphs (Klus and Conrad 2022). This leads to the notion of *coherent sets* in graphs.

5.4.5. Koopman lifting technique

The Koopman generator can also be computed from time-series data only, as shown in [Mauroy and Goncalves \(2016\)](#). The idea is first to estimate the Koopman operator for a fixed lag time τ using EDMD. Then an approximation of the generator can be obtained by computing the matrix logarithm, that is,

$$\widehat{L}_K = \frac{1}{\tau} \log \widehat{K}.$$

The matrix logarithm, however, is not unique, and sufficiently small lag times are required to ensure that fast timescales are not damped out. This derivative-free approach can also be used to identify the governing equations of ordinary and stochastic differential equations.

5.5. Deep learning of reaction coordinates and effective dynamics

In [Section 3.3](#) we discussed the theoretical concepts behind the notion of reaction coordinates of collective variables. As we have seen, in theory the identification of good reaction coordinates can be related to the dominant eigenfunctions and eigenvalues of the transfer operator, the existence of transition manifolds, and the committor functions between the most important metastable sets of the system under consideration. Unfortunately, it seems that the computation of these objects requires us to know the long-term dynamics of the system beforehand or approximating it while computing these objects. In contrast, we want to compute reaction coordinates and the associated effective dynamics in order to study the long-term dynamics.

Recent reviews of computational methods for finding reaction coordinates distinguish between different types of approaches (see [Bhakat 2022](#)): geometric concepts that extend methods such as principal component analysis (PCA) towards reaction coordinates, sampling-oriented ones that are tailored towards efficient sampling of the free energy landscape, and truly dynamics-oriented concepts that seek reaction coordinates in order to compute the minimally complex yet most predictive aspects of a given molecular dynamics trajectory and directly obtain associated thermodynamic and kinetic information. Here we will discuss some algorithmic concepts for computing reaction coordinates from molecular simulation data that try to combine the power of machine learning with the efficiency of algorithms such as MSMs, TICA, EDMD or VAMP.

Several recent approaches for finding reaction coordinates utilize autoencoders. An autoencoder ([Figure 5.3](#)) is a type of deep neural network that is trained in a self-supervised way. The layer structure of the network is usually symmetric with a bottleneck in the middle, and the first half including the bottleneck is called the encoder while the second half is called the decoder. The input data form some kind of molecular dynamics time series $(x_{t_i}), i = 1, \dots, n, t_i = t_0 + i\tau$, which can be all-atom coordinates, or some internal distances, or other degrees of freedom.

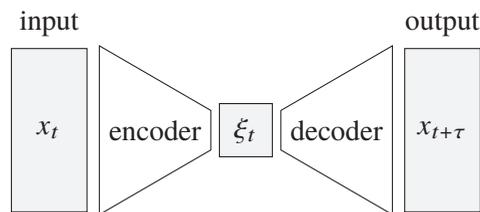


Figure 5.3. Illustration of the autoencoder network architecture. The encoder and decoder are multilayer neural networks.

The autoencoder network must learn to encode a high-dimensional vector x_t as a low-dimensional representation ξ_t to pass the information through the bottleneck and reconstruct the original signal in the decoder.

Time-lagged autoencoder. In the time-lagged autoencoder approach (Wehmeyer and Noé 2018), the idea is put forward that the autoencoder network is trained to minimize the data-based regression loss function

$$\min_{D,E} \mathcal{L}_d(D, E), \quad \mathcal{L}_d(D, E) = \sum_{i=1}^n \|x_{t_{i+1}} - D(E(x_{t_i}))\|^2.$$

Here the autoencoder is employed to learn a nonlinear map $E: \mathbb{X} \rightarrow \mathbb{R}^r$, $\xi_t = E(x_t)$, that reduces dimension in combination with a reconstruction map $D: \mathbb{R}^r \rightarrow \mathbb{X}$, $D(\xi_t) \approx x_{t+\tau}$. The low-dimensional representations ξ_t are called the *latent variables* and are understood to represent reaction coordinates.

Variational autoencoder. Alternatively, the *reweighted autoencoded variational Bayes for enhanced sampling* (RAVE) approach (Ribeiro, Bravo, Wang and Tiwary 2018) and its extensions (Wang, Ribeiro and Tiwary 2019) utilize the autoencoder idea in combination with different loss functionals. In the Bayesian framework, the prior distribution $\pi(\xi)$ on the latent variable space is connected to the joint distribution by

$$p(x, \xi) = p(x | \xi)\pi(\xi)$$

with the likelihood $p(x | \xi)$ describing the decoder. Theoretically, one would then describe the encoder by

$$p(\xi | x) = \frac{p(x | \xi)\pi(\xi)}{p(x)}, \quad p(x) = \int_{\mathbb{X}} p(x | \xi)\pi(\xi) d\xi.$$

However, sampling of $p(x)$ is computationally expensive, so the encoder is treated as being described by another conditional distribution, $q(\xi | x)$, that is approximated

by the autoencoder neural net based on data by minimizing the Kullback–Leibler (KL) divergence³ (Bhakat 2022),

$$\min_{q,p} KL(q(\xi | x) || p(\xi | x)).$$

Minimizing the KL divergence is equivalent to maximising the co-called Evidence Lower Bound (ELBO) which allows us to avoid the hard task of minimising the KL divergence between the approximate $q(\xi | x)$ and the true posterior $p(\xi | r)$. Instead we maximise ELBO,

$$\max_{q,p} (\mathbb{E}_q(\log p(x | \xi)) - KL(q(\xi | x) || \pi(\xi))),$$

where the first term describes the expectation value of the decoder log-likelihood if the latent variable is drawn from the encoder and the second term describes the Kullback–Leibler distance between the encoder and the probability distribution $\pi(\xi)$ on the latent variable space.

The idea of finding reaction coordinates as latent variables in autoencoders has recently been extended in the form of iterative schemes, for example, towards *enhanced sampling simulation* in Bonati, Piccini and Parrinello (2021) by combining it with the *on-the-fly probability-enhanced sampling method* (OPES) (Invernizzi and Parrinello 2020) or, similarly, in *free energy biasing and iterative learning with autoencoders* (FEBILAE) (Belkacemi, Gkeka, Lelièvre and Stoltz 2022), an iterative scheme in which, in each iteration, the learned reaction coordinate is used to perform free energy adaptive biasing to generate new data and learn a new reaction coordinate. While time-lagged autoencoders are a truly dynamics-oriented concept, the other approaches are more sampling-oriented.

Learning the effective dynamics. The *learning the effective dynamics* (LED) approach (Vlachas *et al.* 2022) augments the equation-free methodology (Kevrekidis and Samaey 2009) but takes account of the insight that the effective dynamics will in general be non-Markovian. It employs a probabilistic mapping between reaction coordinate (latent) state space and full-scale molecular dynamics using *mixture density network* (MDN) autoencoders, and evolves the non-Markovian latent dynamics using *long short-term memory models* (LSTMs). The basic idea of LED is as follows.

- (1) Data from short full-scale MD simulations is passed through the encoder network.
- (2) The output (ξ_t) of the encoder is used as time-series input to the LSTM, allowing for the update of its hidden state h_t that captures non-Markovian effects in the

³ Throughout this paper we use the notation $KL(\cdot || \cdot)$ and $KL(\cdot | \cdot)$ to denote the Kullback–Leibler divergence, also known as *relative entropy*, between probability measures (see Definition 6.2), or – assuming they exist – the corresponding Lebesgue densities. Here the double bar between the arguments is used in order to avoid confusion with conditional probabilities.

effective dynamics. The output of the LSTM is a parametrization of the probabilistic non-Markovian latent dynamics $p(\xi_t | h_t)$ based on Gaussian mixture models.

- (3) The LSTM iteratively samples $p(\xi_t | h_t)$ and propagates the low-order latent dynamics up to a medium timescale (much longer than the short MD trajectory used to perform step (1)).

The decoder is employed to map the latent state ξ_t at any desired t back to a high-dimensional representation by sampling $p(\cdot | \xi_t, \xi_{t-\Delta t}, \dots)$. Propagation in the low-order latent space is orders of magnitude cheaper than full-scale MD simulation on the same timescale.

Learning transition manifolds. All the approaches for the identification of reaction coordinates discussed so far show very promising performance when applied to small-to-medium-sized biomolecular systems. However, the underlying theory is rather underdeveloped: we have neither error estimates regarding the quality of the reaction coordinate or with respect to kinetic information, nor theoretical justifications for the hope that they overcome the curse of dimensionality. Here the theory-driven approach via transition manifolds (see Section 3.3.4) is unique in the sense that it permits the construction of a deep learning approach based on the variational formulation of reaction coordinates (see Theorem 3.20) via minimization of the loss functionals \mathcal{F}_D and \mathcal{F}_L given in (3.49) and (3.50). In order to see how this can be realized in a data-based setting, we first rewrite \mathcal{F}_D as

$$\mathcal{F}_D = \int_{\mathbb{X}} f(x) \mu(x) dx,$$

with

$$f(x) = \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\mathbb{L}_z(\vartheta)} \int_{\mathbb{L}_z(\vartheta)} \left| \frac{p(x, y^{(1)})}{\mu(y^{(1)})} - \frac{p(x, y^{(2)})}{\mu(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz,$$

and then utilize a Monte Carlo sampling scheme to approximate the integral

$$\mathcal{F}_D^{(M)} = \frac{1}{M} \sum_{i=1}^M f(x_i),$$

based on data points x_1, \dots, x_M that are distributed according to μ with approximation error (Bitttracher *et al.* 2021)

$$\mathbb{E}[|\mathcal{F}_D - \mathcal{F}_D^{(M)}|] \leq \frac{\text{Var}_{\mu}(f)}{\sqrt{M}}, \quad \text{with} \quad \text{Var}_{\mu}(f) = \mathbb{E}_{\mu}[f^2] - (\mathbb{E}_{\mu}[f])^2. \quad (5.8)$$

The independence of the convergence rate $1/\sqrt{M}$ of the dimension of state space is what gives MC methods an edge over conventional methods, but is only effective if the prefactor $\text{Var}_{\mu}(f)$ does not grow too much with the dimension of state space. The following theorem from Bitttracher *et al.* (2021) shows that this is not the case.

Theorem 5.9 (convergence rate does not depend on full dimension). Assume that the system is ε -lumpable with respect to $\xi: \mathbb{X} \rightarrow \mathbb{Z}$ and the effective density $p_L: \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}$. Define $f_L: \mathbb{Z} \rightarrow \mathbb{R}$ by

$$f_L(z) := \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\mathbb{L}_{z'}(\theta)} \int_{\mathbb{L}_{z'}(\theta)} \left| \frac{p_L(z, y^{(1)})}{\pi(y^{(1)})} - \frac{p_L(z, y^{(2)})}{\pi(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz'. \quad (5.9)$$

Then there exists a constant $C > 0$ that depends only on the dimension of \mathbb{Z} and not on the dimension of \mathbb{X} such that

$$|\text{Var}_\mu(f) - \text{Var}_\mu(f_L \circ \xi)| \leq \varepsilon C \|f_L \circ \xi\|_{L^1_\mu} + O(\varepsilon^2).$$

In words, the variance of f is ε -close to the variance of f_L . As f_L is defined on \mathbb{Z} , the variance of f cannot depend on the full phase space dimension through the dimension of its argument. There is also numerical evidence demonstrating that the variance is quite small. Note, however, that the dimension of \mathbb{X} also indirectly appears in the definition of f_L , via the integration over the $(n-r)$ -dimensional level sets $\mathbb{L}_z(\theta)$; rigorous results on its independence of the level set integration from $\dim(\mathbb{X})$ are still lacking.

The variational principle for reaction coordinates allows us to approximate the reaction coordinate ξ by a neural network and learn the network's parameters from data (ensembles of short MD trajectories). After parametrization of the network, the resulting reaction coordinate ξ is used to learn the (non-Markovian) effective dynamics via the generalized Langevin equation (GLE) (see Section 3.4.2), as outlined in Ayaz *et al.* (2021) and Kappler *et al.* (2018).

ISOKANN: Learning dominant almost invariant subspaces. ISOKANN (invariant subspaces of Koopman operators learned by neural networks) (Rabben, Ray and Weber 2020) tries to approximate membership functions χ as introduced in Section 3.2.4 for the main metastable sets of a system. A membership function χ has to satisfy $K^\tau \chi \approx \chi$, where $K^\tau f(x) = \mathbb{E}_x[f(X_\tau)]$ denotes the Koopman operator with lag time τ . ISOKANN utilizes a re-scaled power iteration

$$\chi_{k+1} = \frac{K^\tau \chi_k - \min(K^\tau \chi_k)}{\|K^\tau \chi_k - \min(K^\tau \chi_k)\|_\infty},$$

where the shifting by $\min(K^\tau \chi_k)$ guarantees that χ_{k+1} has values in $[0, 1]$ (as a membership function should) and convergence to the constant function $\mathbf{1}$ is avoided. Given some initial function χ_0 with values 1 within a metastable core and 0 outside, the re-scaled power iteration converges to the membership function associated with the associated soft metastable state. Repetition of this scheme for all m main metastable cores of a system then results in k membership functions χ_i , $i = 1, \dots, m$, which are expected to span the dominant invariant subspace of the Koopman operator. In ISOKANN each of these χ_i is approximated by a neural network (see Rabben *et al.* 2020 for details), and available numerical evidence indicates that the method does not suffer from the curse of dimensionality.

Furthermore, Rabben *et al.* (2020) claim that transition pathways starting in the i th metastable core can be identified by following the gradient of χ_i , and that this can be used to *learn important transitions*. This claim is justified by means of the holding probability p_{χ_i} associated with the membership function χ_i (see Section 3.2.4 for introduction and discussion of p_{χ_i}): according to Rabben *et al.* (2020), the local direction along the transition pathway is given by $r(x, t) = \nabla p_{\chi_i}(x, t)$, which is proportional to $\nabla \chi_i$. In Rabben *et al.* (2020) this claim is supported by applications to finding important transitions in MD simulations of μ -opioid receptors.

6. Rare event simulation

6.1. Statistics of rare events

In this section we address the problem of efficiently estimating quantities related to the slow MD timescales. To this end, we will first review the basic principles of rare event simulation (RESIM) and explain why the direct numerical simulation of rare events in molecular simulation (such as protein folding) is often infeasible.

6.1.1. Crude Monte Carlo estimator for rare events

To illustrate the problem, suppose we want to compute the probability $\vartheta = \mathbb{P}(X \in E)$ of some event E , such as a folding event before time T or the event of hitting $B \subset \mathbb{X}$ before $A \subset \mathbb{X}$. We assume that $\vartheta \ll 1$, and consider the Monte Carlo approximation of the parameter ϑ : given m independent realizations $X(\omega_1), \dots, X(\omega_m)$ of our stochastic dynamics, the number

$$\vartheta_m = \frac{1}{m} \sum_{i=1}^m \mathbf{1}_{\{X(\omega_i) \in E\}} \quad (6.1)$$

is an unbiased estimator of ϑ that, as $m \rightarrow \infty$, converges to ϑ with probability 1 by the law of large numbers. The variance of the estimator is nicely bounded and decreases with rate $1/m$ since

$$\text{Var}(\vartheta_m) = \frac{1}{m} \vartheta(1 - \vartheta) \leq \frac{1}{4m}.$$

Nevertheless the relative error (or relative standard deviation) is unbounded as a function of ϑ , for

$$\delta_m := \frac{\text{standard deviation of } \hat{\vartheta}_m}{\text{mean of } \hat{\vartheta}_m} \sim \frac{1}{\sqrt{m\vartheta}} \quad \text{as } \vartheta \rightarrow 0. \quad (6.2)$$

As a consequence, the number of Monte Carlo samples required to obtain an accurate estimate grows with $1/\vartheta$ when using a standard Monte Carlo sampling as in (6.1).

The problem of large relative error in RESIM is not restricted to the computation of small probabilities, but it persists also for quantities that are typically large, such as mean first passage times, that suffer from huge relative errors. A strategy to

reduce the error in RESIM is to reduce the variance of the estimator or, equivalently for unbiased estimators such as (6.1), the second moment.

To understand why the second moment of any quantity associated with a rare event is typically much larger than the first moment squared, note that by Jensen's inequality,

$$(\mathbb{E}[\hat{\vartheta}_m])^2 \leq \mathbb{E}[\hat{\vartheta}_m^2],$$

where equality is attained if and only if $\hat{\vartheta}_m$ is almost surely constant (and equal to ϑ), which is equivalent to the estimator having zero variance. It so happens that the dynamics depends on some parameter that controls the smallness of the rare event probability $\vartheta = \vartheta^\sigma$, and the corresponding estimator satisfies a large-deviation principle. For example, for small noise diffusions such as (2.17), with the squared noise coefficient σ^2 being much smaller than the typical energy barrier, ΔV , the estimator for the probability of barrier crossing (see Freidlin and Wentzell 1998, Theorem 1.2, Chapter 4) satisfies

$$\lim_{\sigma \rightarrow 0} \sigma^2 \log \mathbb{E}[\hat{\vartheta}_m] = -\alpha_1 \quad (6.3)$$

and

$$\lim_{\sigma \rightarrow 0} \sigma^2 \log \mathbb{E}[\hat{\vartheta}_m^2] = -\alpha_2, \quad (6.4)$$

where Jensen's inequality implies that

$$\alpha_2 \leq 2\alpha_1.$$

Here equality $\alpha_2 = 2\alpha_1$ can be attained under weaker conditions than zero variance, and an estimator with this property is called (*logarithmically*) *asymptotically efficient* or *log-efficient* (Asmussen, Dupuis, Rubinstein and Wang 2013). Log-efficiency means that the number of samples required to achieve a fixed relative error grows subexponentially as $\sigma \rightarrow 0$. However, the estimator may still have unbounded relative error in this case. Even worse, the typical situation is that α_2 is strictly less than $2\alpha_1$, in which case the necessary sample size m grows exponentially in the limit $\sigma \rightarrow 0$, since (6.3)–(6.4) imply

$$\delta_m = O\left(\exp\left(\frac{2\alpha_1 - \alpha_2}{\sigma^2}\right)\right). \quad (6.5)$$

6.1.2. Variance reduction methods for rate computations

There are two major classes of sampling techniques to reduce the variance in rare event simulation: splitting methods such as RESTART (Villén-Altamirano and Villén-Altamirano 1994) and adaptive multilevel splitting (Cérou and Guyader 2007), which decompose the state space but are still essentially based on the underlying probability distribution, and (force) biasing methods such as importance sampling (L'Ecuyer, Mandjes and Tuffin 2009), which enhance the rare events under consideration by perturbing the underlying probability distribution by forcing the

system and thus altering the rare events statistics; see [Juneja and Shahabuddin \(2006\)](#) for an overview. We should also mention sequential Monte Carlo ([C erou, Del Moral, Furon and Guyader 2012](#)), which combines both worlds and can be embedded into a splitting-like framework.

Popular splitting methods for rate computations in molecular dynamics include milestoning ([West, Elber and Shalloway 2007](#)), transition interface sampling ([Van Erp, Moroni and Bolhuis 2003](#)), forward flux sampling ([Allen, Valeriani and Ten Wolde 2009](#)) and adaptive multilevel splitting ([Aristoff, Leli vre, Mayne and Teo 2015](#)). An advantage of splitting methods is that they are non-invasive and relatively easy to parallelize; a disadvantage is that they typically require some prior knowledge of a low-dimensional reaction coordinate that allows us to monitor the rare event.

Biasing methods in molecular dynamics that are mostly used to compute (static) thermodynamic averages or free energy profiles include umbrella sampling ([Bartels and Karplus 1997](#)), metadynamics ([Bussi, Laio and Tiwary 2020](#)), adaptive biasing force ([Comer *et al.* 2015](#)) and conformational flooding ([Grubm uller 1995](#)), to mention just a few prominent examples. While they mostly do not rely on *a priori* knowledge of a reaction coordinate, they are invasive in that they require us to alter the molecular force field.

Methods that combine the best of both worlds (i.e. they are non-invasive and do not require vast prior knowledge) are replica-based methods, which alter the underlying probability distribution without changing the drift. Methods that belong in this class include replica exchange molecular dynamics ([Swendsen and Wang 1986](#)), also known as parallel tempering ([Earl and Deem 2005](#)), simulated tempering ([Marinari and Parisi 1992](#), [Martinsson, Lu, Leimkuhler and Vanden-Eijnden 2019](#)), infinite swapping ([Dupuis, Liu, Plattner and Doll 2012](#)) and parallel replica ([Voter 1998](#), [Bris *et al.* 2012](#)).

6.2. Adaptive importance sampling of rare events

A biasing method that allows for computing path-dependent (or *dynamical* or *kinetic*) properties, is adaptive importance sampling, which was first introduced by [Dupuis and Wang \(2004, 2007\)](#) in a rather general setting and then formulated by various authors in the context of small-noise diffusions, e.g. [Fleming \(2006\)](#), [Vanden-Eijnden and Weare \(2012\)](#) and [Dupuis, Spiliopoulos and Zhou \(2015\)](#). These methods have in common that they rely on large-deviation techniques and perform optimally in the asymptotic regime of vanishing rare event probability.

Here we describe a closely related yet non-asymptotic adaptive importance sampling approach that is based on a stochastic control framework ([Hartmann and Sch utte 2012](#), [Hartmann, Richter, Sch utte and Zhang 2017](#)). Before we discuss it, we shall briefly explain the general framework of importance sampling. The key idea of importance sampling is to do a *change of measure*. This is done by drawing the samples from another probability measure, say \mathbb{Q} , under which the variance of

the estimator is reduced and, ideally, the event is no longer rare. (We will see later that there are many cases in which these are conflicting goals.)

Suppose that X has distribution \mathbb{P} and let \mathbb{Q} be any measure that is absolutely continuous with respect to \mathbb{P} (i.e. $\mathbb{Q} \ll \mathbb{P}$), so that the likelihood ratio $\mathcal{L} = d\mathbb{Q}/d\mathbb{P}$ is well-defined. We further assume that $\mathcal{L} > 0$ on the relevant set $\{X \in E\}$. Letting

$$\mathbb{E}_{\mathbb{Q}}[Z] = \int Z d\mathbb{Q}$$

denote the expectation of some random variable Z with respect to \mathbb{Q} , we have

$$\mathbb{P}(X \in E) = \mathbb{E}[\mathbf{1}_{\{X \in E\}}] = \mathbb{E}_{\mathbb{Q}}[\mathbf{1}_{\{X \in E\}} \mathcal{L}^{-1}]. \quad (6.6)$$

The idea of importance sampling is to replace the sample mean of $\mathbf{1}_{\{X \in E\}}$ with the sample mean of the new random variable $Z = \mathbf{1}_{\{X \in E\}} \mathcal{L}^{-1}(X)$ over independent draws from the alternative distribution \mathbb{Q} . In other words, (6.1) is replaced by the *importance sampling* (IS) estimator

$$\hat{\vartheta}_m^{\text{IS}} = \frac{1}{m} \sum_{i=1}^m \mathbf{1}_{\{\tilde{X}(\omega_i) \in E\}} \mathcal{L}^{-1}(\tilde{X}(\omega_i)), \quad (6.7)$$

with independent realizations $\tilde{X}(\omega_1), \dots, \tilde{X}(\omega_m) \sim \mathbb{Q}$. It is easy to see that the new estimator is unbiased. Moreover, the choice $\mathbb{Q} = \mathbb{Q}^*$, with

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}}(X) = \frac{\mathbf{1}_{\{X \in E\}}}{\vartheta}, \quad \text{that is, } \mathbb{Q}^* = \mathbb{P}(\cdot \mid X \in E), \quad (6.8)$$

is optimal in that the resulting IS estimator has zero variance under $\mathbb{Q} = \mathbb{Q}^*$, that is,

$$\text{Var}_{\mathbb{Q}^*}(\hat{\vartheta}_m^{\text{IS}}) = 0.$$

As a consequence, we can already get the correct answer for $m = 1$ if we use the optimal proposal distribution. The fact that we can write down the optimal IS distribution explicitly is of limited practical use, however, because the quantity we want to compute, ϑ , appears as a normalization constant.

An important lesson to learn from the previous considerations is that variance reduction is goal-oriented in that it takes into account the random variable or quantity of interest. This observation essentially applies to all variance reduction techniques even though the optimal sampling strategy will depend on the chosen method.

Next we will discuss a particular form of importance sampling for path-dependent quantities in which the optimal change of measure has a particular explicit form that is infeasible from a computational perspective, but that can nevertheless be exploited to devise systematic approximations of the optimal proposal distribution within an exponential family.

6.2.1. A variational formula for the optimal proposal distribution on path space

We now consider continuous path-dependent functionals of the form

$$S(X) = \int_0^\tau f(X_s) ds + g(X_\tau)\mathbf{1}_{\{\tau < \infty\}}, \tag{6.9}$$

of $X = (X_s)_{s \geq t}$, $t \geq 0$, where $f, g: \mathbb{X} \rightarrow \mathbb{R}$ are bounded below and sufficiently smooth real-valued functions and τ is some random or deterministic stopping time.

Definition 6.1. Let $\theta > 0$. We define the free energy of S with respect to \mathbb{P} by

$$\Phi(x, t) = -\theta^{-1} \log \mathbb{E}[\exp(-\theta S(X)) \mid X_t = x]. \tag{6.10}$$

The free energy thus defined is a scaled form of the cumulant generating function (CGF) of the random variable S as a function of the initial data. Similarly, we can interpret $\exp(-\theta\Phi)$ as a moment generating function; it includes path-dependent quantities such as (2.32) or (2.34).

For example, if $f = 1$, $g = 0$ and τ is the first hitting time of some set, then $\exp(-\theta\Phi(x, t)) = \mathbb{E}[\exp(-\theta\tau) \mid X_t = x]$ is the moment generating function of the first hitting time of the process X , conditional on starting at x at time t . By the strong Markov property, Φ will be independent of t whenever the process is time-homogeneous.

One of the key properties of the free energy (6.10) is that it contains information about all moments of the path functional G (provided they exist). By Taylor-expanding Φ about $\theta = 0$, we find that

$$\Phi \approx \mathbb{E}[S] - \frac{\theta}{2} \mathbb{E}[(S - \mathbb{E}[S])^2]$$

for small θ , assuming that the remainder is negligible.

The CGF-like form of the free energy thus couples the first and second moments in a way that has implications when it comes to the question of whether we can control the first moment while reducing variance so as to allow for some control over the total computational cost. We will come back to this issue later, as it turns out to be relevant for the problems that include random stopping times.

Definition 6.2. The relative entropy or Kullback–Leibler divergence of \mathbb{Q} with respect to \mathbb{P} is defined as

$$KL(\mathbb{Q} \mid \mathbb{P}) := \begin{cases} \int \log\left(\frac{d\mathbb{Q}}{d\mathbb{P}}\right) d\mathbb{Q} & \text{if } \mathbb{Q} \ll \mathbb{P}, \log\left(\frac{d\mathbb{Q}}{d\mathbb{P}}\right) \in L^1(\mathbb{Q}), \\ +\infty & \text{otherwise.} \end{cases} \tag{6.11}$$

By Jensen’s inequality $KL(\mathbb{Q} \mid \mathbb{P}) \geq 0$, with equality if and only if $\mathbb{Q} = \mathbb{P}$ except for sets of \mathbb{Q} -measure zero. The fact that the free energy admits a variational form, known as the *Gibbs variational principle*, gives rise to a characterization of the optimal change of measure for computing free energies.

Theorem 6.3 (Gibbs variational principle I). For every measurable functional $S(X)$ that is bounded below, the free energy Φ of S satisfies

$$\Phi = \inf \left\{ \int S \, d\mathbb{Q} + \theta^{-1} KL(\mathbb{Q} | \mathbb{P}) : \mathbb{Q} \ll \mathbb{P} \right\}. \quad (6.12)$$

If $S \exp(-\theta S) \in L^1(\mathbb{P})$, then the infimum is attained at $\mathbb{Q} = \mathbb{Q}^*$ given by

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}} = \frac{\exp(-\theta S)}{\exp(-\theta \Phi)}. \quad (6.13)$$

We refrain from giving a proof of this classical result (e.g. Ellis 1985, Dai Pra, Meneghini and Runggaldier 1996), but we mention that the assertion basically follows once more from an application of Jensen's inequality: assuming that $KL(\mathbb{Q} | \mathbb{P}) < \infty$, we have

$$\begin{aligned} \Phi &= -\theta^{-1} \log \int \exp(-\theta S) \, d\mathbb{P} \\ &= -\theta^{-1} \log \int \exp \left(-\theta S - \log \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right) \right) \, d\mathbb{Q} \\ &\leq \int S + \theta^{-1} \log \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right) \, d\mathbb{Q} \\ &= \int S \, d\mathbb{Q} + \theta^{-1} KL(\mathbb{Q} | \mathbb{P}). \end{aligned}$$

Since the exponential function is *strictly* convex, Jensen's inequality states that equality is attained if and only if the random variable

$$Z = S + \theta^{-1} \log \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right)$$

is \mathbb{Q} -a.s. constant, which entails that the minimizer $\mathbb{Q} = \mathbb{Q}^*$ must satisfy

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}} \propto \exp(-\theta S),$$

The denominator in (6.13) is obtained from normalization and Definition 6.1, noting that

$$\exp(-\theta \Phi) = \int \exp(-\theta S) \, d\mathbb{P}.$$

It is easy to see that \mathbb{Q}^* enjoys the zero variance property: for $X^* \sim \mathbb{Q}^*$, with probability 1 we have

$$\mathbb{E}[\exp(-\theta S)] = \exp(-\theta S(X^*)) \frac{d\mathbb{P}}{d\mathbb{Q}^*}. \quad (6.14)$$

Limitations. In almost all situations of practical relevance, the optimal change of measure must be approximated numerically. Importance sampling can be fragile with respect to bad approximations of \mathbb{Q}^* , and therefore the computational gain

that can be achieved by importance sampling is very sensitive to the fidelity of the numerical method. To see this, recall the definition (6.2) of the relative error of the standard Monte Carlo estimator and call

$$R^2(\mathbb{Q}) := \delta_{\text{IS}}^2 = \frac{\text{Var}_{\mathbb{Q}}(\exp(-\theta S)\mathcal{L}^{-1})}{\Psi^2}, \quad \Psi = \mathbb{E}[\exp(-\theta S)]$$

the relative importance sampling error per sample point (i.e. for $m = 1$) and under a change of measure from \mathbb{P} to $\mathbb{Q} \ll \mathbb{P}$, with likelihood ratio $\mathcal{L} = d\mathbb{Q}/d\mathbb{P}$. We further assume that $\mathbb{Q}^* \ll \mathbb{Q}$. Then

$$\begin{aligned} R^2(\mathbb{Q}) &= \frac{\text{Var}_{\mathbb{Q}}(\exp(-\theta S)\mathcal{L}^{-1})}{\Psi^2} \\ &= \text{Var}_{\mathbb{Q}}\left(\frac{d\mathbb{Q}^*}{d\mathbb{Q}}\right) \\ &= \mathbb{E}_{\mathbb{Q}}\left[\left(\frac{d\mathbb{Q}^*}{d\mathbb{Q}}\right)^2 - 1\right], \end{aligned}$$

where the rightmost expression,

$$\chi^2(\mathbb{Q}^* | \mathbb{Q}) := \mathbb{E}_{\mathbb{Q}}\left[\left(\frac{d\mathbb{Q}^*}{d\mathbb{Q}}\right)^2 - 1\right], \tag{6.15}$$

is the χ^2 -divergence between \mathbb{Q}^* and \mathbb{Q} . Using the fact that the χ^2 -divergence is an upper bound for the KL divergence, we can prove the following non-asymptotic upper and lower bounds for the relative error that hold away from the large-deviation regime (see Hartmann and Richter 2021, Proposition 2.7):

$$R(\mathbb{Q}) \geq \sqrt{\exp(M_- \text{KL}(\mathbb{Q} | \mathbb{Q}^*) + \text{KL}(\mathbb{Q}^* | \mathbb{Q})) - 1}, \tag{6.16a}$$

$$R(\mathbb{Q}) \leq \sqrt{\exp(M_+ \text{KL}(\mathbb{Q} | \mathbb{Q}^*) + \text{KL}(\mathbb{Q}^* | \mathbb{Q})) - 1}, \tag{6.16b}$$

where $\mathbb{Q} \ll \mathbb{P}$ and the constants M_{\pm} are given by

$$M_- = \inf_{E \in \mathcal{E}} \frac{\mathbb{Q}^*(E)}{\mathbb{Q}(E)}, \quad M_+ = \sup_{E \in \mathcal{E}} \frac{\mathbb{Q}^*(E)}{\mathbb{Q}(E)}. \tag{6.17}$$

Here \mathcal{E} is an appropriate σ -algebra of measurable sets over which the infimum and the supremum are taken. Note that $M_- \in [0, 1]$ and $M_+ \in [1, \infty]$, which implies that there is a gap between upper and lower bound whenever $M_- \neq M_+$; if $M_- = 0$ and $M_+ = \infty$, then (6.16) states that

$$\sqrt{\exp(\text{KL}(\mathbb{Q}^* | \mathbb{Q})) - 1} \leq R(\mathbb{Q}) \leq \infty. \tag{6.18}$$

When $\mathbb{P}, \mathbb{Q}, \mathbb{Q}^*$ are probability measures with compact support on $\mathbb{R}^{\mathcal{N}}$, say, the constants can assume non-trivial values $M_- \neq 0$ and $M_+ \neq \infty$. However, to our knowledge no such results are known in the infinite-dimensional setting, for example when $\mathbb{P}, \mathbb{Q}, \mathbb{Q}^*$ are measures on the space of continuous paths.

Hartmann and Richter (2021, Section 3) have derived computable upper and lower error bounds for the relative error when uniform approximations to the optimal change of measure on path space are available.

Example 6.4 (suboptimal Gaussian change of measure). Equations (6.16) or (6.18) can be used to systematically derive bounds for a suboptimal change of measure $\mathbb{Q} \approx \mathbb{Q}^*$. These bounds then show that there is often an exponential growth of the relative error. For example, when \mathbb{P} is a Gaussian law on $\mathbb{R}^{\mathcal{N}}$ with mean 0, covariance $\Sigma = \sigma^2 I_{\mathcal{N} \times \mathcal{N}}$ and $S(X) = \alpha \cdot X$ for some $\alpha \in \mathbb{R}^{\mathcal{N}}$, then \mathbb{Q}^* is the Gaussian tilted measure, with density

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}}(x) = \exp\left(-\alpha \cdot x - \frac{\sigma^2}{2} |\alpha|^2\right),$$

that is, \mathbb{Q}^* is Gaussian with mean $\mathbf{m} = -\sigma^2 \alpha$ and covariance $\Sigma = \sigma^2 I_{n \times n}$. For simplicity, we have set $\theta = 1$. Now consider a Gaussian measure \mathbb{Q} with the same covariance but with the perturbed mean $\mathbf{m}_\epsilon = \mathbf{m} + \epsilon \mathbf{1}$. Then, using that the KL divergence between Gaussians can be explicitly computed,

$$R(\mathbb{Q}) \geq \sqrt{\exp\left(\frac{\mathcal{N} \epsilon^2}{2\sigma^2}\right) - 1},$$

showing an exponential growth of the relative error in the (squared) perturbation parameter that becomes even more severe for larger dimension \mathcal{N} or lower ‘temperature’ σ^2 .

One can think of the last example as a finite-dimensional version of an importance sampling scheme that works by changing the drift of a system, where the perturbation comes from an error in the drift term that may render the resulting importance sampling estimator useless for a high-dimensional system. (Note that this is the typical situation in MD simulations.)

The curse of dimensionality for importance sampling has been analysed in the seminal works by Li, Bengtsson and Bickel (2005) and Bengtsson, Bickel and Li (2008); see also Agapiou, Papaspiliopoulos, Sanz-Alonso and Stuart (2015) and Sanz-Alonso (2018).

6.2.2. Diffusion processes and exponential change of measure

Depending on the concrete problem under consideration, the Gibbs variational principle assumes different forms that then give rise to different computational methods to approximate the optimal proposal distribution. For example, for Markov chains or Markov jump processes, the minimization problem in (6.12) turns into a Markov decision problem (Banisch and Hartmann 2016), whereas it turns into a finite-dimensional optimization problem for static sampling tasks (Valsson and Parrinello 2014). Interestingly, for small noise diffusions the variational formula turns into a *deterministic* control problem that is associated with the Freidlin–Wentzell-type

large-deviation rate function for the respective rare event probabilities (Boué and Dupuis 1998); see also Dupuis, Katsoulakis, Pantazis and Rey-Bellet (2020) for a related variance reduction approach for sensitivity analysis of rare event probabilities.

To spell out the variational formula when the dynamics is diffusive, recall the definition (2.28) of a generic Itô stochastic differential equation:

$$dX_s = b(X_s) ds + \sigma(X_s) dW_s. \tag{6.19}$$

Here and in the following, $b: \mathbb{X} \rightarrow \mathbb{R}^N$ is a smooth vector field, $\sigma: \mathbb{X} \rightarrow \mathbb{R}^{N \times s}$ is a smooth matrix field and W is an s -dimensional Brownian motion.⁴ For simplicity we assume that the coefficients b and σ are time-homogeneous; the generalization to time-dependent problems is straightforward. Our standard example will be (2.17) with $b(x) = -\nabla V(x)$ and $\sigma(x) = \sqrt{2\beta^{-1}}$ being constant, but we stress that none of the results in this section are tied to the reversible setting.

We regard the distribution \mathbb{P} of X as a probability measure on the space $\Omega = C([0, \infty), \mathbb{X})$ of continuous trajectories in $\mathbb{X} \subset \mathbb{R}^N$ that is induced by the Brownian motion in (6.19). We call \mathbb{P} a *path space measure*. A change of measure from \mathbb{P} to $\mathbb{Q} \ll \mathbb{P}$ then corresponds to a change of drift from b to $b^u = b + \sigma u$, where $u = (u_t)_{t \geq 0}$ is some \mathbb{R}^s -valued stochastic process that is adapted to the filtration generated by the Brownian motion. Now, defining a controlled process X^u as the solution to

$$dX_s^u = b(X_s^u) ds + \sigma(X_s^u) u_s ds + \sigma(X_s^u) dW_s, \tag{6.20}$$

it turns out that the relevant candidates for a control u to solve the minimization problem in (6.12) are Markovian, so u_s can be written as a function of X_s^u and s . Let us suppose that u is adapted and satisfies the *Novikov condition* (see Krylov 2019)

$$\mathbb{E} \left[\exp \left(\frac{1}{2} \int_0^\infty |u_s|^2 ds \right) \right] < \infty. \tag{6.21}$$

We call controls u that satisfy (6.21) and for which (6.20) has a unique strong solution *admissible*, and we let \mathcal{A} denote the class of admissible controls. We define the controlled Brownian motion

$$W_t^u = W_t - \int_0^t u_s ds,$$

and an auxiliary process $(Z_s^u)_{s \geq 0}$, with

$$Z_t^u = \int_0^t u_s \cdot dW_s - \frac{1}{2} \int_0^t |u_s|^2 ds, \tag{6.22}$$

⁴ We assume throughout that drift and diffusion coefficients satisfy suitable Lipschitz and growth conditions, so that an SDE like (6.19) has a unique strong solution.

or equivalently

$$Z_t^u = \int_0^t u_s \cdot dW_s^u + \frac{1}{2} \int_0^t |u_s|^2 ds. \quad (6.23)$$

Girsanov's theorem (e.g. Øksendal 2003, Theorem 8.6.4) now states that $(W_s^u)_{0 \leq s \leq \tau}$ is a standard Brownian motion under the probability measure \mathbb{Q} with likelihood ratio

$$\mathcal{L}_\tau^u := \frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{[0, \tau]} = \exp(Z_\tau^u) \quad (6.24)$$

with respect to \mathbb{P} ; the Novikov condition (6.21) guarantees that $\mathbb{E}[\exp(Z_\tau^u)] = 1$, i.e. that \mathbb{Q} is a probability measure.⁵ Inserting (6.23)–(6.24) into the variational formula (6.12), using that W_s^u is a Brownian motion with respect to \mathbb{Q} , we obtain the SDE representation of the right-hand side of (6.12) for $u \in \mathcal{A}$:

$$\begin{aligned} & \mathbb{E}_{\mathbb{Q}}[S(X) + \theta^{-1} KL(\mathbb{Q} | \mathbb{P})] \\ &= \mathbb{E}_{\mathbb{Q}} \left[\int_0^\tau \left(f(X_s) + \frac{1}{2\theta} |u_s|^2 \right) ds + g(X_\tau) \mathbf{1}_{\{\tau < \infty\}} \right]. \end{aligned}$$

Using that the law of X under \mathbb{Q} is the same as the law of X^u under \mathbb{P} , we obtain the following; for a proof, see Boué and Dupuis (1998) or Dai Pra *et al.* (1996).

Corollary 6.5 (Gibbs variational principle II). Consider the functional S given in (6.9) for the process X given by the uncontrolled SDE (6.19) and a random stopping time τ for X . Then, under the Novikov condition (6.21), the Gibbs variational formula (6.12) for the free energy Φ of S reads

$$\Phi(x, t) = \inf_{u \in \mathcal{A}} \mathbb{E}_{x,t} \left[\int_t^{\tau^u} \left(f(X_s^u) + \frac{1}{2\theta} |u_s|^2 \right) ds + g(X_{\tau^u}^u) \mathbf{1}_{\{\tau^u < \infty\}} \right], \quad (6.25)$$

where X^u is the solution of the controlled SDE (6.20), τ^u is a stopping time for X^u and we have used the shorthand $\mathbb{E}_{x,t}[\cdot] = \mathbb{E}[\cdot | X_t^u = x]$ for the expectation over realizations of X^u starting at $X_t^u = x$.

To illustrate the previous statement, we consider two simple examples.

Example 6.6 (exit from a set). Consider the one-dimensional scaled Brownian motion

$$X_t = x + \sigma W_t \quad (6.26)$$

starting at $X_0 = x$. We let $D = (a, b)$ for some $a < b$ and define $\tau = \tau(D)$ to be the first exit time from D . We set $f = 1$ and $g = 0$ in (6.9) such that $S(X) = \tau(D)$. Let

⁵ Informally, the restriction of the likelihood ratio to $[0, \tau]$ denotes the restriction to trajectories of length τ , or more specifically, the restriction to the filtration generated by the Brownian motion up to time τ .

Φ again be the associated free energy. From (2.36)–(2.37) we conclude that (see Section 2.3.2)

$$\mathbb{E}[\tau \mid X_0 = x] = \begin{cases} \sigma^{-2}(b-x)(x-a) & a < x < b, \\ 0 & \text{else.} \end{cases} \tag{6.27}$$

Hence $\mathbb{E}_{x,0}[\tau]$ is nicely bounded on \bar{D} for every $\sigma \neq 0$ and, as a consequence, τ is almost surely finite. By the Feynman–Kac theorem (see Pham 2009, Theorem 1.3.17), it follows that the function

$$\Psi(x) = \exp(-\theta\Phi(x)) = \mathbb{E}[\exp(-\theta\tau) \mid X_0 = x], \quad x \in \mathbb{R} \tag{6.28}$$

solves the linear elliptic boundary value problem

$$\frac{\sigma^2}{2}\Psi''(x) - \theta\Psi(x) = 0, \quad \Psi(a) = \Psi(b) = 1, \tag{6.29}$$

with unique positive solution for $a \leq x \leq b$ given by

$$\Psi(x) = \exp(-\theta\Phi(x)) = \frac{e^{-\gamma x}(e^{\gamma(a+b)} + e^{2\gamma x})}{e^{\gamma a} + e^{\gamma b}}, \quad \gamma = \sqrt{\frac{2\theta}{\sigma^2}}.$$

Even though the dynamics is not metastable and $\mathbb{P}(\tau < \infty) = 1$, the mean first exit time (6.27) diverges for all $x \in D$ as $\sigma \rightarrow 0$. Therefore exiting from D (i.e. hitting the boundary of D) is a rare event when σ is small.

It is instructive to take a look at the relative error $\delta_m = \delta_m(\hat{\Psi}_m; \sigma)$ of the standard Monte Carlo estimator $\hat{\Psi}_m$ of the moment generating function $\Psi = \exp(-\theta\Phi)$. It has the property

$$\limsup_{\sigma \rightarrow 0} \log \delta_m = \infty,$$

since, for $x \in D$, the second moment diverges at a (strictly) higher exponential rate than the first moment squared as $\sigma \rightarrow 0$. As a consequence, the estimator is not log efficient (which implies that it has unbounded relative error as $\sigma \rightarrow 0$). Figure 6.1 shows the function $\Psi(\cdot; \theta = 1)$ and the relative error for the corresponding standard Monte Carlo estimator. (The divergence of the relative error for $\sigma \rightarrow 0$ can be considered as the infinite-dimensional analogue of the relative error bound in Example 6.4 that diverges exponentially as $\sigma \rightarrow 0$.) On the other hand, Corollary 6.5 implies that

$$\Phi(x) = -\theta^{-1} \log \mathbb{E}[\exp(-\theta\tau)] = \inf_{u \in \mathcal{A}} \mathbb{E} \left[\tau^u + \frac{1}{2\theta} \int_0^{\tau^u} |u_t|^2 dt \right], \tag{6.30}$$

where τ^u denotes the exit time from D , i.e. the first hitting time of either a or b , under the controlled process

$$X_t^u = x + \sigma W_t^u = x + \sigma \left(\int_0^t u_s ds + W_t \right).$$

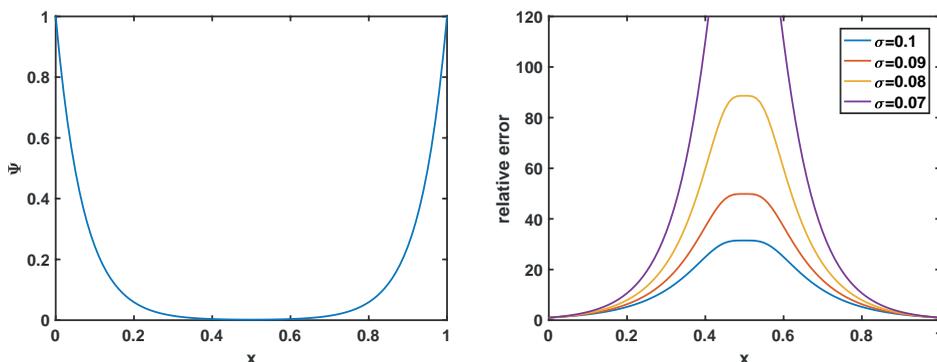


Figure 6.1. Moment generating function $\Psi(\cdot; \theta = 1)$ of the first exit time as a function of the initial condition $X_0 = x$ and relative error δ_n for the corresponding standard Monte Carlo estimator. For simplicity we have set the sample size to $m = 1$ so that the relative error shown can be regarded as the relative error per sample point.

Therefore a control u^* that minimizes the right-hand side of (6.30) will seek to minimize the mean first hitting time of the boundary ∂D , while controlling the variance of the estimator. As a consequence, the corresponding optimal change of measure from \mathbb{P} to $\mathbb{Q}^* = \mathbb{Q}^{u^*}$ has the zero variance property and the property that the rare event under consideration (i.e. hitting the boundary of D) is no longer rare. We will characterize the minimizer of the underlying stochastic optimal control in the next section.

Example 6.7 (committor probability). Under the assumptions of the previous example, we let $\rho(A)$, $\rho(B)$ be the first hitting times of the sets $A = (-\infty, a]$ and $B = [b, \infty)$, so that $\tau = \min\{\rho(A), \rho(B)\}$. Further, let

$$g(x) = \begin{cases} +\infty & \text{if } x \in A, \\ 0 & \text{else.} \end{cases}$$

Then

$$\mathbb{E}[\exp(-\theta g(X_\tau))] = \mathbb{P}(\rho(B) < \rho(A)), \quad (6.31)$$

which implies that

$$-\theta^{-1} \log \mathbb{P}(\rho(B) < \rho(A)) \leq \mathbb{E} \left[g(X_{\tau^u}) + \frac{1}{2\theta} \int_0^{\tau^u} |u_t|^2 dt \right]. \quad (6.32)$$

By inspecting the right-hand side of the last inequality, we conclude that a control $u \in \mathcal{A}$ that minimizes the right-hand side (assuming it exists) seeks to drive the process towards stopping at $x = b$ rather than $x = a$, since the process for $u = 0$ will eventually hit one of the boundary points, but hitting the leftmost boundary point has an infinite penalty. In other words, under the optimal change of measure \mathbb{Q}^* ,

all trajectories will almost surely exit from D through $\partial B \subset \partial D$, while avoiding $\partial A \subset \partial D$.

We will now discuss importance sampling of rare events in a non-asymptotic setting, specifically for diffusions with small yet non-vanishing noise. In order to turn the Gibbs principle into a workable numerical method, we will interpret the variational principle as a stochastic optimal control problem, with the unique optimal control force (or bias) generating the zero variance probability measure.

6.2.3. Stochastic optimal control

Let $O \subset \mathbb{X}$ be a bounded open set with smooth (at least C^3) boundary and let $\tau = \min\{\rho(O), T\}$ be the minimum of the first exit time from O and some finite terminal time $T \in (0, \infty)$. We let X^u be the solution of (6.20) with initial condition $X_t^u = x$.⁶ We consider the cost functional

$$J(u; x, t) = \mathbb{E}_{x,t} \left[\int_t^{\tau^u} \left(f(X_s^u) + \frac{1}{2\theta} |u_s|^2 \right) ds + g(X_{\tau^u}^u) \mathbf{1}_{\{\tau^u < \infty\}} \right], \tag{6.33}$$

where we further assume that the running cost f and the terminal cost g are bounded continuous functions on $O \subset \mathbb{X}$. Our aim is to minimize J over $u \in \mathcal{A}$ and subject to the controlled dynamics (6.20).

The *value function* or *optimal cost-to-go* is defined as

$$v(x, t) = \inf_{u \in \mathcal{A}} J(u; x, t). \tag{6.34}$$

It can be shown (e.g. Fleming and Soner 2006, Section IV.3) that the optimal control is a Markovian feedback control of the form

$$u_t = c(X_t^u, t), \tag{6.35}$$

with c being a suitable feedback policy. The following theorem gives sufficient conditions for optimality in terms of the *dynamic programming* equation or *Hamilton–Jacobi–Bellman* (HJB) equation associated with (6.20) and (6.33)–(6.34); see Fleming and Soner (2006, Theorem 3.1, Section IV.3).

Theorem 6.8 (verification theorem). Let $v \in C^{2,1}(O \times [0, T])$ be a classical solution of the HJB equation

$$-\frac{\partial v}{\partial t} + \mathcal{H}(x, \nabla v, \nabla^2 v) = 0, \quad (x, t) \in O \times [0, T] \tag{6.36}$$

with boundary data

$$v(x, t) = g(x), \quad (x, t) \in (\partial O \times [0, T]) \cup (O \times \{T\}), \tag{6.37}$$

⁶ Following the relevant literature (e.g. Fleming and Soner 2006), we suppose throughout that b, σ are both continuously differentiable with bounded spatial derivatives on $\bar{O} \times [0, \infty) \subset \mathbb{R}^N \times [0, \infty)$.

and

$$\mathcal{H}(x, y, z) = \min_{h \in \mathbb{R}^s} \left\{ \frac{1}{2} \sigma \sigma^\top : z + b \cdot y + (\sigma h) \cdot y + \frac{1}{2\theta} |h^2| + f(x) \right\}. \quad (6.38)$$

If $|\nabla_x v(x, t)| \leq M(1 + |x|)$ for some constant $M > 0$, then

$$u_t^* = -\theta(\sigma(X_t^u))^\top \nabla_x v(X_t^u, t) \quad (6.39)$$

is an optimal control, i.e. $J(u^*; x, t) \leq J(u; x, t)$ for all admissible controls $u \in \mathcal{A}$, and moreover

$$v(x, t) = J(u^*; x, t). \quad (6.40)$$

Sketch of proof. We shall briefly explain the idea behind the verification theorem. To this end, we note that

$$-\theta \sigma^\top y \in \arg \min_{h \in \mathbb{R}^s} \left\{ \frac{1}{2} \sigma \sigma^\top : z + b \cdot y + (\sigma h) \cdot y + \frac{1}{2\theta} |h^2| + f(x) \right\}$$

is the unique minimizer of the right-hand side of (6.38). Therefore

$$u_t^* = -\theta(\sigma(X_t^u))^\top \nabla_x v(X_t^u, t)$$

is a candidate for the optimal control. Let

$$L^u = \frac{1}{2} a : \nabla^2 + b \cdot \nabla + (\sigma u) \cdot \nabla, \quad \text{with } a = \sigma \sigma^\top,$$

denote the generator of the controlled SDE (6.20), and consider an arbitrary admissible control $u \in \mathcal{A}$, where we suppose that $\mathbb{Q}(\tau^u < \infty) = 1$ and

$$\mathbb{E} \left[\int_t^{\tau^u} (\sigma^\top \nabla_x) v(X_s^u, s) \cdot dW_s \mid X_t^u = x \right] = 0.$$

Then Itô's formula applied to the function v yields

$$\mathbb{E}[v(X_{\tau^u}^u, \tau^u) \mid X_t^u = x] - v(x, t) = \mathbb{E} \left[\int_t^{\tau^u} \left(\frac{\partial}{\partial s} + L^u \right) v(X_s^u, s) ds \mid X_t^u = x \right].$$

Using that $v(X_{\tau^u}^u, \tau^u) = g(X_{\tau^u}^u)$ and substituting $-\partial v / \partial s$ by the right-hand side of the HJB equation (6.36), i.e. the Hamiltonian (6.38), we obtain, after dropping the \min_h term and setting $h = u$,

$$\begin{aligned} v(x, t) &= \mathbb{E} \left[- \int_t^{\tau^u} \left(\frac{\partial}{\partial s} + L^u \right) v(X_s^u, s) ds + g(X_{\tau^u}^u) \mid X_t^u = x \right] \\ &\leq \mathbb{E} \left[\int_t^{\tau^u} \left(f(X_s^u) + \frac{1}{2\theta} |u_s|^2 \right) ds + g(X_{\tau^u}^u) \mid X_t^u = x \right] \\ &= J(u; x, t), \end{aligned}$$

which shows that $v(x, t) \leq J(u; x, t)$ for all controls with finite stopping time τ^u . Repeating the calculation with the unique minimizer $h = u^*$ gives $v(x, t) = J(u^*; x, t)$, which proves the assertion. \square

Interpretation. The combination of Corollary 6.5 and Theorem 6.8 shows that the free energy $\Phi(x) = -\theta^{-1} \log \mathbb{E}[\exp(-\theta S(X)) \mid X_t = x]$ of our path functional S is identical to the value function $v(x, t) = \inf_{u \in \mathcal{A}} J(u; x, t)$ of the associated stochastic control problem, with cost function J given by (6.33) and subject to the controlled SDE (6.20).

The optimal control itself satisfies $u_t^* = -\theta(\sigma(X_t^u))^\top \nabla_x \Phi(X_t^u, t)$. Moreover, it generates the optimal change of measure, \mathbb{Q}^* , on path space (see Hartmann *et al.* 2017, Theorem 2 and Appendix D for details).

Corollary 6.9 (zero variance property). The likelihood

$$\mathcal{L}_{t,\tau}^* := \frac{d\mathbb{Q}^*}{d\mathbb{P}} \Big|_{[t,\tau^*]} = \exp(Z_{t,\tau^*}^{u^*})$$

with

$$Z_{t,\tau^*}^{u^*} := \int_t^{\tau^*} u_s^* \cdot dW_s + \frac{1}{2} \int_t^{\tau^*} |u_s^*|^2 ds$$

has the zero variance property. In particular, it holds with probability 1 that (see (6.14))

$$\mathbb{E}[\exp(-\theta S(X)) \mid X_t = x] = \exp(-Z_{t,\tau^*}^{u^*} - \theta S(X^{u^*})),$$

or equivalently

$$\Phi(x, t) = S(X^{u^*}) + \theta^{-1} Z_{t,\tau^*}^{u^*}.$$

Gradient descent interpretation. Broadly speaking, the optimal control increases the likelihood of the rare event by adding a possibly time-dependent gradient force to the dynamics, that is,

$$b \mapsto b - \sigma \sigma^\top \nabla v,$$

and it does so in such a way that the variance of the corresponding importance sampling estimator is minimized. The controlled process therefore does a gradient descent in the free energy landscape $\Phi = v$ that is only perturbed by the drift b of the uncontrolled dynamics. In the case of diffusive molecular dynamics with $b = -\nabla V$ and $\sigma = \sqrt{2\epsilon}$, the control has the effect of *tilting* the potential:

$$V \mapsto V + 2\epsilon v.$$

Note, however, that the value function v is the quantity of interest, Φ , which means that the optimal control is not directly available. Numerical methods to approximate the value function are discussed in Section 6.5 below.

We first illustrate the idea behind the optimal control formulation by again considering the previous two toy examples.

Example 6.10 (exit from a set, continued). We again consider the variational representation of the first exit time from a set. We set $O = D$ and let $T \rightarrow \infty$. Then

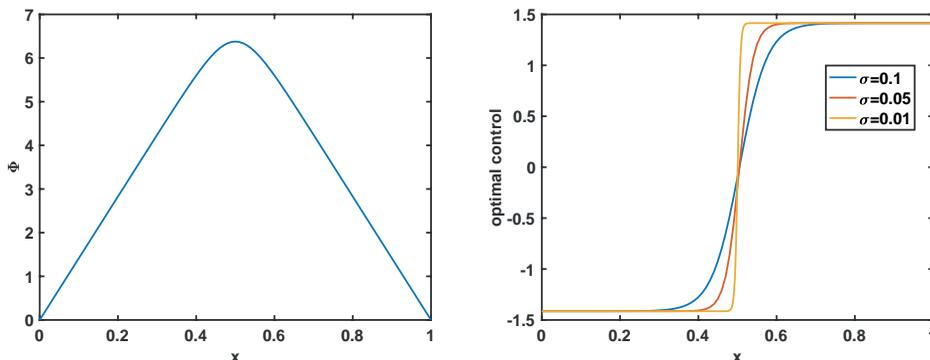


Figure 6.2. Value function $v = \Phi(\cdot; \theta = 1)$ of the exit problem and the corresponding optimal control for different temperatures σ . In the limit $\sigma \rightarrow 0$, the control converges to a Heaviside-like function.

$\min\{\tau(D), T\} \rightarrow \tau(D) =: \tau$, and Corollary 6.9 yields

$$-\theta^{-1} \log \mathbb{E}[\exp(-\theta\tau)] = \mathbb{E}\left[\tau^* + \frac{1}{2\theta} \int_0^{\tau^*} |u_s^*|^2 ds\right].$$

Here the optimal control is stationary (i.e. not explicitly time-dependent) since the CGF

$$\Phi(x) = -\theta^{-1} \log \mathbb{E}[\exp(-\theta\tau) \mid X_t = x]$$

is independent of time (by the strong Markov property of the process); it is given by

$$u_s^* = -\theta\sigma \Phi'(X_s^{*u}).$$

For $\sigma = 0.1$, $\theta = 1$ and $D = (0, 1)$, the function Φ and the resulting feedback law $c(x) = -0.1\Phi'(x)$, with a sigmoid-like shape, are shown in Figure 6.2. It can be observed that by pushing the process to the nearest boundary, the control increases the likelihood of the rare event while minimizing the variance of the corresponding importance sampling estimator.

That the likelihood of the rare event (i.e. hitting the boundary of D) is increased simply follows from the fact that the mean first hitting time is reduced, since the drift induced by the optimal control is strictly positive for $x > 0.5$ and strictly negative for $x < 0.5$.

Example 6.11 (committor probability, continued). Under the assumptions of Example 6.7, the committor equation (3.21) has the unique solution

$$q_{AB}(x) = \frac{x - a}{b - a}.$$

Setting $q_{AB} = \exp(-\theta v)$ and inserting the expression into (3.21), it readily follows that v is a classical solution of the HJB equation (6.36) which after some

simplifications reads

$$v''(x) - \theta|v'(x)|^2 = 0, \quad a < x < b.$$

Note that $v = -\theta^{-1} \log q_{AB}$ satisfies the boundary conditions $v(a) = +\infty$ and $v(b) = 0$, so it is the value function of our optimal control problem and thus

$$-\theta^{-1} \log q_{AB}(x) = \min_{u \in \mathcal{A}} \mathbb{E} \left[g(X_{\tau^u}) + \frac{1}{2\theta} \int_0^{\tau^u} |u_t|^2 dt \right]. \quad (6.41)$$

The corresponding controlled process

$$dX_t^* = \sigma u_t^* dt + \sigma dW_t,$$

with initial condition $X_0^u \in (a, b)$ and under the optimal control

$$u_t^* = -\theta \sigma v'(X_t^u) = \frac{\sigma}{X_t^u - a},$$

has a unique strong solution (see Gyöngy and Martínez 2001) with the property that the process terminates with probability 1 in finite time by hitting the target state $x = b$. Since the optimal control generates the repulsive force that pushes the particle to the target boundary $\partial B = \{b\}$ while avoiding $\partial A = \{a\}$, the control not only increases the likelihood of the rare event $\{\rho(B) < \rho(A)\}$ and minimizes the variance, but it also reduces the average length $\bar{\tau} = \mathbb{E}[\min\{\rho(A), \rho(B)\}]$ of the sampled trajectories.

Limitations. The last example reveals a phenomenon that is in some sense characteristic of problems that involve unbounded random stopping times (see Awad, Glynn and Rubinstein 2013) and that will be analysed more closely in Section 6.4.1. Under the original probability measure \mathbb{P} , the process can exit in finite time and with positive probability through either ∂A or ∂B , that is, for any $T > 0$ we have

$$\mathbb{P}(\rho(A) < T) > 0 \quad \text{and} \quad \mathbb{P}(\rho(B) < T) > 0.$$

On the other hand, the controlled process cannot leave its domain through ∂A , so we obtain

$$\mathbb{Q}^*(\rho(A) < T) = 0 \quad \text{for all } T > 0.$$

As a consequence, \mathbb{P} is not absolutely continuous with respect to \mathbb{Q}^* , even when restricted to paths of finite length T . For the problem at hand, however, the loss of absolute continuity does not lead to a breakdown of IS because in

$$q_{AB}(x) = \mathbb{E}_{\mathbb{Q}^*} \left[\exp(-\theta g(X_{\tau^*})) \frac{d\mathbb{P}}{d\mathbb{Q}^*} \right]$$

we almost surely have $\tau^* = \rho(B)$ under \mathbb{Q}^* , and therefore the events for which the inverse likelihood ratio is singular have zero probability under the measure \mathbb{Q}^* from which the IS samples are drawn. Nevertheless, the fact that the inverse likelihood ratio becomes singular indicates that the importance weights (i.e. the inverse

likelihood ratios), which have to be estimated using a numerical approximation of the controlled process, may have a large variance, which may render IS difficult in some situations.

6.3. Duality of control and estimation

The HJB equation (6.36)–(6.38) associated with the Gibbs variational principle, Corollary 6.5, has a straightforward interpretation in terms of the celebrated Feynman–Kac formula (e.g. Del Moral 2004). Assuming sufficient regularity of the solution, the semilinear HJB equation is equivalent to a linear partial differential equation where the two equations are related by a *logarithmic transformation*.

To show the equivalence, recall the definition of the second-order differential operator (2.30) associated with the *uncontrolled* dynamics (6.19):

$$L = \frac{1}{2}a : \nabla_x^2 + b \cdot \nabla_x, \quad \text{with } a = \sigma \sigma^\top,$$

Assumption 6.12. For any $T \in [0, \infty)$ and bounded open set $O \subset \mathbb{X}$ with smooth boundary ∂O , the exit time

$$\tau = \inf\{s \geq t : (X_s, s) \notin Q\}, \quad t \geq 0$$

of the set $Q = O \times [0, T]$ is almost surely finite.

Assumption 6.13. There exists a function $\psi : \mathbb{X} \times [0, T] \rightarrow \mathbb{R}$ that is $C^{2,1}(Q)$, bounded on \bar{Q} and strictly positive on Q , and that is a solution of the linear parabolic equation

$$\begin{aligned} \left(\frac{\partial}{\partial t} + L\right)\psi &= \theta f \psi && \text{in } Q, \\ \psi &= \exp(-\theta g) && \text{on } \partial Q^+, \end{aligned} \tag{6.42}$$

where $\partial Q^+ = (\partial O \times [0, T]) \cup (O \times \{T\})$ is the terminal set of the augmented process $(X_s, s)_{s \geq t}$.

By the Feynman–Kac theorem (e.g. Fleming and Soner 2006, Appendix D), the function ψ is equal to the exponential of the free energy (or moment generating function), i.e. $\psi = \Psi$ with

$$\Psi(x, t) := \exp(-\theta \Phi(x, t)) = \mathbb{E}[\exp(-\theta S(X)) \mid X_t = x]. \tag{6.43}$$

Logarithmic transformation. To show that (6.42) is equivalent to (6.36)–(6.37), we introduce the logarithmic transformation of ψ by

$$\phi(x, t) = -\theta^{-1} \log \psi(x, t),$$

which is well-defined since $\psi > 0$ on Q . Moreover, by the chain rule,

$$-\theta^{-1} \psi^{-1} \left(\frac{\partial}{\partial t} + L\right)\psi = \left(\frac{\partial}{\partial t} + L\right)\phi + \frac{\theta}{2} |\nabla_x \phi|_\alpha^2,$$

where $|y|_a^2 = y^\top ay$, $a = \sigma\sigma^\top$ denotes a weighted Euclidean norm. As a consequence, (6.42) is equivalent to the nonlinear parabolic equation

$$\begin{aligned} \left(\frac{\partial}{\partial t} + L\right)\phi + \frac{\theta}{2}|\nabla_x\phi|_a^2 + f &= 0 \quad \text{in } Q, \\ \phi &= g \quad \text{on } \partial Q^+. \end{aligned} \tag{6.44}$$

Noting that

$$-\frac{\theta}{2}|y|_a^2 = \min_{h \in \mathbb{R}^k} \left\{ (\sigma h) \cdot y + \frac{1}{2\theta}|h|^2 \right\},$$

it follows that (6.44) is equivalent to (6.36)–(6.37). Moreover ϕ is a classical solution that is equal to the free energy, that is, we have $\phi = \Phi$. We briefly discuss two notable special cases.

Stochastic control with indefinite time horizon. Letting $T \rightarrow \infty$, we have

$$\tau = \min\{T, \tau(O)\} \rightarrow \tau(O)$$

with probability 1, where $\tau(O) = \inf\{s \geq 0: X_s \notin O\}$ is the first exit time of the set $O \subset \mathbb{X}$. Since the uncontrolled process X is time-homogeneous, so is the value function, and therefore (6.44) turns into the boundary value problem

$$\begin{aligned} LF - \frac{\theta}{2}|\nabla F|_a^2 + f &= 0 \quad \text{in } O, \\ F &= g \quad \text{on } \partial O, \end{aligned} \tag{6.45}$$

for the value function F associated with a stochastic control problem with indefinite time horizon:

$$\inf_{u \in \mathcal{A}} \mathbb{E} \left[\int_0^{\rho^u} f(X_s^u) + \frac{1}{2\theta}|u_s|^2 ds + g(X_{\rho^u}^u) \right]. \tag{6.46}$$

Here we let $\tau^u = \tau^u(O)$ denote the first exit time with respect to the controlled process. In this case the value function $F = F(x)$ is time-independent; the stochastic control representations of the exit problem, Example 6.6, and the committor problem, Example 6.7, belong to this category.

Stochastic control with finite time horizon. We let $r > 0$ denote the maximum radius of any open ball $\mathcal{B}_r(\cdot) \subset O$ contained in O . If we keep $T < \infty$ fixed while letting r grow, such that $O \uparrow \mathbb{R}^N$, it follows that with probability 1

$$\tau = \min\{T, \tau(O)\} \rightarrow T.$$

In this case the solution to (6.44) converges to the solution of the backward evolution HJB equation

$$\begin{aligned} -\frac{\partial G}{\partial t} &= LG - \frac{\theta}{2}|\nabla_x G|_a^2 + f, \quad (x, t) \in \mathbb{R}^N \times [0, T), \\ G(x, T) &= g(x), \quad x \in \mathbb{R}^N \end{aligned} \tag{6.47}$$

which is associated with the finite-time stochastic control problem:

$$\inf_{u \in \mathcal{A}} \mathbb{E} \left[\int_0^T f(X_s^u) + \frac{1}{2\theta} |u_s|^2 ds + g(X_T^u) \right]. \quad (6.48)$$

Here the value function $G = G(x, t)$ will typically be time-dependent, even for time-homogeneous processes; problems involving occupation measures of measurable sets $A \subset \mathbb{R}^N$ and expressions such as

$$g(X_T) = \mathbf{1}_A(X_T), \quad \int_t^T \mathbf{1}_A(X_s) ds, \quad g(X_T) = X_T, \quad g(X_T) = X_T^2, \dots,$$

some of which appear in connection with the estimation of MSM transition matrices (see Section 4.2.3), belong in this category.

6.3.1. Nonlinear Feynman–Kac formula and backward SDE

The representation of the solution to a semilinear HJB equation in terms of a free energy or cumulant generating function is sometimes called the *nonlinear Feynman–Kac formula*. We will now discuss yet another representation in terms of a pair of forward and backward stochastic differential equations. For simplicity, we confine the subsequent discussion to the finite-time horizon case; we will comment on the random stopping time case whenever it is necessary.

A backward SDE (BSDE) is an equation of the form

$$dY_s = -h(X_s, Y_s, Z_s) ds + Z_s \cdot dW_s, \quad Y_T = g(X_T) \quad (6.49)$$

for a pair of processes $(Y, Z) = (Y_s, Z_s)_{t \leq s \leq T}$. Here h is some suitable function, called the *driver* or *generator*, and $X = (X_s)_{t \leq s \leq T}$ is the solution of the uncontrolled SDE (6.19); the role of the auxiliary process Z will be explained below.

Our aim is to derive a BSDE representation for the classical solution $G \in C^{2,1}([0, T] \times \mathbb{R}^N)$ of the semilinear HJB equation (6.47). To this end, we define the processes

$$Y_s = G(X_s, s), \quad Z_s = \theta(\sigma(X_s))^\top \nabla_x G(X_s, s). \quad (6.50)$$

By Itô's formula,

$$dG(X_s, s) = \left(\frac{\partial}{\partial s} + L \right) G(X_s, s) ds + (\sigma^\top \nabla_x G)(X_s, s) \cdot dW_s, \quad (6.51)$$

which upon inserting (6.47) and (6.50) yields the forward–backward SDE (FBSDE) system

$$\begin{aligned} dX_s &= b(X_s) ds + \sigma(X_s) \cdot dW_s, \quad X_t = x, \\ dY_s &= -f(X_s) ds + \frac{1}{2\theta} |Z_s|^2 ds + \theta^{-1} Z_s \cdot dW_s, \quad Y_T = g(X_T) \end{aligned} \quad (6.52)$$

for the triple (X, Y, Z) . Note that, by definition, Y is continuous and adapted to the filtration generated by the Brownian motion W (like X). Moreover, the process Z is

predictable and square-integrable, in accordance with the interpretation of $Z_s = -u_s^*$ as a control.

Further, note that the equation for Y in (6.52) must be understood as a *backward* SDE rather than a *time-reversed* SDE, since, by definition, Y_s at time $s \in [t, T]$ is measurable with respect to the filtration generated by the Brownian motion $(W_r)_{0 \leq r \leq s}$, whereas a time-reversed version of Y_s would depend on W_T via the terminal condition $Y_T = g(X_T)$, which would require a larger filtration. The solution to (6.52) is therefore a triplet (X, Y, Z) , and since Y is adapted, it follows that Y_t is a deterministic function of the initial data $X_t = x$ only, that is,

$$Y_t = G(x, t) = \Phi(x, t).$$

We summarize the observations in the following lemma (e.g. Pham 2009).

Lemma 6.14 (FBSDE representation of the free energy). Let the function $G \in C^{2,1}([0, T] \times \mathbb{R}^N) \cap C([0, T] \times \mathbb{R}^N)$ denote the solution of (6.47) such that

$$|\nabla_x G(x, t)| \leq C(1 + |x|^r)$$

for some $C, r > 0$ uniformly in $t \in [0, T]$. Then (Y, Z) defined by

$$Y_s = G(X_s, s), \quad Z_s = \theta(\sigma(X_s))^\top \nabla_x G(X_s, s), \quad t \leq s \leq T,$$

is the solution to the BSDE

$$dY_s = -f(X_s) ds + \frac{1}{2\theta} |Z_s|^2 ds + \theta^{-1} Z_s \cdot dW_s, \quad Y_T = g(X_T). \quad (6.53)$$

In particular, $Y_t = G(x, t)$.

The representation of semilinear PDEs by BSDEs goes back to the seminal work of Pardoux and Peng (1990), who have even shown that the equivalence between BSDEs and HJB equations extends to the case when the HJB equation has only a viscosity solution; see Pham (2009, Theorem 6.3.3). When the terminal data g and the driver satisfy some Lipschitz condition, existence and uniqueness of the BSDE can be proved by standard fixed-point arguments; see e.g. Pardoux and Tang (1999).

Under suitable regularity assumptions, the statement of the theorem remains valid when the finite time T is replaced by a random stopping time τ , e.g. the first exit time from a bounded set. Existence and uniqueness of FBSDEs with unbounded stopping time and bounded terminal costs have been studied in Kobylanski (2000), whereas the case of unbounded terminal cost has been treated in Delbaen, Hu and Richou (2011).

Remark 6.15. A remark on the role of the control variate Z is in order, since (6.52) has $m + 1$ equations but $m + k + 1$ unknowns. In (6.49), let $h = 0$ and consider the BSDE

$$dY_s = Z_s \cdot dW_s, \quad Y_T = X_T.$$

Then both $(Y, Z) = (X_t, 0)$ and $(\tilde{Y}, \tilde{Z}) \equiv (X_T, 0)$ satisfy the equation, and in fact there are infinitely many possible pairs with this property. Clearly, \tilde{Y} is not adapted to the filtration $\mathcal{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ generated by the Brownian motion, because $\tilde{Y}_t = X_T$ is not measurable with respect to $\mathcal{F}_t \subset \mathcal{F}_T$ for $t < T$. An adapted version of Y can be obtained by replacing $\tilde{Y}_t = X_T$ with its best approximation in L^2 , i.e. $\tilde{Y}_t \mapsto \mathbb{E}[\tilde{Y}_t | \mathcal{F}_t]$. Since

$$Y_t = \mathbb{E}[Y_t | \mathcal{F}_t] = \mathbb{E}\left[X_T - \int_t^T \xi_s \cdot dW_s \mid \mathcal{F}_t\right] = \mathbb{E}[X_T | \mathcal{F}_t]$$

for any square-integrable and adapted process ξ , we conclude that Y with $Y_t = \mathbb{E}[\tilde{Y}_t | \mathcal{F}_t]$ is a martingale.

As a consequence, the martingale representation theorem (e.g. Øksendal 2003, Chapter 4.3) states that there exists a unique predictable process Z such that

$$Y_t = Y_0 + \int_0^t Z_s \cdot dW_s,$$

which implies

$$Y_t = Y_T - \int_t^T Z_s \cdot dW_s = X_T - \int_t^T Z_s \cdot dW_s.$$

Hence we can think of the auxiliary process Z as a control that guarantees that Y is adapted, where the number of components of Z agrees with the number of independent components of the Brownian motion.

Zero variance property. The role of the process Z in the FBSDE representation of the HJB equation is not only to guarantee that Y is adapted, but it can be literally interpreted as a control since

$$Z_s = \theta(\sigma(X_s))^\top \nabla G(X_s, s),$$

even though it is evaluated along the uncontrolled process X rather than the controlled process X^u . Here the control Z_s plays the role of a control variate that produces a zero variance estimator. To see this, consider the solution (X, Y, Z) of (6.52) and set

$$\mathcal{L}_{t,T}^{-Z} = - \int_t^T Z_s \cdot dW_s - \frac{1}{2} \int_t^T |Z_s|^2 ds.$$

Using (6.52), the last expression can be recast as

$$\theta^{-1} \mathcal{L}_{t,T}^{-Z} = Y_t - \int_t^T f(X_s) ds - g(X_T)$$

where we have used that $Y_T = g(X_T)$. Therefore, using the identification of Y_t with the free energy $\Phi(x, t)$, we have with probability 1

$$\Phi(x, t) = S(X) + \theta^{-1} \mathcal{L}_{t,T}^{-Z}. \quad (6.54)$$

Later we will explain how (6.54) can be used to systematically devise robust free energy estimators when a numerical approximation of Z is available.

Importance sampling for FBSDEs. Even though \mathcal{L}^{-Z} has the form of the optimal log-likelihood in (6.14), it is important to realize that equation (6.54) does not involve any control, in that all quantities depend on the uncontrolled forward process X . This may be advantageous when the existence of the Radon–Nikodým derivative is not guaranteed. However, from a computational perspective there may be situations in which it is difficult to sample the terminal condition $g(X_\tau)$ by forward trajectories, in which case it may be advantageous to bias (i.e. control) the forward dynamics.

Specifically, we consider a change of drift of the form

$$b \mapsto b + \sigma w \tag{6.55}$$

for some adapted, square-integrable process $w = (w_s)_{t \leq s \leq T}$ that may or may not depend on the state of the process X^w with the new drift. Under this change of drift, using the identification

$$Y_s^w = G(X_s^w, s), \quad Z_s^w = \theta(\sigma(X_s^w))^\top \nabla_x G(X_s^w, s),$$

the original FBSDE (6.52) turns into

$$\begin{aligned} dX_s^w &= b(X_s^w) ds + \sigma(X_s^w) w_s ds + \sigma(X_s^w) dW_s, & X_t^w &= x, \\ dY_s^w &= -h^w(X_s^w, Y_s^w, Z_s^w) ds + \theta^{-1} Z_s^w \cdot dW_s, & Y_T^w &= g(X_T^w), \end{aligned} \tag{6.56}$$

with the driver

$$h^w(x, y, z) = -\frac{1}{2\theta} |z|^2 - \theta^{-1} z \cdot w + f(x). \tag{6.57}$$

It can be readily seen that (6.56)–(6.57) and (6.52) represent the same HJB equation (6.47). The change of drift furnishes an exponential change of measure in the BSDE solution (i.e. in the free energy functional). As a consequence, *every* estimator based on the expression

$$\Phi(x, t) = S(X^w) + \theta^{-1} \mathcal{L}_{t,T}^{-Z^w} - \theta^{-1} w \cdot Z^w \tag{6.58}$$

has zero variance under the probability measure $\mathbb{Q} = \mathbb{Q}^w$ generated by the controlled process X^w , where w is any adapted and square-integrable process.

The main conclusion is that we can change the drift of the forward SDE by modifying the control without affecting the variance of the free energy estimator (Kebiri, Neureither and Hartmann 2019, Hartmann, Kebiri, Neureither and Richter 2019). Having a zero variance estimator is of course only useful under the assumption that it is possible to (approximately) solve the BSDE associated with (6.52) or (6.56). Similar ideas along these lines have been suggested by Bender and Moseler (2010), who use a change of drift with the aim of reducing the variance of the BSDE simulation.

6.3.2. Conditioning and Doob's h -transform

We saw on page 620 that the optimal change of measure has the form of a conditional probability. We will now argue that the zero variance IS estimator, conditioned on the rare event, is consistent with the optimal control point of view.

For simplicity we confine our attention to deterministic stopping times and terminal costs only and thus set $f = 0$. Specifically, we consider only events of the form $\{X \in E\} = \{X_T \in B\}$ for some measurable set $B \subset \mathbb{X}$ deterministic terminal time $T > 0$. (See Theorem 6.26 below for a generalization.) We call

$$\mathbb{P}_{x,t}(\cdot) = \mathbb{P}(\cdot \mid X_t = x),$$

and we define the function (see Section 4.2.3)

$$h(x, t) = \mathbb{P}_{x,t}(X_T \in B). \quad (6.59)$$

Then, for any $s \in (t, T]$, it follows by the Markov property of X that

$$\begin{aligned} h(x, t) &= \int_{\mathbb{X}} \mathbb{P}_{x,t}(X_T \in B \mid X_s = y) d\mathbb{P}_{x,t}(X_s = y) \\ &= \int_{\mathbb{X}} \mathbb{P}_{y,s}(X_T \in B) d\mathbb{P}_{x,t}(X_s = y) \\ &= \mathbb{E}_{x,t}[h(X_s, s)], \end{aligned}$$

where $\mathbb{E}_{x,t}[\cdot] = \mathbb{E}[\cdot \mid X_t = x]$ denotes the expectation over all paths starting at $X_t = x$. As a consequence,

$$\left(\frac{\partial}{\partial t} + L\right)h(x, t) = \lim_{s \searrow t} \frac{1}{s-t} (\mathbb{E}_{x,t}[h(X_s, s)] - h(x, t)) = 0, \quad (6.60)$$

in other words, h is space–time harmonic. Itô's formula implies that the process $L_s = h(X_s, s)$ is a positive (local) martingale that can be interpreted as the likelihood ratio of a change of measure; we refrain from going into details and refer to [Rogers and Williams \(2000, Chapter IV.39–40\)](#) for further reading.

Here we shall only provide some intuition and explain the connection to stochastic optimal control. To this end, we call $\mathbb{Q} = \mathbb{P}(\cdot \mid X_T \in B)$ the conditioning of the reference measure \mathbb{P} , and define the expectation of some bounded and measurable function g conditional on $\{X_T \in B\}$ as

$$\mathbb{E}_{\mathbb{Q}}[g(X_s) \mid X_t = x] = \mathbb{E}[g(X_s) \mid X_T \in B, X_t = x]. \quad (6.61)$$

We seek an explicit expression for the associated semigroup. To this end, we suppose that X has a transition density that for simplicity, but with a slight abuse of the notation in (2.1), is denoted by p :

$$\mathbb{P}(X_s \in C \mid X_t = x) = \int_C p(s-t, x, y) dy, \quad s > t,$$

where C is any measurable set. Calling $q = q(t, s, x, y)$ the density of the conditioned process (assuming it exists for any $s > t$), it follows after disintegration and

using the Markov property that

$$\begin{aligned} \mathbb{P}(X_s \in C \mid X_T \in B, X_t = x) &= \frac{\mathbb{P}(X_s \in C, X_T \in B \mid X_t = x)}{\mathbb{P}(X_T \in B \mid X_t = x)} \\ &= \frac{1}{h(x, t)} \int_C p(s - t, x, y) h(s, y) \, dy. \end{aligned}$$

In other words,

$$q(t, s, x, y) = \frac{p(t - s, x, y) h(y, s)}{h(x, t)} \tag{6.62}$$

is the transition density associated with the conditioned process with law \mathbb{Q} . This implies

$$\mathbb{E}_{\mathbb{Q}}[g(X_s) \mid X_t = x] = \frac{\mathbb{E}_{x,t}[g(X_s)h(X_s, s)]}{h(x, t)},$$

and we can formally compute the generator of the conditioned semigroup at time t :

$$\begin{aligned} (A^t g)(x) &= \lim_{s \searrow t} \frac{1}{s - t} (\mathbb{E}_{\mathbb{Q}}[g(X_s) \mid X_t = x] - g(x)) \\ &= \lim_{s \searrow t} \frac{1}{(s - t)} \frac{\mathbb{E}[g(X_s)h(X_s, s) \mid X_t = x] - h(x, t)g(x)}{h(x, t)} \\ &= \frac{1}{h(x, t)} \left(\frac{\partial}{\partial t} + L \right) (h(x, t)g(x)). \end{aligned}$$

Using that h is space–time harmonic, the last equality implies that

$$A^t g = Lg + (\sigma\sigma^\top \nabla_x \log h) \cdot \nabla g. \tag{6.63}$$

The above construction is called *Doob’s h -transform*. The corresponding SDE reads

$$dX_s^h = b^h(X_s^h, s) \, ds + \sigma(X_s^h) \, dW_s, \quad X_t^h = x \tag{6.64}$$

with drift

$$b^h = b + \sigma\sigma^\top \nabla_x \log h. \tag{6.65}$$

Zero variance property revisited. Under an additional uniform integrability condition, the previous considerations can be generalized to rare events of the form $\{X \in E\} = \{X_\tau \in B\}$, where τ is some random stopping time; this includes committor probabilities or the probability of reaching a target set before some finite time T ; for example, when $\tau = \min\{\rho(A), \rho(B)\}$, then

$$h(x) = \mathbb{P}(\rho(B) < \rho(A) \mid X_t = x)$$

does not depend on t . As a consequence $h(x) = q_{AB}(x)$ is the committor probability of hitting B before A , and the controlled dynamics (6.64) with drift

$$b^h = b + \sigma\sigma^\top \nabla \log q_{AB}(x)$$

agrees with the solution to the optimal control problem for the committor probability (see Example 6.7). This observation can be turned into the statement that

$$\mathcal{L}_T := \frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{[0,T]} = \frac{h(X_T)}{h(X_0)}$$

is the likelihood ratio between \mathbb{P} and its conditioned version $\mathbb{Q} = \mathbb{P}(\cdot \mid X_\tau \in B)$ when restricted to paths of length $T < \tau$. Since $h(X_\tau^h) = 1$, because all controlled paths end up in B , the somewhat trivial consequence is that an IS estimator based on the reweighted expectation

$$\mathbb{P}(\rho(B) < \rho(A) \mid X_0 = x) = \mathbb{E}_x[\mathbf{1}_{\{X_\tau \in B\}}] = \mathbb{E}_x[\mathbf{1}_{\{X_\tau^h \in B\}} \mathcal{L}_\tau^{-1}] = h(x)$$

is a zero variance estimator. We summarize the last steps.

Lemma 6.16. Let \mathbb{Q} be the law of paths of the process (6.64) with drift $b^h = b + \sigma \sigma^\top \nabla_x h$, where $h(x, t) = \mathbb{P}(X_\tau \in B \mid X_t = x)$ for some measurable set B and some a.s. finite stopping time τ . Then

$$\mathbb{Q}(X_\tau \in B \mid X_0 = x) = 1 \quad \text{for all } x \in \mathbb{X}. \tag{6.66}$$

Moreover, the law of X^h is the law of X conditioned on stopping in B in finite time, that is,

$$\mathbb{Q} = \mathbb{P}(\cdot \mid X_\tau \in B).$$

6.4. Connections and equivalences

In the last two subsections we have touched upon several formulations of the central RESIM problem: compute the free energy Φ of a path functional

$$S(X) = \int_0^\tau f(X_s) ds + g(X_\tau) \mathbf{1}_{\{\tau < \infty\}},$$

as defined in (6.10) with respect to a reference probability measure \mathbb{P} by a change of measure to $\mathbb{Q} \ll \mathbb{P}$. We will now review several equivalent computational tasks, all of which characterize the zero variance measure \mathbb{Q}^* that, under appropriate integrability conditions on $\exp(-\theta S)$, is defined by

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}}(X) = \exp(\theta \Phi - \theta S(X)).$$

Problem 6.17 (exponential change of measure). Find an admissible $u^* \in \mathcal{A}$ such that the path measure \mathbb{Q}^{u^*} induced by (6.20) on the space of continuous curves coincides with \mathbb{Q}^* .

Problem 6.18 (variance minimization). Find $u^* \in \mathcal{A}$ such that

$$\text{Var} \left(\exp(-\theta S(X^{u^*})) \frac{d\mathbb{P}}{d\mathbb{Q}^*} \right) = \inf_{u \in \mathcal{A}} \text{Var} \left(\exp(-\theta S(X^u)) \frac{d\mathbb{P}}{d\mathbb{Q}^u} \right).$$

Problem 6.19 (optimal control). Find $u^* \in \mathcal{A}$ that minimizes the cost functional J defined by (6.33):

$$J(u^*) = \inf_{u \in \mathcal{A}} J(u). \tag{6.67}$$

Problem 6.20 (dynamic programming equation). Find a solution v to the HJB equation (or *dynamic programming PDE*)

$$-\frac{\partial v}{\partial t} + \mathcal{H}(x, \nabla v, \nabla^2 v) = 0.$$

Problem 6.21 (Feynman–Kac representation). Solve the FBSDE

$$\begin{aligned} dX_s &= b(X_s) ds + \sigma(X_s) \cdot dW_s, & X_t &= x \\ dY_s &= -f(X_s) ds + \frac{1}{2\theta} |Z_s|^2 ds + \theta^{-1} Z_s \cdot dW_s, & Y_T &= g(X_T). \end{aligned}$$

The connections between Problems 6.17–6.21 are summarized in the following theorem, which goes back to Nüsken and Richter (2021) and which we state in a slightly modified form.

Theorem 6.22 (connections and equivalences). Problems 6.17–6.21 are equivalent in the following sense: let $v \in C^{2,1}(O \times [0, T])$ be a classical solution to Problem 6.20 and define

$$u_s^* = -\theta(\sigma^T \nabla_x v)(X_s^u, s).$$

Then we have the following.

- (i) The measure \mathbb{Q}^{u^*} induced by (6.20) with the control $u = u^*$ coincides with \mathbb{Q}^* , that is, u^* is an admissible control that yields the solution to Problem 6.17. Moreover,

$$\exp(-\theta v(x, t)) = \mathbb{E}[\exp(-\theta S(X) \mid X_t = x)]$$

is the normalization constant for \mathbb{Q}^* . Thus the solution to Problem 6.20 also yields a solution to Problem 6.17.

- (ii) The measure \mathbb{Q}^* is a zero variance measure for Problem 6.18. The random variable

$$\exp(-\theta S(X^{u^*})) \frac{d\mathbb{P}}{d\mathbb{Q}^*}$$

is \mathbb{Q}^* -a.s. constant. In other words, the solution to Problem 6.17 also solves Problem 6.18.

- (iii) There is an admissible control determined by \mathbb{Q}^* (via the likelihood ratio $d\mathbb{Q}^*/d\mathbb{P}$ and Girsanov’s theorem) that almost surely agrees with u^* and that minimizes the cost functional of Problem 6.19:

$$v(x, t) = J(u^*; x, t).$$

This is equivalent to v solving the HJB equation of Problem 6.20. Hence variance minimization (Problem 6.18) provides the solution to Problem 6.19 and, equivalently, Problem 6.20.

(iv) The pair of processes

$$Y_s = v(X_s, s), \quad Z_s = \theta(\sigma^\top \nabla_x v)(X_s, s)$$

is a solution to Problem 6.21. Conversely, the BSDE solution (Y, Z) defines the unique viscosity solution of the HJB equation in Problem 6.20 via the relation

$$Y_t = v(x, t).$$

Under Assumptions 6.12 and 6.13, the viscosity solution agrees with the classical solution. As a consequence, the solution to Problem 6.20 provides the solution to Problem 6.21 and vice versa.

We will argue below in Section 6.5 that solving any of Problems 6.17–6.21 can be reduced to approximating a path space probability measure with respect to an appropriate divergence.

6.4.1. Further duality relations

Theorem 6.22 is strongly tied to the logarithmic scale in the free energy. The duality between control and estimation, however, basically relies on convexity arguments, first and foremost Jensen's inequality. Following Dai Pra *et al.* (1996), a variational characterization of expected values similar to Theorem 6.3 holds for non-negative random variables $S = S(X) \geq 0$.

Lemma 6.23. Let S be non-negative and measurable. Then, for all $p \in [1, \infty)$, it holds that

$$(\mathbb{E}[(S(X))^p])^{1/p} = \sup \left\{ \mathbb{E}_{\mathbb{Q}} \left[S(X) \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right)^{-1/p} \right] : \mathbb{Q} \ll \mathbb{P} \right\}. \quad (6.68)$$

If the expectation on the left is finite and if S is not (almost surely) identically zero, then the supremum in (6.68) is attained at

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}} = \frac{(S(X))^p}{\mathbb{E}[(S(X))^p]}. \quad (6.69)$$

Proof. The proof due to Dai Pra *et al.* (1996, Proposition 2.5) is along the lines of the proof of the Gibbs variational principle; we give it for the reader's convenience.

In order to avoid trivial statements, we assume $S^p \in L^1(\mathbb{P})$ and $S \neq 0$ (almost surely). Again using the shorthand $\mathcal{L} = d\mathbb{Q}/d\mathbb{P}$ for the likelihood ratio and applying Jensen's inequality, it readily follows that

$$(\mathbb{E}[(S(X))^p])^{1/p} = (\mathbb{E}_{\mathbb{Q}}[(S(X))^p \mathcal{L}^{-1}])^{1/p} \geq \mathbb{E}_{\mathbb{Q}}[S(X) \mathcal{L}^{-1/p}].$$

To show the reverse inequality, we assume that there exists $n_0 \in \mathbb{N}$ such that the truncated observable

$$S_n(X) = S(X)\mathbf{1}_{\{S(X) < n\}}, \quad n \geq n_0$$

is strictly positive on a set of positive \mathbb{P} -measure, and we define a probability measure \mathbb{Q}_n by

$$\mathcal{L}_n := \frac{d\mathbb{Q}_n}{d\mathbb{P}} = \frac{(S_n(X))^p}{\mathbb{E}[(S_n(X))^p]}.$$

By definition, $\mathbb{Q}_n(0 < S < n) = 1$ for all $n \geq n_0$, and therefore

$$\mathbb{E}_{\mathbb{Q}_n}[S(X)\mathcal{L}_n^{-1/p}] = \int_{\{0 < S < n\}} (\mathbb{E}[(S_n(X))^p])^{1/p} d\mathbb{Q}_n = (\mathbb{E}[(S_n(X))^p])^{1/p}.$$

Then, by monotone convergence,

$$\sup_{n \geq n_0} \mathbb{E}_{\mathbb{Q}_n}[S(X)\mathcal{L}_n^{-1/p}] \geq \lim_{n \rightarrow \infty} (\mathbb{E}[(S_n(X))^p])^{1/p} = (\mathbb{E}[(S(X))^p])^{1/p},$$

which implies that

$$\sup \left\{ \mathbb{E}_{\mathbb{Q}} \left[S(X) \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right)^{-1/p} \right] : \mathbb{Q} \ll \mathbb{P} \right\} \geq (\mathbb{E}[(S(X))^p])^{1/p}.$$

Inserting the expression (6.69) shows that equality is attained when $S^p \in L^1(\mathbb{P})$ and $\mathbb{Q} = \mathbb{Q}^*$. □

Zero variance property. It is again a straightforward consequence of Jensen’s inequality and the strict convexity of the power function $f(x) = |x|^p$ for $p > 1$ that \mathbb{Q}^* defines a zero variance change if $p > 1$. If $p = 1$, a zero variance change of measure can be characterized by the following theorem, which is a slight variation of Awad *et al.* (2013, Theorem 4). Before we state the theorem, we adjust Assumption 6.13 according to our needs.

Assumption 6.24. Let Q be defined as in Assumption 6.12. We suppose there exists a function $h \in C^{2,1}(Q) \cap C(\bar{Q})$ that is a solution to

$$\begin{aligned} \left(\frac{\partial}{\partial t} + L \right) h &= -f \quad \text{in } Q, \\ h &= g \quad \text{on } \partial Q^+, \end{aligned} \tag{6.70}$$

where $f, g \geq 0$ are continuous and $\partial Q^+ = (\partial O \times [0, T]) \times (O \times \{T\})$ is the terminal set of $(X_s, s)_{s \geq t}$.

Assumption 6.25. Let h satisfy Assumption 6.24. We further assume that h is strictly positive inside Q , and that for any $\delta > 0$ the function

$$x \mapsto (\sigma(x))^T \nabla_x \log h(x, t)$$

is bounded on the set $G_\delta = \{(x, t) \in Q : h(x, t) > \delta\}$.

The next statement is due to [Awad *et al.* \(2013\)](#).

Theorem 6.26 (zero variance estimator for $p = 1$). Let

$$\eta(x, t) = \mathbb{E} \left[\int_t^\tau f(X_s) ds + g(X_\tau) \mid X_t = x \right],$$

where τ is a bounded stopping time according to Assumption 6.12. If h satisfies Assumptions 6.24 and 6.25, then $h = \eta$. Furthermore, letting \mathbb{Q}^* be the path space measure generated by the SDE

$$dX_s^h = b^h(X_s^h, s) ds + \sigma(X_s^h) dW_s, \quad X_t^h = x \quad (6.71)$$

with drift

$$b^h = b + \sigma \sigma^\top \nabla_x \log h, \quad (6.72)$$

then

$$\eta(x, t) = \int_t^\tau f(X_s^h) \mathcal{L}_{t,s}^{-1} ds + g(X_\tau^h) \mathcal{L}_{t,\tau}^{-1} \quad \mathbb{Q}^* \text{-a.s.},$$

where \mathcal{L} denotes the likelihood ratio

$$\mathcal{L}_{t,s} = \exp \left(\int_t^s u_s^* \cdot dW_s + \frac{1}{2} \int_t^s |u_s^*|^2 ds \right),$$

with

$$u_s^* = (\sigma(X_s^h))^\top \nabla_x \log h(X_s^h, s), \quad s \geq t.$$

Theorem 6.26 is essentially a generalization of the h -transform trick discussed in Section 6.3.2 to functionals that involve a non-zero running cost term f . Note that in this case the likelihood ratio of the running cost term appears *inside* the integral. Under additional integrability assumptions (e.g. [Awad *et al.* 2013](#), Section 4, Assumption 3) it is possible to pass to the limit $T \rightarrow \infty$ and obtain a problem with a random time horizon. Neither the formulation of Theorem 6.26 nor the corresponding hitting time problem for $T \rightarrow \infty$ have a straightforward stochastic control interpretation when $f \neq 0$.

6.4.2. Issues for unbounded random stopping times

One important feature of the control representation of the free energy is that it allows for the control of several moments at once. For example, the cost functional associated with the sampling of $S(X) = \tau$ when τ is the potentially long first exit time from a metastable set involves the minimization of $\mathbb{E}_{\mathbb{Q}}[\tau]$ under the controlled dynamics while minimizing sample variance.

The importance sampling estimator associated with Theorem 6.26 shows a rather different behaviour when it comes to problems that involve unbounded random stopping times, such as first exit or hitting times. In extreme cases, the importance sampling estimator may even have infinite simulation time, as has been argued by [Awad *et al.* \(2013\)](#) and is shown by the following example.

Example 6.27 (exit from a set, continued). We again consider the Brownian motion, $X_t = x + \sigma W_t$ starting at x and exiting from the set $D = (0, 1)$. The boundary value problem

$$\frac{\sigma^2}{2} h''(x) = -1, \quad h(0) = h(1) = 0, \tag{6.73}$$

for the *mean first exit time* (MFET), $h(x) = \mathbb{E}[\tau \mid X_0 = x]$, has the unique solution

$$h(x) = \frac{x(1-x)}{\sigma^2}. \tag{6.74}$$

According to Theorem 6.26, the zero variance importance sampling measure \mathbb{Q}^* is generated by the Doob-transformed SDE with drift $b^h(x) = \sigma^2(\log h(x))'$:

$$\dot{X}_s^h = \sigma^2 \frac{1 - 2X_s^h}{X_s^h(1 - X_s^h)} + \sigma \dot{W}_s, \quad X_0^h = x. \tag{6.75}$$

The extra drift is singular at the domain boundaries $x \in \{0, 1\}$, which implies that $\tau = \infty$ with probability 1 (under \mathbb{Q}^*). In other words, the controlled process cannot exit from the domain D . Even worse, since for any $T > 0$ the event $\{\tau < T\}$ has positive probability under P but zero probability under \mathbb{Q}^* , the reference measure is not absolutely continuous with respect to the new measure \mathbb{Q}^* , even when restricted to paths of finite length.

Note that, formally, \mathbb{Q}^* is nevertheless a zero variance change of measure, but at the expense of generating paths that are \mathbb{Q}^* -a.s. infinitely long, that is, τ^h is almost surely constant (i.e. it has zero variance) but with the constant being equal to $+\infty$.

Possible fixes. A simple resolution of the problem in Example 6.27 is to penalize long trajectories by adding a terminal cost g . In the example here, we can just add a constant $g = c > 0$, so that we replace the MFET with $h_c(x) = \mathbb{E}[\tau + c \mid X_0 = x] = h(x) + c$. By this simple trick, the MFET remains almost surely finite, because the term in the denominator in (6.75) is bounded away from 0.

The original control and the control resulting from the penalized sampling problem are shown in Figure 6.3(a). It can be observed that the control force remains finite at the set boundaries, which implies that the MFET stays finite for all initial values $x \in D$. Under \mathbb{Q}_c^* , all trajectories almost surely terminate in finite time. Since the extra terminal cost $g = c$ simply shifts the quantity of interest by a constant that can be trivially removed, \mathbb{Q}_c^* also generates a zero variance estimator for h .

As an alternative, we can replace the mean with its certainty-equivalence, i.e. the free energy

$$\Phi(x; \theta) := -\theta^{-1} \log \mathbb{E}[\exp(-\theta\tau) \mid X_0 = x].$$

Note that we have dropped the second argument t , since our process is time-homogeneous.

It is possible to extract moments of τ from the free energy by using the properties or the associated moment generating function: if there exists an $\epsilon > 0$ such that

$$\Psi(x, \theta) := \mathbb{E}[\exp(-\theta\tau) \mid X_0 = x]$$

is finite for all $\theta \in (-\epsilon, \epsilon)$, then all moments of τ exist and

$$\mathbb{E}[\tau^k \mid X_0 = x] = \frac{d}{d\theta} \Psi(x; \theta) \Big|_{\theta=0}.$$

Note that $\Psi = \exp(-\theta v)$ can be expressed in terms of the value function v of our control problem. In principle it is possible to approximate the derivatives with respect to θ by finite differences or more elaborate techniques from sensitivity analysis (e.g. Dupuis, Katsoulakis, Pantazis and Plecháč 2016, Dupuis *et al.* 2020, Tsourtis, Pantazis, Katsoulakis and Harmandaris 2015), yet the control becomes heavily penalized for small θ as equations (6.33) and (6.39) show. As a consequence the control converges to zero as $\theta \rightarrow 0$ and hence no longer leads to a reduction of the simulation time.

Remark 6.28. The reader may wonder whether the optimal control for the moment generating function can be used to directly control the variance of the estimator of the first moment for finite values of $\theta > 0$. Unfortunately the answer is negative. To see this, let

$$\tau^* = \tau \frac{d\mathbb{P}}{d\mathbb{Q}^*},$$

and suppose we want to estimate $\mathbb{E}[\tau]$ by sampling τ^* under the controlled dynamics. Using that the optimal likelihood ratio $\mathcal{L}^* = d\mathbb{Q}^*/d\mathbb{P}$, for which equality in (6.12) is attained, can be recast as

$$\mathcal{L}^* = \frac{e^{-\theta\tau}}{\Psi}, \quad \Psi = \mathbb{E}[\exp(-\theta\tau)],$$

it follows that

$$\begin{aligned} \text{Var}(\tau^*) &= \mathbb{E}_{\mathbb{Q}^*} \left[\left(\tau \frac{d\mathbb{P}}{d\mathbb{Q}^*} \right)^2 \right] - (\mathbb{E}[\tau])^2 \\ &= \mathbb{E} \left[\tau^2 \frac{d\mathbb{P}}{d\mathbb{Q}^*} \right] - (\mathbb{E}[\tau])^2 \\ &= \mathbb{E}[\exp(-\theta\tau)] \mathbb{E}[\tau^2 \exp(\theta\tau)] - (\mathbb{E}[\tau])^2. \end{aligned}$$

Now, since the random variables $U(\tau) = \exp(-\theta\tau)$ and $V(\tau) = \tau^2 \exp(\theta\tau)$ are anticorrelated, for U is decreasing whereas V is increasing, it follows that

$$\begin{aligned} \mathbb{E}[\tau^2] &= \mathbb{E}[U(\tau)V(\tau)] \\ &\leq \mathbb{E}[U(\tau)] \mathbb{E}[V(\tau)] \\ &= \mathbb{E}[\exp(-\theta\tau)] \mathbb{E}[\tau^2 \exp(\theta\tau)]. \end{aligned}$$

As a consequence,

$$\text{Var}(\tau^*) \geq \text{Var}(\tau),$$

where equality is attained for $\theta = 0$, in which case the optimal control vanishes. By an analogous argument (see [Badowski 2016](#)), it can be shown that

$$\text{Var}(\tau^*) \mathbb{E}[\tau^*] \geq \text{Var}(\tau) \mathbb{E}[\tau],$$

which implies that the reduction of the mean trajectory length due to the control does not compensate for the increase of the variance.

6.4.3. Control variate limit of the IS estimator

From a control perspective, the limit dynamics as $\theta \rightarrow 0$ is not particularly relevant, since the control vanishes in the limit. Nevertheless the IS estimator has a non-trivial limit as a control variate estimator, which we will briefly discuss with our standard example.

Example 6.29 (exit from a set, continued). Recall from the proof of Theorem 6.3 that by the strict convexity of the exponential function, equality is attained if and only if the optimal control generates a zero variance change of measure \mathbb{Q}^* that leaves the random variable $\tau + \theta^{-1} \log d\mathbb{Q}^*/d\mathbb{P}$ almost surely constant, in which case

$$\Psi(x; \theta) = e^{-\theta\tau^*} \frac{d\mathbb{P}}{d\mathbb{Q}^*} \quad \mathbb{Q}^*\text{-a.s.} \tag{6.76}$$

for all $\theta > 0$, where τ^* denotes the stopping time under the optimal control u^* . Taking logarithms on both sides and letting $\theta \rightarrow 0$, we obtain a zero variance estimator for the mean that has the form

$$\mathbb{E}_x[\tau] = \tau + \left. \frac{d}{d\theta} \right|_{\theta=0} \log \mathcal{L}^*. \tag{6.77}$$

Inserting the likelihood ratio $\mathcal{L}^* = \mathcal{L}_{\tau^*}^{u^*}$ under the optimal control $u^* = O(\theta)$ and taking the limit $\theta \rightarrow 0$, we obtain

$$\mathbb{E}_x[\tau] = \tau + \sigma^{-1} \int_0^\tau (2X_t - 1) dW_t, \tag{6.78}$$

where the expression inside the Itô integral is exactly $-\sigma h'(x)$, with h being the MFET as defined in (6.74). This should be compared to the result of the Doob h -transform, for which an extra term proportional to $(\log h(x))'$ is causing the drift to become singular at the boundaries of the domain.

Figure 6.3(b) shows a comparison of crude Monte Carlo (CMC) and control variate estimators of the mean first exit time, which illustrates the variance reduction by the additive control variate term.

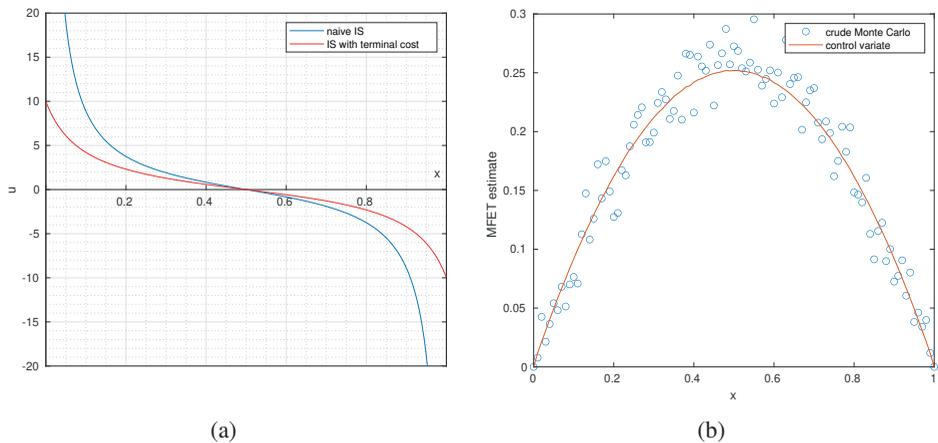


Figure 6.3. (a) Naive optimal drift for zero variance importance sampling (blue) and its regularized version including a small terminal cost (red). (b) Comparison of crude Monte Carlo (blue) and control variate (red) estimators of the mean first exit time $\mathbb{E}_x[\tau]$ of the set $D = (0, 1)$ for $N = 100$ independent realizations. The observed small bias of the control variate estimator can be attributed to the time discretization error of the underlying Euler–Maruyama discretization, which leads to an overestimation of τ .

The main conclusion of the last example is that while the control disappears from the equation in the limit $\theta \rightarrow 0$, and hence can no longer lead to a reduction of the average simulation time, the IS estimator turns into a zero variance control variate estimator.

This behaviour is what one would generally expect from the fact that the nonlinearity in the associated HJB equation vanishes as $\theta \rightarrow 0$, which leads to a linear equation for the first moment. Assuming sufficient regularity of the coefficients, the convergence of the estimator then follows from the convergence of the control, i.e. the derivative of the value function. To our knowledge, this connection has not yet been discussed in the literature on variance reduction methods for SDEs (e.g. Graham and Talay 2013).

6.5. Computational aspects

For realistic scenarios, solving the nonlinear HJB equations to compute optimal controls that can be used in an adaptive IS framework is not an option. There are two main lines of attack to approximate either the value function (i.e. the free energy or cumulant generating function) or the corresponding optimal control. Roughly, one can distinguish methods based on stochastic optimization of an appropriate loss (e.g. the cost functional or an auxiliary entropy-based functional) or time-stepping methods using the probabilistic FBSDE representation of the HJB equation. It turns out that the two approaches are intimately connected, because loss functions

also play a key role as regularization terms in the FBSDE approach, as has been pointed out in Nüsken and Richter (2021) and Richter (2022).

6.5.1. Path space approximations for finite time horizon

What the stochastic optimization methods and the FBSDE approach have in common is that they can be thought of approximation problems for the optimal path space measure \mathbb{Q}^* (see Theorem 6.22). Given a parametric family $\{\hat{\mathbb{Q}}(\alpha) : \alpha \in \mathbb{R}^n\}$ of path space measures $\hat{\mathbb{Q}}(\alpha)$, these approximation problems can then be solved by standard optimization routines, including stochastic gradient descent schemes in connection with approximations by deep neural networks (e.g. Hartmann et al. 2019, Richter 2022).

Before we point out further connections, we shall briefly introduce the key concepts regarding loss functions based on divergences between probability measures.

Divergences and loss functions. We call \mathcal{M}_1 the set of probability measures $\mathbb{Q} \ll \mathbb{P}$ on the space of continuous curves $X : [0, T] \rightarrow \mathbb{X}$, and notice that the controlled SDE (6.20) induces a map

$$\mathcal{A} \ni u \mapsto \mathbb{Q}^u \in \mathcal{M}_1,$$

which can be made explicit in terms of Radon–Nikodým derivatives via Girsanov’s theorem (6.24). As a consequence, we can elevate divergences between path measures to loss functions on vector fields. Namely, let $D : \mathcal{M}_1 \times \mathcal{M}_1 \rightarrow [0, \infty]$ be a divergence, that is,

$$D(\mathbb{P}, \mathbb{Q}) \geq 0 \quad \text{and} \quad D(\mathbb{P}, \mathbb{Q}) = 0 \quad \text{if and only if} \quad \mathbb{P} = \mathbb{Q},$$

where it is in most cases sufficient that the rightmost equality holds only up to \mathbb{Q} -null sets. Prominent examples include the KL divergence (6.11) or, more generally, f -divergences such as the χ^2 -divergence (6.15).

Using the shorthand $\mathbb{Q} = \mathbb{Q}^u$ and setting

$$\mathcal{R}_D(u) = D(\mathbb{Q}, \mathbb{Q}^*), \quad u \in \mathcal{A}, \tag{6.79}$$

we immediately see that $\mathcal{R}_D \geq 0$ and $\mathcal{R}_D(u) = 0$ if and only if $u = u^*$. It is therefore plausible that an approximation of the optimal control vector field u^* can be found by minimizing the loss function \mathcal{R}_D over the set of admissible controls $u \in \mathcal{A}$.

Definition 6.30 (loss functions). For $\mathbb{Q}, \mathbb{Q}^* \in \mathcal{M}_1$, equivalent to \mathbb{P} , we define

- (i) the relative entropy loss $\mathcal{R}_{RE} = KL(\mathbb{Q} | \mathbb{Q}^*)$,
- (ii) the cross-entropy loss $\mathcal{R}_{CE} = KL(\mathbb{Q}^* | \mathbb{Q})$,
- (iii) the variance loss $\mathcal{R}_{Var}(u) = \text{Var} \left(\frac{d\mathbb{Q}^*}{d\mathbb{Q}} \right)$,
- (iv) the log-variance loss $\mathcal{R}_{Var}^{\log}(u) = \text{Var} \left(\log \frac{d\mathbb{Q}^*}{d\mathbb{Q}} \right)$.

The above loss functions can be explicitly expressed in terms of the original IS control problem. We highlight two notable special cases and confine ourselves to finite stopping times T , for which all loss functions are well-defined.

Cross-entropy loss. To establish a connection between optimal control and cross-entropy minimization, note that (6.14) and (6.43) together with Theorem 6.8 imply that

$$S - \log \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \frac{d\mathbb{Q}}{d\mathbb{Q}^*} \right) = J(u^*).$$

Taking the expectation with respect to \mathbb{Q} and using that both \mathbb{Q} and \mathbb{Q}^* are absolutely continuous with respect to \mathbb{P} and vice versa yields

$$J(u) = J(u^*) + KL(\mathbb{Q} | \mathbb{Q}^*) \quad (6.80)$$

where u is any admissible control, u^* is the optimal control, and $\mathbb{Q} = \mathbb{Q}^u$ and $\mathbb{Q}^* = \mathbb{Q}^{u^*}$ are the corresponding path space measures. In other words, $J(u)$ equals the relative entropy loss $\mathcal{R}_{RE}(u)$ up to an additive constant.

The idea now is to seek a minimizer of $KL(\mathbb{Q} | \mathbb{Q}^*)$ in the set of probability measures $\hat{\mathbb{Q}} = \hat{\mathbb{Q}}(\alpha) \in \mathcal{M}_1$ that are generated by a parametric family of controls $\hat{u} = \hat{u}(\alpha)$, $\alpha \in \mathbb{R}^n$. By (6.14) the optimal change of measure is known up to a normalizing factor, which enters (6.80) only as an additive constant (i.e. a function of the initial data (x, t) but that does not depend on the realizations of the process). Nevertheless, minimizing $\mathcal{R}_{RE}(\alpha) = KL(\hat{\mathbb{Q}}(\alpha) | \mathbb{Q}^*)$ over $\alpha \in \mathbb{R}^n$ is not easily achieved since the functional may have several local minima.

With a little trick, however, we can turn the minimization of (6.80) into a feasible minimization problem, simply by flipping the arguments. To this end, we define

$$\mathcal{R}_{CE}(\alpha) = KL(\mathbb{Q}^* | \hat{\mathbb{Q}}(\alpha)). \quad (6.81)$$

Clearly, the relation (6.80) does not hold with arguments in the Kullback–Leibler (or KL) divergence term reversed, since $KL(\cdot | \cdot)$ is not symmetric; nevertheless,

$$\mathcal{R}_{RE}(\alpha) \geq 0, \quad \mathcal{R}_{CE}(\alpha) \geq 0 \quad \text{and} \quad \mathcal{R}_{RE}(\alpha) = 0 \quad \text{if and only if} \quad \mathcal{R}_{CE}(\alpha) = 0, \quad (6.82)$$

where the minimum is attained if and only if $\hat{\mathbb{Q}} = \mathbb{Q}^*$. Hence, by continuity of the relative entropy, we may expect that by minimizing the ‘wrong’ functional \mathcal{R}_{CE} we get close to the optimal change of measure, provided that the optimal \mathbb{Q}^* can be approximated by our parametric family $\hat{\mathbb{Q}}$.

Ignoring additive constants, the cross-entropy minimization problem is readily seen to be equivalent to the maximization of the functional (Zhang *et al.* 2014):

$$CE(\alpha) = \mathbb{E}[\log \mathcal{L}(\alpha) \exp(-S(X))], \quad (6.83)$$

where $\log \mathcal{L} = \log(d\hat{\mathbb{Q}}/d\mathbb{P})$ denotes the the log likelihood ratio between controlled and uncontrolled trajectories. If \hat{u} is a linear combination of suitable basis functions

$\phi_1, \dots, \phi_n: \mathbb{X} \times [0, \infty) \rightarrow \mathbb{R}$, for example

$$\hat{u}_s(\alpha) = -\theta \sigma(X_s)^\top \sum_{i=1}^n \alpha_i \nabla_x \phi_i(X_s, s), \tag{6.84}$$

then by Girsanov’s theorem $\text{CE}(\alpha)$ is quadratic in the unknown $\alpha = (\alpha_1, \dots, \alpha_n)$. This then implies that the necessary optimality condition is of the form

$$K\alpha = r, \tag{6.85}$$

where $K = (K_{ij})_{1 \leq i, j \leq N}$ and $r = (r_i)_{1 \leq i \leq N}$ are given by

$$\begin{aligned} K_{ij} &= \mathbb{E} \left[\exp(-S(X)) \int_0^T (\sigma^\top \nabla_x \phi_i)(X_s, s) (\sigma^\top \nabla_x \phi_j)(X_s, s) ds \right], \\ r_i &= -\mathbb{E} \left[\exp(-S(X)) \int_0^T (\sigma^\top \nabla_x \phi_i)(X_s, s) \cdot dW_s \right]. \end{aligned} \tag{6.86}$$

Remark 6.31 (iterative sampling scheme). Note that the average in equation (6.86) is over the uncontrolled realizations X . It is easy to see that the matrix S is positive definite if the basis functions ϕ_i are linearly independent, which implies that equation (6.85) has a unique solution and our necessary condition is in fact sufficient. Nevertheless it may happen in practice that the coefficient matrix K is badly conditioned, in which case it may be advisable to evaluate the coefficients using importance sampling or a suitable annealing strategy; see Zhang *et al.* (2014) and Hartmann, Schütte and Zhang (2016) for further details.

Log-variance loss. The log-variance loss has an interpretation in terms of the solution to the BSDE (6.53), as has been noted in Nüsken and Richter (2021) and Richter (2022). The idea is to replace the auxiliary variable $Z_s = -u_s^*$ by an arbitrary parametric auxiliary variable $\hat{Z}_s = -\hat{u}$ with a Markovian control $\hat{u}_s(\alpha) = c(X_s, s; \alpha)$ and interpret the BSDE solution

$$Y_t - y_0 = - \int_0^t f(X_s) ds - \frac{1}{\theta} \int_0^t c(X_s, s) \cdot dW_s + \frac{1}{2\theta} \int_0^t |c(X_s, s)|^2 ds \tag{6.87}$$

as a controlled forward process for any given initial condition $Y_0 = y_0$. Here, unlike the previous linear regression formulation in the cross-entropy method, the parametrization of the control can be most efficiently done using deep neural networks and automatic differentiation tools; see E, Han and Jentzen (2017), Hartmann *et al.* (2019) and Nüsken and Richter (2021). (Nonetheless, we use the same parameter α to denote the parametric family.)

Clearly $Y_T = g(X_T)$ if $\hat{u} = u^*$ is the optimal control; for an arbitrary admissible control $\hat{u}_s(\alpha) = c(X_s, s, \alpha)$, using (6.87), Girsanov’s theorem and the explicit form

(6.14) of the optimal change of measure, we obtain

$$\begin{aligned}\mathcal{R}_{\text{Var}}^{\log}(\alpha) &= \text{Var}\left(\log\left(\frac{d\mathbb{Q}^*}{d\hat{\mathbb{Q}}}\right)\right) \\ &= \text{Var}\left(\log\left(\frac{d\hat{\mathbb{Q}}}{d\mathbb{P}}\frac{d\mathbb{P}}{d\mathbb{Q}^*}\right)\right) \\ &= \text{Var}\left(\theta\left(y_0 - Y_T - \int_0^t f(X_s) ds\right) + \theta(S(X) + \Phi)\right) \\ &= \theta^2 \text{Var}(Y_T - g(X_T)),\end{aligned}$$

As a consequence, ignoring the multiplicative constant $\theta^2 > 0$,

$$\mathcal{R}_{\text{Var}}^{\log}(\alpha) := \text{Var}(Y_T - g(X_T)), \quad (6.88)$$

Note that the log-variance loss does not depend on y_0 , since the variance of a random variable is independent of constant shifts, in contrast to the quadratic loss

$$\mathcal{R}_2(\alpha, y_0) = \mathbb{E}[(Y_T - g(X_T))^2]$$

that has been used in [E *et al.* \(2017\)](#) to compute FBSDE solutions using deep neural network approximations and that has an explicit dependence on y_0 .

Remark 6.32 (importance sampling for FBSDEs). The variance and the log-variance losses are computed by taking averages over the reference measure \mathbb{P} , in accordance with the fact that forward process X in the FBSDE (6.52) is uncontrolled. There may be situations, however, in which it is advisable to bias the forward dynamics and sample the terminal condition from a controlled forward process X^w as in (6.56)–(6.57).

For example, the log-variance under a controlled forward process X^w reads

$$\mathcal{R}_{\text{Var}^w}^{\log}(\alpha) = \text{Var}_{\mathbb{P}^w}\left(\log\frac{d\mathbb{Q}^*}{d\hat{\mathbb{Q}}}\right) = \text{Var}_{\mathbb{P}^w}(Y_T^w - g(X_T^w)),$$

with \mathbb{P}^w being the path measure induced by X^w , and

$$Y_T^w - y_0 = -\frac{1}{\theta} \int_0^T w_s \cdot c_s ds - \int_0^T f_s ds - \frac{1}{\theta} \int_0^T c_s \cdot dW_s + \frac{1}{2\theta} \int_0^T |c_s|^2 ds, \quad (6.89)$$

where we have used the abbreviations $c_s = c(X_s^w, s)$ and $f_s = f(X_s^w)$. Using an additional bias w can be useful in situations in which the terminal condition is difficult to sample, resulting in a large sample variance of the loss function or its gradient (see [Kebiri *et al.* 2019](#), [Hartmann *et al.* 2019](#)).

(Non-)robustness of loss functions. Even though the aforementioned loss functions can all be used inside a regression scheme to approximate FBSDE solutions, and even though they share many variational properties, their Monte Carlo estimators have rather different properties.

One of the key observations in Nüsken and Richter (2021) is that the derivatives of the log-variance loss and the relative entropy are equivalent when the former is evaluated at $w = u$. Letting

$$\delta\mathcal{R}(u) \cdot \zeta := \left. \frac{d}{d\epsilon} \mathcal{R}(u + \epsilon\zeta) \right|_{\epsilon=0} \tag{6.90}$$

denote the directional (Gâteaux) derivative of the loss \mathcal{R} at u along the test functions $\zeta \in \mathcal{U}$ (e.g. square-integrable and adapted to the filtration generated by the Brownian motion), and further letting

$$\mathcal{R}_{\text{Var}_w}^{\text{log}}(u) = \text{Var}_{\mathbb{P}^w} \left(\log \frac{d\mathbb{Q}^*}{d\mathbb{Q}} \right)$$

denote the log-variance loss under the reference measure \mathbb{P}^w (see Remark 6.32), it holds that

$$\frac{1}{2} \delta\mathcal{R}_{\text{Var}_w}^{\text{log}}(u) \cdot \zeta \Big|_{w=u} = \delta\mathcal{R}_{\text{RE}}(u) \cdot \zeta, \tag{6.91}$$

assuming that $u, w \in \mathcal{A}$ are admissible and the two loss functions are Gâteaux-differentiable.

The last equation complements the observation that the relative entropy loss equals the cost functional (6.33) up to an additive constant, in that it establishes a connection between the minimization of the cost functional and the variational (forward) formulation of the FBSDE in terms of the log-variance loss. The surprising bit is not that both uniquely determine the optimal control $u = u^*$, but that the gradients of the corresponding loss functions agree for $u \neq u^*$ (up to a multiplicative factor).

Note, moreover, that computing the derivatives of the relative entropy loss requires differentiating both the SDE solution, X^u , and running and terminal costs, f and g , whereas differentiating f and g is not necessary to compute the derivative of the log-variance loss. This opens the door to gradient-free implementations when using the log-variance loss, as has been pointed out in Richter (2022).

The log-variance loss has other remarkable properties, one of which is that the variance of the gradient of the Monte Carlo estimator vanishes at the optimum; the next lemma, with slight modifications, is taken from Nüsken and Richter (2021).

Lemma 6.33 (robustness of log-variance estimator). Calling $\widehat{\mathcal{R}}_{\text{Var}_v}^{\text{log},m}$ the crude Monte Carlo estimator of the log-variance loss based on m independent realizations of X_T^w and Y_T^w as given by (6.89), then, independently of $w \in \mathcal{A}$,

$$\text{Var}(\delta\widehat{\mathcal{R}}_{\text{Var}_v}^{\text{log},m}(u^*) \cdot \zeta) = 0 \quad \text{for all } \zeta \in \mathcal{U}. \tag{6.92}$$

Intriguingly, cross-entropy and relative entropy losses do *not* have this property. The log-variance loss estimator moreover has the feature that its relative error is

uniformly bounded in the state space dimension, \mathcal{N} , if the target measure \mathbb{Q}^* is a product measure – a property that is again not shared by the cross-entropy loss (Nüsken and Richter 2021, Section 5).

Related work. We should mention that the connection between HJB equations and FBSDEs is not new and not specific to the situation here; see e.g. E *et al.* (2017), Sirignano and Spiliopoulos (2018), Huré, Pham, Bachouch and Langrené (2021) and Bachouch, Huré, Langrené and Pham (2022) for a treatment of semilinear partial differential equations of HJB type, with a focus on the approximation by deep neural networks, or Beck *et al.* (2019) and Pham, Warin and Germain (2021) for the fully nonlinear case; see also E, Han and Jentzen (2021) for a review.

Regarding the backward time discretization of uncoupled FBSDE on finite time-horizon, like the ones considered here, we refer to the rich literature on least-squares Monte Carlo, e.g. Gobet, Lemor and Warin (2005), Bender and Moseler (2010), Bender and Steiner (2012), Turkedjiev (2013) and Gobet and Turkedjiev (2016).

6.5.2. Stochastic gradient descent for unbounded stopping times

The cross-entropy method and the FBSDE-based approach have their merits when the stopping time is bounded. For problems with unbounded random time horizon, sampling of the terminal condition can be difficult, especially when the stopping time distribution is very broad. (Backward time-stepping schemes such as least-squares Monte Carlo can be even more problematic.)

In this case the minimization of the cost functional can be a viable alternative (Hartmann and Schütte 2012). We describe the core idea. To this end, we use the notation of Section 6.3 assuming that $\tau = \rho(O)$ is some possibly unbounded random stopping time (e.g. a first hitting time). We further assume that the dynamics is time-homogeneous which, by the strong Markov property, implies that the value function does not depend on t . We suppose that v can be approximated in some suitable norm by the Galerkin ansatz

$$\hat{v}(x) = \sum_{i=1}^n \alpha_i \phi_i(x), \quad (6.93)$$

where $\phi_1, \dots, \phi_n: \bar{O} \rightarrow \mathbb{R}$ are smooth basis functions whose choice depends on the problem at hand. The optimal control can then be approximated by the linear combination of the $\sigma^\top \nabla \phi_i$:

$$\hat{u}_s = -\theta \sigma(X_s^u)^\top \sum_{i=1}^n \alpha_i \nabla \phi_i(X_s^u). \quad (6.94)$$

Plugging the above representation into (6.33) yields the following finite-dimensional optimization problem: minimize

$$J(\hat{u}) = \mathbb{E} \left[\int_0^{\hat{\tau}} \left(f(X_s^{\hat{u}}) + \frac{1}{2\theta} |\hat{u}_s|^2 \right) ds + g(X_{\hat{\tau}}^{\hat{u}}) \mathbf{1}_{\{\hat{\tau} < \infty\}} \right] \quad (6.95)$$

over the controls \hat{u} , where $X^{\hat{u}}$ is the solution of the SDE (6.20) with control $u = \hat{u}$ and $\hat{\tau} = \tau^{\hat{u}}$ denotes the stopping time under the controlled dynamics.

Let us define $\hat{J}(\alpha) = J(\hat{u}(\alpha))$, with the shorthand $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{R}^n$. Because of the dependence of the process X^α and the random stopping time $\hat{\tau} = \tau^{\hat{u}(\alpha)}$ on the parameter α , the functional \hat{J} is not quadratic in α , but it has been shown (Lie 2016) that it is strongly convex if the ansatz functions ϕ_i are non-overlapping. In this case \hat{J} has a unique minimum, which suggests doing a gradient descent in the parameter α :

$$\alpha^{(m+1)} = \alpha^{(m)} - h_m \nabla \hat{J}(\alpha^{(m)}). \tag{6.96}$$

Here $(h_m)_{m \geq 0}$ is a sequence of step sizes (or *learning rates*) that go to zero as $m \rightarrow \infty$, and the gradient $\nabla \hat{J}(\alpha)$ must be interpreted in the sense of (6.90). Then the gradient $\nabla \hat{J}(\alpha)$ has the components

$$\frac{\partial \hat{J}}{\partial \alpha_k} = -\delta J(\hat{u}(\alpha)) \cdot (\theta \sigma^\top \nabla \phi_k). \tag{6.97}$$

Gradient computation. The explicit expression for the gradient can be computed by Girsanov’s theorem. To this end, we assume that $\mathbb{Q}(\hat{\tau} < \infty) = 1$ and introduce the shorthand

$$\ell(X^{\hat{u}}, \hat{u}) = \int_0^{\hat{\tau}} \left(f(X_s^{\hat{u}}, s) + \frac{1}{2\theta} |\hat{u}_s|^2 \right) ds + g(X_{\hat{\tau}}^{\hat{u}}).$$

Then there exists $\mathbb{Q}^\epsilon \ll \mathbb{P}$, such that

$$\begin{aligned} \frac{d}{d\epsilon} J(\hat{u} + \epsilon \zeta) \Big|_{\epsilon=0} &= \frac{d}{d\epsilon} \mathbb{E}[\ell(X^{\hat{u} + \epsilon \zeta}, \hat{u} + \epsilon \zeta)] \Big|_{\epsilon=0} \\ &= \frac{d}{d\epsilon} \mathbb{E} \left[\ell(X, \hat{u} + \epsilon \zeta) \frac{d\mathbb{Q}^\epsilon}{d\mathbb{P}} \right] \Big|_{\epsilon=0}, \end{aligned}$$

where

$$\frac{d\mathbb{Q}^\epsilon}{d\mathbb{P}} \Big|_{[0, \tau]} = \exp(Z_\tau^{u + \epsilon \zeta}).$$

If we formally differentiate under the expectation and then switch back to the controlled process $X^{\hat{u}}$ under the reference measure \mathbb{P} , we obtain (see Lie 2021 for details)

$$\delta J(\hat{u}) \cdot \zeta = \mathbb{E} \left[\ell(X^{\hat{u}}, \hat{u}) \int_0^{\hat{\tau}} \zeta_s \cdot dW_s + \frac{1}{\theta} \int_0^{\hat{\tau}} \hat{u}_s \cdot \zeta_s ds \right].$$

Under the optimal control u^* , we have (see Hartmann *et al.* 2017, Lemma 1)

$$\delta J(u^*) \cdot \zeta = 0 \quad \text{for all } \zeta \in \mathcal{U}.$$

Finally, we get the explicit expression for the gradient from (6.97):

$$\frac{\partial \hat{J}}{\partial \alpha_k} = -\mathbb{E} \left[\theta \ell(X^{\hat{u}}, \hat{u}) \int_0^{\hat{\tau}} (\sigma^\top \nabla \phi_k)(X_s^{\hat{u}}) \cdot dW_s + \int_0^{\hat{\tau}} \hat{u}_s \cdot (\sigma^\top \nabla \phi_k)(X_s^{\hat{u}}) ds \right]. \quad (6.98)$$

Hartmann and Schütte (2012) have derived a discrete representation of the gradient for diffusive molecular dynamics using Euler's method.

Limitations. In principle the gradient can be estimated by Monte Carlo. In practice, however, controlling the sample variance is extremely tricky for path functionals that involve very long or even unbounded stopping times (Birrell and Rey-Bellet 2020, Birrell, Katsoulakis and Rey-Bellet 2021). Therefore the robust estimation of the gradient (6.98) requires suitable variance reduction techniques for the estimation of the gradient (see Arampatzis, Katsoulakis and Rey-Bellet 2016).

Another difficulty is the choice of basis functions. Even though the cost functional (6.95) has nice convexity properties when the basis functions have non-overlapping supports, the canonical choice for high-dimensional problems is radial basis functions, such as Gaussians (Hartmann and Schütte 2012, Quer, Donati, Keller and Weber 2018). In this case, due to a lack of convexity, one cannot expect convergence of the gradient flow (6.96) to the global minimizer, even with sophisticated control strategies for the learning rate. This explains why a good initial guess for the discretized control (6.94) is crucial for the convergence of the method. Methods to generate good initial guesses can be based on biased MD simulations (e.g. using metadynamics (Bussi *et al.* 2020) or variational autoencoders (Belkacemi *et al.* 2022)) that generate bias potentials from simulation data in an automated fashion; see Lie and Quer (2017) and Borrell, Quer, Richter and Schütte (2022) for an application of biasing methods to generate initial guesses for path space importance sampling.

7. Concluding remarks

We have seen how the timescale gap in molecular dynamics can be tackled when switching from long direct MD simulations in atomistic simulation to the level of transfer operators and linear evolution equations, where long timescales are accessible via dimension reduction and techniques based on variational formulations. Moreover, we have discussed how these dimension reduction techniques can be combined with modern machine learning strategies to yield efficient algorithms suitable for realistic molecular systems and biologically relevant timescales.

In addition, we have also approached rare event simulation from another angle by asking how to optimally reduce the variance of direct numerical simulation techniques. Here we have seen that several approaches (importance sampling, optimal control, forward–backward SDEs, etc.) have equivalent variational formulations too, and that, again, modern machine learning techniques can be used to find

efficient algorithms. This family of low-variance rare event simulation techniques is not as prominent in molecular dynamics as the transfer operator-based approaches, but the connections between the two fields are inspiring and promising.

Acknowledgements

This work has been partially funded by the Deutsche Forschungsgemeinschaft (DFG) through grant CRC 1114 *Scaling Cascades in Complex Systems* (project no. 235221301) and under Germany's Excellence Strategy through grant EXC-2046 *The Berlin Mathematics Research Center MATH+* (project no. 390685689).

References

- S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso and A. M. Stuart (2015), Importance sampling: Computational complexity and intrinsic dimension, *Statist. Sci.* **32**, 405–431.
- D. J. Alford-Lago, C. W. Curtis, A. T. Ihler and O. Issan (2022), Deep learning enhanced dynamic mode decomposition, *Chaos* **32**, 033116.
- R. J. Allen, D. Frenkel and P. R. Ten Wolde (2006), Forward flux sampling-type schemes for simulating rare events: Efficiency analysis, *J. Chem. Phys.* **124**, 194111.
- R. J. Allen, C. Valeriani and P. R. Ten Wolde (2009), Forward flux sampling for rare event simulations, *J. Phys. Condensed Matter* **21**, 463102.
- M. R. Allshouse and T. Peacock (2015), Lagrangian based methods for coherent structure detection, *Chaos* **25**, 097617.
- G. Andrew, R. Arora, J. Bilmes and K. Livescu (2013), Deep canonical correlation analysis, in *Proceedings of the 30th International Conference on Machine Learning*, Vol. 28 of Proceedings of Machine Learning Research, PMLR, pp. 1247–1255.
- G. Arampatzis, M. A. Katsoulakis and L. Rey-Bellet (2016), Efficient estimators for likelihood ratio sensitivity indices of complex stochastic dynamics, *J. Chem. Phys.* **144**, 104107.
- D. Aristoff, T. Lelièvre, C. G. Mayne and I. Teo (2015), Adaptive multilevel splitting in molecular dynamics simulations, *ESAIM Proc. Surveys* **48**, 215–225.
- S. Asmussen, P. Dupuis, R. Y. Rubinstein and H. Wang (2013), Rare event simulation, in *Encyclopedia of Operations Research and Management Science* (S. I. Gass and M. C. Fu, eds), Springer, pp. 1264–1279.
- A. M. Avila and I. Mezić (2020), Data-driven analysis and forecasting of highway traffic dynamics, *Nature Commun.* **11**, 2090.
- H. P. Awad, P. W. Glynn and R. Y. Rubinstein (2013), Zero-variance importance sampling estimators for Markov process expectations, *Math. Oper. Res.* **38**, 358–388.
- C. Ayaz, L. Scalfi, B. A. Dalton and R. R. Netz (2022), Generalized Langevin equation with a nonlinear potential of mean force and nonlinear memory friction from a hybrid projection scheme, *Phys. Rev. E* **105**, 054138.
- C. Ayaz, L. Tepper, F. N. Brünic, J. Kappler, J. O. Daldrop and R. R. Netz (2021), Non-Markovian modeling of protein folding, *Proc. Nat. Acad. Sci. USA* **118**, e2023856118.
- A. Bachouch, C. Huré, N. Langrené and H. Pham (2022), Deep neural networks algorithms for stochastic control problems on finite horizon: Numerical applications, *Methodol. Comput. Appl. Probab.* **24**, 143–178.

- T. Badowski (2016), Adaptive importance sampling via minimization of estimators of cross-entropy, mean square, and inefficiency constant. Doctoral thesis, Freie Universität Berlin.
- R. Banisch and C. Hartmann (2016), A sparse Markov chain approximation of LQ-type stochastic control problems, *Math. Control Related Fields* **6**, 363–389.
- R. Banisch and P. Koltai (2017), Understanding the geometry of transport: Diffusion maps for Lagrangian trajectory data unravel coherent sets, *Chaos* **27**, 035804.
- R. Banisch, N. D. Conrad and C. Schütte (2015), Reactive flows and unproductive cycles for random walks on complex networks, *Eur. Phys. J. Spec. Top.* **224**, 2369–2387.
- C. Bartels and M. Karplus (1997), Multidimensional adaptive umbrella sampling: Applications to main chain and side chain peptide conformations, *J. Comput. Chem.* **18**, 1450–1462.
- K. A. Beauchamp, G. R. Bowman, T. J. Lane, L. Maibaum, I. S. Haque and V. S. Pande (2011), MSMBuilder2: Modeling conformational dynamics at the picosecond to millisecond scale, *J. Chem. Theory Comput.* **7**, 3412–3419.
- S. Beccara, T. Skrbic, R. Covino and P. Faccioli (2012), Dominant folding pathways of a WW domain, *Proc. Nat. Acad. Sci. USA* **109**, 2330–2335.
- C. Beck, A. Jentzen *et al.* (2019), Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations, *J. Nonlinear Sci.* **29**, 1563–1619.
- Z. Belkacemi, P. Gkeka, T. Lelièvre and G. Stoltz (2022), Chasing collective variables using autoencoders and biased trajectories, *J. Chem. Theory Comput.* **18**, 59–78.
- J. Bello-Rivas and R. Elber (2015), Exact milestoning, *J. Chem. Phys.* **142**, 094102.
- J. Beltran and C. Landim (2010), Tunneling and metastability of continuous time Markov chains, *J. Statist. Phys.* **140**, 1065–1114.
- J. Beltran and C. Landim (2013), A martingale approach to metastability, *Probab. Theory Related Fields* **161**, 267–307.
- C. Bender and T. Moseler (2010), Importance sampling for backward SDEs, *Stoch. Anal. Appl.* **28**, 226–253.
- C. Bender and J. Steiner (2012), Least-squares Monte Carlo for backward SDEs, in *Numerical Methods in Finance* (R. A. Carmona *et al.*, eds), Springer, pp. 257–289.
- T. Bengtsson, P. Bickel and B. Li (2008), 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, pp. 316–334.
- A. M. Berezhkovskii and A. Szabo (2019), Committors, first-passage times, fluxes, Markov states, milestones, and all that, *J. Chem. Phys.* **150**, 054106.
- N. Berglund (2013), Kramers' law: Validity, derivations and generalisations, *Markov Process. Related Fields* **19**, 459–490.
- S. Bhakat (2022), Collective variable discovery in the age of machine learning: Reality, hype and everything in between, *RSC Adv.* **12**, 25010–25024.
- A. Bianchi and A. Gaudillièrè (2016), Metastable states, quasi-stationary distributions and soft measures, *Stoch. Process. Appl.* **126**, 1622–1680.
- J. Birrell and L. Rey-Bellet (2020), Uncertainty quantification for Markov processes via variational principles and functional inequalities, *SIAM/ASA J. Uncertain. Quantif.* **8**, 539–572.
- J. Birrell, M. A. Katsoulakis and L. Rey-Bellet (2021), Quantification of model uncertainty on path-space via goal-oriented relative entropy, *ESAIM Math. Model. Numer. Anal.* **55**, 131–169.

- A. Bittracher and C. Schütte (2020), A weak characterization of slow variables in stochastic dynamical systems, in *Advances in Dynamics, Optimization and Computation (SON 2020)* (O. Junge *et al.*, eds), Springer, pp. 132–150.
- A. Bittracher, P. Koltai, S. Klus, R. Banisch, M. Dellnitz and C. Schütte (2018), Transition manifolds of complex metastable systems, *J. Nonlinear Sci.* **28**, 471–512.
- A. Bittracher, M. Mollenhauer, P. Koltai and C. Schütte (2021), Optimal reaction coordinates: Variational characterization and sparse computation. Available at [arXiv:2107.10158](https://arxiv.org/abs/2107.10158) (to appear in *SIAM J. Multiscale Model. Simul.*).
- P. G. Bolhuis and D. W. H. Swenson (2021), Transition path sampling as Markov chain Monte Carlo of trajectories: Recent algorithms, software, applications, and future outlook, *Adv. Theory Simul.* **4**, 2000237.
- P. G. Bolhuis, D. Chandler, C. Dellago and P. Geissler (2002), Transition path sampling: Throwing ropes over mountain passes, in the dark, *Annu. Rev. Phys. Chem.* **59**, 291–318.
- L. Bonati, G. Piccini and M. Parrinello (2021), Deep learning the slow modes for rare events sampling, *Proc. Nat. Acad. Sci. USA* **118**, e2113533118.
- S. D. Bond, B. B. L. Benedict and J. Leimkuhler (1999), The Nosé–Poincaré method for constant temperature molecular dynamics, *J. Comput. Phys.* **151**, 114–134.
- E. R. Borrell, J. Quer, L. Richter and C. Schütte (2022), Improving control based importance sampling strategies for metastable diffusions via adapted metadynamics. Available at [arXiv:2206.06628](https://arxiv.org/abs/2206.06628).
- N. Bou-Rabee and E. Vanden-Eijnden (2010), Pathwise accuracy and ergodicity of Metropolized integrators for SDEs, *Commun. Pure Appl. Math.* **63**, 655–696.
- M. Boué and P. Dupuis (1998), A variational representation for certain functionals of Brownian motion, *Ann. Probab.* **26**, 1641–1659.
- A. Bovier and F. Den Hollander (2016), *Metastability: A Potential-Theoretic Approach*, Vol. 351 of Grundlehren der mathematischen Wissenschaften, Springer.
- A. Bovier, M. Eckhoff, V. Gayrard and M. Klein (2002a), Metastability and low lying spectra in reversible Markov chains, *Comm. Math. Phys.* **228**, 219–255.
- A. Bovier, M. Eckhoff, V. Gayrard and M. Klein (2002b), Metastability in reversible diffusion processes I: Sharp asymptotics for capacities and exit times, *J. Eur. Math. Soc.* **6**, 399–424.
- A. Bovier, V. Gayrard and M. Klein (2002c), Metastability in reversible diffusion processes II: Precise asymptotics for small eigenvalues, *J. Eur. Math. Soc.* **7**, 69–99.
- G. R. Bowman, V. S. Pande and F. Noé, eds (2014), *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation*, Vol. 797 of Advances in Experimental Medicine and Biology, Springer.
- G. R. Bowman, V. Volez and V. S. Pande (2011), Taming the complexity of protein folding, *Curr. Opin. Struct. Biol.* **21**, 4–11.
- C. L. Bris, T. Lelièvre, M. Luskin and D. Perez (2012), A mathematical formalization of the parallel replica dynamics, *Monte Carlo Methods Appl.* **18**, 119–146.
- S. L. Brunton, J. L. Proctor and J. N. Kutz (2016), Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proc. Nat. Acad. Sci. USA* **113**, 3932–3937.
- G. Bussi, A. Laio and P. Tiwary (2020), Metadynamics: A unified framework for accelerating rare events and sampling thermodynamics and kinetics, in *Handbook of Materials Modeling – Methods: Theory and Modeling* (W. Andreoni and S. Yip, eds), Springer, pp. 565–595.

- J. D. Carroll and J. J. Chang (1970), Analysis of individual differences in multidimensional scaling via an N-way generalization of 'Eckart–Young' decomposition, *Psychometrika* **35**, 283–319.
- F. Cérou and A. Guyader (2007), Adaptive multilevel splitting for rare event analysis, *Stoch. Anal. Appl.* **25**, 417–443.
- F. Cérou, P. Del Moral, T. Furon and A. Guyader (2012), Sequential Monte Carlo for rare event estimation, *Statist. Comput.* **22**, 795–808.
- D. Chandler (1998), *Finding Transition Pathways: Throwing Ropes Over Rough Mountain Passes, in the Dark*, World Scientific.
- J. D. Chodera, W. C. Swope, J. W. Pitera and K. A. Dill (2006), Long-time protein folding dynamics from short-time molecular dynamics simulations, *Multiscale Model. Simul.* **5**, 1214–1226.
- A. J. Chorin, O. H. Hald and R. Kupferman (2000), Optimal prediction and the Mori–Zwanzig representation of irreversible processes, *Proc. Nat. Acad. Sci. USA* **97**, 2968–2973.
- J. Comer, J. C. Gumbart, J. Hénin, T. Lelièvre, A. Pohorille and C. Chipot (2015), The adaptive biasing force method: Everything you always wanted to know but were afraid to ask, *J. Phys. Chem. B* **119**, 1129–1151.
- N. D. Conrad, M. Sarich and C. Schütte (2012), Estimating the eigenvalue error of Markov state models, *Multiscale Model. Simul.* **10**, 61–81.
- N. D. Conrad, M. Weber and C. Schütte (2015), Finding dominant structures of nonreversible Markov processes, *SIAM J. Mult. Model. Simul.* **14**, 1319–1340.
- G. Crooks (1999), Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences, *Phys. Rev. E* **60**, 2721–2726.
- P. Dai Pra, L. Meneghini and W. Runggaldier (1996), Connections between stochastic control and dynamic games, *Math. Control Signals Systems* **9**, 303–326.
- E. B. Davies (1982a), Metastable states of symmetric Markov semigroups I, *Proc. London Math. Soc.* **s3-45**, 133–150.
- E. B. Davies (1982b), Metastable states of symmetric Markov semigroups II, *J. London Math. Soc.* **s2-26**, 541–556.
- C. Davis and W. M. Kahan (1970), The rotation of eigenvectors by a perturbation III, *SIAM J. Numer. Anal.* **7**, 1–46.
- P. Del Moral (2004), *Feynman–Kac Formulae: Genealogical and Interacting Particle Systems with Applications*, Probability and its Applications, Springer.
- F. Delbaen, Y. Hu and A. Richou (2011), On the uniqueness of solutions to quadratic BSDEs with convex generators and unbounded terminal conditions, *Ann. Inst. H. Poincaré Probab. Statist.* **47**, 559–574.
- M. Dellnitz and O. Junge (1998), An adaptive subdivision technique for the approximation of attractors and invariant measures, *Comput. Vis. Sci.* **1**, 63–68.
- M. Dellnitz and O. Junge (1999), On the approximation of complicated dynamical behavior, *SIAM J. Numer. Anal.* **36**, 491–515.
- P. Deuffhard and M. Weber (2005), Robust Perron cluster analysis in conformation dynamics, *Linear Algebra Appl.* **398**, 161–184.
- P. Deuffhard, M. Dellnitz, O. Junge and C. Schütte (1999), Computation of essential molecular dynamics by subdivision techniques, in *Computational Molecular Dynamics: Challenges, Methods, Ideas*, Vol. 4 of Lecture Notes in Computational Science and Engineering, Springer, pp. 98–115.

- G. Di Gesu, T. Lelièvre, D. Peutrec and B. Nectoux (2016), Jump Markov models and transition state theory: The quasi-stationary distribution approach, *Faraday Discuss.* **195**, 469–495.
- L. Donati, M. Heida, B. G. Keller and M. Weber (2018), Estimation of the infinitesimal generator by square-root approximation, *J. Phys. Condensed Matter* **30**, 425201.
- L. Donati, M. Weber and B. G. Keller (2021), Markov models from the square root approximation of the Fokker–Planck equation: Calculating the grid-dependent flux, *J. Phys. Condensed Matter* **33**, 115902.
- M. D. Donsker and S. R. S. Varadhan (1975), On a variational formula for the principal eigenvalue for operators with maximum principle, *Proc. Nat. Acad. Sci. USA* **72**, 780–783.
- J. L. Doob (1984), *Classical Potential Theory and its Probabilistic Counterpart*, Vol. 262 of Grundlehren der Mathematischen Wissenschaften, Springer.
- D. Down, S. P. Meyn and R. L. Tweedie (1995), Exponential and uniform ergodicity of Markov processes, *Ann. Probab.* **23**, 1671–1691.
- P. Dupuis and H. Wang (2004), Importance sampling, large deviations, and differential games, *Stochastics* **76**, 481–508.
- P. Dupuis and H. Wang (2007), Subsolutions of an Isaacs equation and efficient schemes for importance sampling, *Math. Oper. Res.* **32**, 723–757.
- P. Dupuis, M. A. Katsoulakis, Y. Pantazis and P. Plecháč (2016), Path-space information bounds for uncertainty quantification and sensitivity analysis of stochastic dynamics, *SIAM/ASA J. Uncertain. Quantif.* **4**, 80–111.
- P. Dupuis, M. A. Katsoulakis, Y. Pantazis and L. Rey-Bellet (2020), Sensitivity analysis for rare events based on Rényi divergence, *Ann. Appl. Probab.* **30**, 1507–1533.
- P. Dupuis, Y. Liu, N. Plattner and J. D. Doll (2012), On the infinite swapping limit for parallel tempering, *Multiscale Model. Simul.* **10**, 986–1022.
- P. Dupuis, K. Spiliopoulos and X. Zhou (2015), Escaping from an attractor: Importance sampling and rest points I, *Ann. Appl. Probab.* **25**, 2909–2958.
- W. E and E. Vanden-Eijnden (2004), Metastability, conformation dynamics, and transition pathways in complex systems, in *Multiscale Modelling and Simulation*, Vol. 39 of Lecture Notes in Computational Science and Engineering, Springer, pp. 35–68.
- W. E and E. Vanden-Eijnden (2006), Towards a theory of transition paths, *J. Statist. Phys.* **123**, 503–523.
- W. E and E. Vanden-Eijnden (2010), Transition-path theory and path-finding algorithms for the study of rare events, *Annu. Rev. Phys. Chem.* **61**, 391–420.
- W. E, J. Han and A. Jentzen (2017), Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations, *Commun. Math. Statist.* **5**, 349–380.
- W. E, J. Han and A. Jentzen (2021), Algorithms for solving high dimensional PDEs: From nonlinear Monte Carlo to machine learning, *Nonlinearity* **35**, 278–310.
- D. J. Earl and M. W. Deem (2005), Parallel tempering: Theory, applications, and new perspectives, *Phys. Chem. Chem. Phys.* **7**, 3910–3916.
- R. S. Ellis (1985), *Entropy, Large Deviations and Statistical Mechanics*, Springer.
- H. Eyring (1935), The activated complex in chemical reactions, *J. Chem. Phys.* **3**, 107–115.
- P. Faccioli, A. Lonardi and H. Orland (2010), Dominant reaction pathways in protein folding: A direct validation against molecular dynamics simulations, *J. Chem. Phys.* **133**, 045104.

- A. K. Faradjian and R. Elber (2004), Computing time scales from reaction coordinates by milestoning, *J. Chem. Phys.* **120**, 10880–10889.
- H. Federer (1969), *Geometric Measure Theory*, Springer.
- W. H. Fleming (2006), Risk sensitive stochastic control and differential games, *Commun. Inf. Syst.* **6**, 161–177.
- W. H. Fleming and H. M. Soner (2006), *Controlled Markov Processes and Viscosity Solutions*, Springer.
- A.-S. Frank, A. Sikorski and S. Röblitz (2022), Spectral clustering of Markov chain transition matrices with complex eigenvalues. Available at [arXiv:2206.14537](https://arxiv.org/abs/2206.14537).
- M. Freidlin and A. D. Wentzell (1998), *Random Perturbations of Dynamical Systems*, Springer.
- G. Froyland (2013), An analytic framework for identifying finite-time coherent sets in time-dependent dynamical systems, *Phys. D* **250**, 1–19.
- G. Froyland and O. Junge (2018), Robust FEM-based extraction of finite-time coherent sets using scattered, sparse, and incomplete trajectories, *SIAM J. Appl. Dyn. Syst.* **17**, 1891–1924.
- G. Froyland, G. Gottwald and A. Hammerlindl (2013), A computational method to extract macroscopic variables and their dynamics in multiscale systems, *SIAM J. Appl. Dyn. Syst.* **13**, 1816–1846.
- Y. Ge and V. A. Voelz (2021), Markov state models to elucidate ligand binding mechanism, *Methods Mol. Biol.* **2266**, 239–259.
- P. Geiß (2017), The tensor-train format and its applications: Modeling and analysis of chemical reaction networks, catalytic processes, fluid flows, and Brownian dynamics. Doctoral thesis, Freie Universität Berlin.
- E. Gobet and P. Turkedjiev (2016), Linear regression MDP scheme for discrete backward stochastic differential equations under general conditions, *Math. Comp.* **85**, 1359–1391.
- E. Gobet, J.-P. Lemor and X. Warin (2005), A regression-based Monte Carlo method to solve backward stochastic differential equations, *Ann. Appl. Probab.* **15**, 2172–2202.
- H. Grabert (1982), *Projection Operator Techniques in Nonequilibrium Statistical Mechanics*, Springer Tracts in Modern Physics, Springer.
- C. Graham and D. Talay (2013), *Stochastic Simulation and Monte Carlo Methods*, Springer.
- H. Grubmüller (1995), Predicting slow structural transitions in macromolecular systems: Conformational flooding, *Phys. Rev. E* **52**, 2893.
- I. Gyöngy and T. Martínez (2001), On stochastic differential equations with locally unbounded drift, *Czechoslovak Math. J.* **51**, 763–783.
- W. Hackbusch (2014), Numerical tensor calculus, *Acta Numer.* **23**, 651–742.
- D. Hamelberg, J. Mongan and J. A. McCammon (2004), Accelerated molecular dynamics: A promising and efficient simulation method for biomolecules, *J. Chem. Phys.* **120**, 11919.
- J. Han, J. Lu and M. Zhou (2020), Solving high-dimensional eigenvalue problems using deep neural networks: A diffusion Monte Carlo like approach, *J. Comput. Phys.* **423**, 109792.
- C. Hartmann and L. Richter (2021), Nonasymptotic bounds for suboptimal importance sampling. Available at [arXiv:2102.09606](https://arxiv.org/abs/2102.09606).
- C. Hartmann and C. Schütte (2012), Efficient rare event simulation by optimal nonequilibrium forcing, *J. Statist. Mech. Theory Exp.* **2012**, P11004.

- C. Hartmann, O. Kebiri, L. Neureither and L. Richter (2019), Variational approach to rare event simulation using least-squares regression, *Chaos* **29**, 063107.
- C. Hartmann, L. Neureither and U. Sharma (2020), Coarse graining of nonreversible stochastic differential equations: Quantitative results and connections to averaging, *SIAM J. Math. Anal.* **52**, 2689–2733.
- C. Hartmann, L. Richter, C. Schütte and W. Zhang (2017), Variational characterization of free energy: Theory and algorithms, *Entropy* **19**, 626.
- C. Hartmann, C. Schütte and W. Zhang (2016), Model reduction algorithms for optimal control and importance sampling of diffusions, *Nonlinearity* **29**, 2298.
- M. Heida (2018), Convergences of the squareroot approximation scheme to the Fokker–Planck operator, *Math. Models Methods Appl. Sci.* **28**, 2599–2635.
- F. Hérou, M. Hitrik and J. Sjöstrand (2008), Tunnel effect for Kramers–Fokker–Planck type operators: Return to equilibrium and applications, *Int. Math. Res. Not.* **2008**, rnn057.
- F. Hérou, M. Hitrik and J. Sjöstrand (2010), Tunnel effect and symmetries for Kramers–Fokker–Planck type operators. Available at [arXiv:1007.0838v1 \[math.SP\]](https://arxiv.org/abs/1007.0838v1).
- C. Hijón, P. Español, E. Vanden-Eijnden and R. Delgado-Buscalioni (2010), Mori–Zwanzig formalism as a practical computational tool, *Faraday Discuss.* **144**, 301–322.
- M. Hoffmann, M. Scherer, T. Hempel, A. Mardt, B. de Silva, B. E. Husic, S. Klus, H. Wu, N. Kutz, S. L. Brunton and F. Noé (2021), Deeptime: A Python library for machine learning dynamical models from time series data, *Mach. Learn. Sci. Technol.* **3**, 015009.
- H. Hotelling (1936), Relations between two sets of variates, *Biometrika* **28**, 321–377.
- J. Hua, F. Noorian, D. Moss, P. H. W. Leong and G. H. Gunaratne (2017), High-dimensional time series prediction using kernel-based Koopman mode regression, *Nonlinear Dynam.* **90**, 1785–1806.
- X. Huang, Y. Yao, G. Bowman, J. Sun, L. J. Guibas, G. Carlsson and V. Pande (2010), Constructing multi-resolution Markov state models (MSMs) to elucidate RNA hairpin folding mechanisms, in *Pacific Symposium on Biocomputing 2010 (PSB 2010)* (R. B. Altman *et al.*, eds), World Scientific, pp. 228–239.
- W. Huisinga (2001), Metastability of Markovian systems: A transfer operator approach in application to molecular dynamics. Doctoral thesis, Freie Universität Berlin.
- W. Huisinga and B. Schmidt (2002), Metastability and dominant eigenvalues of transfer operators, in *Advances in Algorithms for Macromolecular Simulation* (C. Chipot *et al.*, eds), Vol. 49 of Lecture Notes in Computational Science and Engineering, Springer.
- W. Huisinga, S. Meyn and C. Schütte (2004), Phase transitions & metastability in Markovian and molecular systems, *Ann. Appl. Probab.* **14**, 419–458.
- C. Huré, H. Pham, A. Bachouch and N. Langrené (2021), Deep neural networks algorithms for stochastic control problems on finite horizon: Convergence analysis, *SIAM J. Numer. Anal.* **59**, 525–557.
- B. E. Husic and V. S. Pande (2018), Markov state models: From an art to a science, *J. Amer. Chem. Soc.* **140**, 2386.
- S. Hussain and A. Haji-Akbari (2020), Studying rare events using forward-flux sampling: Recent breakthroughs and future outlook, *J. Chem. Phys.* **152**, 060901.
- M. Invernizzi and M. Parrinello (2020), Rethinking metadynamics: From bias potentials to probability distributions, *J. Phys. Chem. Lett.* **11**, 2731–2736.
- C. Jarzynski (1997), Nonequilibrium equality for free energy differences, *Phys. Rev. Lett.* **78**, 2690–2693.

- S. Juneja and P. Shahabuddin (2006), Rare-event simulation techniques: An introduction and recent advances, in *Simulation* (S. G. Henderson and B. L. Nelson, eds), Vol. 13 of Handbooks in Operations Research and Management Science, Elsevier, pp. 291–350.
- J. Kappler, J. O. Daldrop, F. N. Brüning, M. D. Boehle and R. Netz (2018), Memory-induced acceleration and slowdown of barrier crossing, *J. Chem. Phys.* **148**, 014903.
- I. Karatzas and S. E. Shreve (1991), *Brownian Motion and Stochastic Calculus*, Graduate Texts in Mathematics, Springer.
- T. Kato (1995), *Perturbation Theory for Linear Operators*, Springer.
- O. Kebiri, L. Neureither and C. Hartmann (2019), Adaptive importance sampling with forward–backward stochastic differential equations, in *Stochastic Dynamics Out of Equilibrium (IHPStochDyn 2017)* (G. Giacomini *et al.*, eds), Vol. 282 of Proceedings in Mathematics & Statistics, Springer, pp. 265–281.
- B. Keller, J.-H. Prinz and F. Noé (2011), Markov models and dynamical fingerprints: Unraveling the complexity of molecular kinetics, *Chem. Phys.* **396**, 92–107.
- I. G. Kevrekidis and G. Samaey (2009), Equation-free multiscale computation: Algorithms and applications, *Annu. Rev. Phys. Chem.* **60**, 321–344.
- S. Klus and N. D. Conrad (2022), Koopman-based spectral clustering of directed and time-evolving graphs, *J. Nonlinear Sci.* **33**, 8.
- S. Klus, A. Bittracher, I. Schuster and C. Schütte (2018a), A kernel-based approach to molecular conformation analysis, *J. Chem. Phys.* **149**, 244109.
- S. Klus, B. E. Husic, M. Mollenhauer and F. Noé (2019a), Kernel methods for detecting coherent structures in dynamical data, *Chaos* **29**, 123112.
- S. Klus, P. Koltai and C. Schütte (2016), On the numerical approximation of the Perron–Frobenius and Koopman operator, *J. Comput. Dyn.* **3**, 51–79.
- S. Klus, F. Nüske and B. Hamzi (2020a), Kernel-based approximation of the Koopman generator and Schrödinger operator, *Entropy* **22**, 722.
- S. Klus, F. Nüske and S. Peitz (2022), Koopman analysis of quantum systems, *J. Phys. A Math. Theor.* **55**, 314002.
- S. Klus, F. Nüske, P. Koltai, H. Wu, I. Kevrekidis, C. Schütte and F. Noé (2018b), Data-driven model reduction and transfer operator approximation, *J. Nonlinear Sci.* **28**, 985–1010.
- S. Klus, F. Nüske, S. Peitz, J.-H. Niemann, C. Clementi and C. Schütte (2020b), Data-driven approximation of the Koopman generator: Model reduction, system identification, and control, *Phys. D* **406**, 132416.
- S. Klus, I. Schuster and K. Muandet (2019b), Eigendecompositions of transfer operators in reproducing kernel Hilbert spaces, *J. Nonlinear Sci.* **30**, 283–315.
- M. Kobylanski (2000), Backward stochastic differential equations and partial differential equations with quadratic growth, *Ann. Probab.* **28**, 558–602.
- P. Koltai, H. Wu, F. Noé and C. Schütte (2018), Optimal data-driven estimation of generalized Markov state models for non-equilibrium dynamics, *Computation* **6**, 22.
- I. Kontoyiannis and S. P. Meyn (2003), Spectral theory and limit theorems for geometrically ergodic Markov processes, *Ann. Appl. Probab.* **13**, 304–362.
- I. Kontoyiannis and S. P. Meyn (2012), Geometric ergodicity and the spectral gap of non-reversible Markov chains, *Probab. Theory Related Fields* **154**, 327–339.
- M. Korda and I. Mezić (2018a), Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control, *Automatica* **93**, 149–160.

- M. Korda and I. Mezic (2018*b*), On convergence of extended dynamic mode decomposition to the Koopman operator, *J. Nonlinear Sci.* **28**, 687–710.
- D. Kressner (2006), Block algorithms for reordering standard and generalized Schur forms, *ACM Trans. Math. Softw.* **32**, 521–532.
- S. V. Krivov (2018), Protein folding free energy landscape along the committer: The optimal folding coordinate, *J. Chem. Theory Comput.* **14**, 3418–3427.
- N. V. Krylov (2019), A few comments on a result of A. Novikov and Girsanov’s theorem, *Stochastics* **91**, 1186–1189.
- O. F. Lange and H. Grubmüller (2006), Collective Langevin dynamics of conformational motions in proteins, *J. Chem. Phys.* **124**, 214903.
- A. Lasota and M. C. Mackey (1994), *Chaos, Fractals and Noise*, Vol. 97 of Applied Mathematical Sciences, second edition, Springer.
- J. Latorre, P. Metzner, C. Hartmann and C. Schütte (2011), A structure-preserving numerical discretization of reversible diffusions, *Commun. Math. Sci.* **9**, 1051–1072.
- P. L’Ecuyer, M. Mandjes and B. Tuffin (2009), *Importance Sampling in Rare Event Simulation*, Wiley, pp. 17–38.
- F. Legoll and T. Lelièvre (2010), Effective dynamics using conditional expectations, *Nonlinearity* **23**, 2131.
- B. Li, T. Bengtsson and P. Bickel (2005), Curse-of-dimensionality revisited: Collapse of importance sampling in very high-dimensional systems. Technical report 696, Department of Statistics, UC Berkeley.
- H. C. Lie (2016), On a strongly convex approximation of a stochastic optimal control problem for importance sampling of metastable diffusions. Doctoral thesis, Freie Universität Berlin.
- H. C. Lie (2021), Fréchet derivatives of expected functionals of solutions to stochastic differential equations. Available at [arXiv:2106.09149](https://arxiv.org/abs/2106.09149).
- H. C. Lie and J. Quer (2017), Some connections between importance sampling and enhanced sampling methods in molecular dynamics, *J. Chem. Phys.* **147**, 194107.
- J. Lu and E. Vanden-Eijnden (2014), Exact dynamical coarse-graining without time-scale separation, *J. Chem. Phys.* **141**, 044109.
- M. Lücke and F. Nüske (2022), tgEDMD: Approximation of the Kolmogorov operator in tensor train format, *J. Nonlinear Sci.* **32**, 44.
- A. Mardt, L. Pasquali, H. Wu and F. Noé (2018), VAMPnets for deep learning of molecular kinetics, *Nature Commun.* **9**, 5.
- E. Marinari and G. Parisi (1992), Simulated tempering: A new Monte Carlo scheme, *Europhys. Lett.* **19**, 451.
- N. Marrouch, J. Slawinska, D. Giannakis and H. L. Read (2020), Data-driven Koopman operator approach for computational neuroscience, *Ann. Math. Artif. Intell.* **88**, 1155–1173.
- A. Martinsson, J. Lu, B. Leimkuhler and E. Vanden-Eijnden (2019), The simulated tempering method in the infinite switch limit with adaptive weight learning, *J. Statist. Mech. Theory Exp.* **2019**, 013207.
- J. C. Mattingly, A. M. Stuart and D. J. Higham (2002), Ergodicity for SDEs and approximations: Locally Lipschitz vector fields and degenerate noise, *Stoch. Process. Appl.* **101**, 185–232.
- A. Mauroy and J. Goncalves (2016), Linear identification of nonlinear systems: A lifting technique based on the Koopman operator, in *2016 IEEE 55th Conference on Decision and Control (CDC)*, IEEE, pp. 6500–6505.

- T. Melzer, M. Reiter and H. Bischof (2001), Nonlinear feature extraction using generalized canonical correlation analysis, in *Artificial Neural Networks (ICANN 2001)* (G. Dorffner *et al.*, eds), Springer, pp. 353–360.
- P. Metzner (2007), Transition path theory for Markov processes: Application to molecular dynamics. Doctoral thesis, Freie Universität Berlin.
- P. Metzner, F. Noé and C. Schütte (2009a), Estimating the sampling error: Distribution of transition matrices and functions of transition matrices for given trajectory data, *Phys. Rev. E* **80**, 021106.
- P. Metzner, C. Schütte and E. Vanden-Eijnden (2006), Illustration of transition path theory on a collection of simple examples, *J. Chem. Phys.* **125**, 084110.
- P. Metzner, C. Schütte and E. Vanden-Eijnden (2009b), Transition path theory for Markov jump processes, *Multiscale Model. Simul.* **7**, 1192–1219.
- P. Metzner, M. Weber and C. Schütte (2010), Observation uncertainty in reversible Markov chains, *Phys. Rev. E* **82**, 031114.
- S. Meyn and R. Tweedie (1993), *Markov Chains and Stochastic Stability*, Springer.
- L. Molgedey and H. G. Schuster (1994), Separation of a mixture of independent signals using time delayed correlations, *Phys. Rev. Lett.* **72**, 3634–3637.
- M. Mollenhauer (2022), On the statistical approximation of conditional expectation operators. Doctoral thesis, Freie Universität Berlin.
- M. Mollenhauer and P. Koltai (2020), Nonparametric approximation of conditional expectation operators. Available at [arXiv:2012.12917](https://arxiv.org/abs/2012.12917).
- M. Mollenhauer, N. Mücke and T. J. Sullivan (2022), Learning linear operators: Infinite-dimensional regression as a well-behaved non-compact inverse problem. Available at [arXiv:2211.08875](https://arxiv.org/abs/2211.08875).
- H. Mori (1965), Transport, collective motion, and Brownian motion, *Prog. Theor. Phys.* **33**, 423–455.
- D. Moroni, T. van Erp and P. Bolhuis (2004), Investigating rare events by transition interface sampling, *Phys. A* **340**, 395–401.
- F. Noé (2008), Probability distributions of molecular observables computed from Markov models, *J. Chem. Phys.* **128**, 244103.
- F. Noé and F. Nüske (2013), A variational approach to modeling slow processes in stochastic dynamical systems, *Multiscale Model. Simul.* **11**, 635–655.
- F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich and T. Weigl (2009), Constructing the full ensemble of folding pathways from short off-equilibrium trajectories, *Proc. Nat. Acad. Sci. USA* **106**, 19011–19016.
- F. Nüske, P. Gelß, S. Klus and C. Clementi (2021), Tensor-based computation of metastable and coherent sets, *Phys. D* **427**, 133018.
- F. Nüske, R. Schneider, F. Vitalini and F. Noé (2016), Variational tensor approach for approximating the rare-event kinetics of macromolecular systems, *J. Chem. Phys.* **144**, 054105.
- F. Nüske, H. Wu, J.-H. Prinz, C. Wehmeyer, C. Clementi and F. Noé (2017), Markov state models from short non-equilibrium simulations: Analysis and correction of estimation bias, *J. Chem. Phys.* **146**, 094104.
- N. Nüsken and L. Richter (2021), Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: Perspectives from the theory of controlled diffusions and measures on path space, *Partial Differ. Equ. Appl.* **2**, 48.

- B. Øksendal (2003), *Stochastic Differential Equations: An Introduction with Applications*, Springer.
- E. Olivieri and M. E. Vares (2005), *Large Deviations and Metastability*, Encyclopedia of Mathematics and its Applications, Cambridge University Press.
- I. Oseledets (2011), Tensor-train decomposition, *SIAM J. Sci. Comput.* **33**, 2295–2317.
- I. Oseledets and E. Tyrtyshnikov (2010), TT-cross approximation for multidimensional arrays, *Linear Algebra Appl.* **432**, 70–88.
- V. Pande, K. Beauchamp and G. Bowman (2010), Everything you wanted to know about Markov state models but were afraid to ask, *Methods* **52**, 99–105.
- E. Pardoux and S. Peng (1990), Adapted solution of a backward stochastic differential equation, *Systems Control Lett.* **14**, 55–61.
- E. Pardoux and S. Tang (1999), Forward–backward stochastic differential equations and quasilinear parabolic PDEs, *Probab. Theory Related Fields* **114**, 123–150.
- F. Paul, C. Wehmeyer, E. T. Abualrous, H. Wu, M. D. Crabtree, J. Schöneberg, J. Clarke, C. Freund, T. R. Weikl and F. Noé (2017), Protein–peptide association kinetics beyond the seconds timescale from atomistic simulations, *Nature Commun.* **8**, 1095.
- S. Peitz and S. Klus (2019), Koopman operator-based model reduction for switched-system control of PDEs, *Automatica* **106**, 184–191.
- G. Perez-Hernandez, F. Paul, T. Giorgino, G. De Fabritiis and F. Noé (2013), Identification of slow molecular order parameters for Markov model construction, *J. Chem. Phys.* **139**, 015102.
- D. Pfau, J. S. Spencer, A. G. D. G. Matthews and W. M. C. Foulkes (2020), *Ab initio* solution of the many-electron Schrödinger equation with deep neural networks, *Phys. Rev. Res.* **2**, 033429.
- H. Pham (2009), *Continuous-Time Stochastic Control and Optimization with Financial Applications*, Vol. 61 of Stochastic Modelling and Applied Probability, Springer.
- H. Pham, X. Warin and M. Germain (2021), Neural networks-based backward scheme for fully nonlinear PDEs, *SN Partial Differ. Equ. Appl.* **2**, 16.
- G. Pinamonti, J. Zhao, D. E. Condon, F. Paul, F. Noé, D. H. Turner and G. Bussi (2017), Predicting the kinetics of RNA oligonucleotides using Markov state models, *J. Chem. Theory Comput.* **13**, 926–934.
- F. Pinski and A. Stuart (2010), Transition paths in molecules: Gradient descent in path space, *J. Chem. Phys.* **132**, 184104.
- R. G. Pinsky (1985), On the convergence of diffusion processes conditioned to remain in a bounded region for large time to limiting positive recurrent diffusion processes, *Ann. Probab.* **13**, 363–378.
- J. H. Prinz, B. Keller and F. Noé (2011), Probing molecular kinetics with Markov models: Metastable states, transition pathways and spectroscopic observables, *Phys. Chem. Chem. Phys.* **13**, 16912–16927.
- J. Quer, L. Donati, B. G. Keller and M. Weber (2018), An automatic adaptive importance sampling algorithm for molecular dynamics in reaction coordinates, *SIAM J. Sci. Comput.* **40**, A653–A670.
- R. J. Rabben, S. Ray and M. Weber (2020), ISOKANN: Invariant subspaces of Koopman operators learned by a neural network, *J. Chem. Phys.* **153**, 114109.
- W. Ren and E. Vanden-Eijnden (2002), String method for the study of rare events, *Phys. Rev. B* **66**, 052301.

- J. Ribeiro, P. Bravo, Y. Wang and P. Tiwary (2018), Reweighted autoencoded variational Bayes for enhanced sampling (RAVE), *J. Chem. Phys.* **149**, 072301.
- L. Richter (2022), Solving high-dimensional PDEs, approximation of path space measures and importance sampling of diffusions. Doctoral thesis, BTU Cottbus–Senftenberg.
- H. Risken (1996), *The Fokker–Planck Equation*, second edition, Springer.
- S. Röblitz and M. Weber (2013), Fuzzy spectral clustering by PCCA+: Application to Markov state models and data classification, *Adv. Data Anal. Classif.* **7**, 147–179.
- K. Röder and D. J. Wales (2022), The energy landscape perspective: Encoding structure and function for biomolecules, *Front. Molecular Biosci.* **9**, 820792.
- L. C. G. Rogers and D. Williams (2000), *Diffusions, Markov Processes and Martingales*, Vol. 2: *Itô Calculus*, Cambridge University Press.
- B. Roux (2021), String method with swarms-of-trajectories, mean drifts, lag time, and committor, *J. Phys. Chem. A* **125**, 7558–7571.
- B. Roux (2022), Transition rate theory, spectral analysis, and reactive paths, *J. Chem. Phys.* **156**, 134111.
- D. Sanz-Alonso (2018), Importance sampling and necessary sample size: An information theory approach, *SIAM/ASA J. Uncertain. Quantif.* **6**, 867–879.
- M. Sarich (2011), Projected transfer operators. Doctoral thesis, Freie Universität Berlin.
- M. Sarich, F. Noé and C. Schütte (2010), On the approximation quality of Markov state models, *Multiscale Model. Simul.* **8**, 1154–1177.
- M. K. Scherer, B. Trendelkamp-Schroer, F. Paul, G. Pérez-Hernández, M. Hoffmann, N. Plattner, C. Wehmeyer, J.-H. Prinz and F. Noé (2015), PyEMMA 2: A software package for estimation, validation, and analysis of Markov models, *J. Chem. Theory Comput.* **11**, 5525.
- P. J. Schmid (2010), Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5–28.
- B. Schölkopf and A. J. Smola (2001), *Learning with Kernels: Support Vector Machines, Regularization, Optimization and Beyond*, MIT Press.
- C. Schütte (1998), Conformational dynamics: Modelling, theory, algorithm, and application to biomolecules. Habilitation thesis, Freie Universität Berlin.
- C. Schütte and W. Huisinga (2000), On conformational dynamics induced by Langevin processes, in *EQUADIFF 99: International Conference on Differential Equations* (K. Fiedler, Gröger and J. Sprekels, eds), World Scientific, pp. 1247–1262.
- C. Schütte and M. Sarich (2014), *Metastability and Markov State Models in Molecular Dynamics: Modeling, Analysis, Algorithmic Approaches*, Vol. 32 of Courant Lecture Notes, American Mathematical Society.
- C. Schütte, A. Fischer, W. Huisinga and P. Deuffhard (1999), A direct approach to conformational dynamics based on hybrid Monte Carlo, *J. Comput. Phys.* **151**, 146–168.
- C. Schütte, W. Huisinga and S. Meyn (2003), Metastability of diffusion processes, in *IUTAM Symposium on Nonlinear Stochastic Dynamics* (N. S. Namachchivaya and Y. K. Lin, eds), Springer, pp. 71–81.
- C. Schütte, F. Noé, J. Lu, M. Sarich and E. Vanden-Eijnden (2011), Markov state models based on milestoning, *J. Chem. Phys.* **134**, 204105.
- C. Schwantes and V. Pande (2013), Improvements in Markov state model construction reveal many non-native interactions in the folding of NTL9, *J. Chem. Theory Comput.* **9**, 2000–2009.

- C. R. Schwantes and V. S. Pande (2015), Modeling molecular kinetics with tICA and the kernel trick, *J. Chem. Theory Comput.* **11**, 600–608.
- M. Senne, B. Trendelkamp-Schroer, A. Mey, C. Schütte and F. Noé (2012), EMMA: A software package for Markov model building and analysis, *J. Chem. Theory Comput.* **8**, 2223–2238.
- J. Shawe-Taylor and N. Cristianini (2004), *Kernel Methods for Pattern Analysis*, Cambridge University Press.
- H. Sidky, W. Chen and A. L. Ferguson (2020), Machine learning for collective variable discovery and enhanced sampling in biomolecular simulation, *Molecular Phys.* **118**, e1737742.
- A. Sikorski (2015), PCCA+ and its application to spatial time series clustering. Bachelor thesis, Freie Universität Berlin.
- A. Sikorski (2023), Reduced dynamics of high dimensional stochastic systems. Doctoral thesis, Freie Universität Berlin.
- N. Singhal and V. S. Pande (2005), Error analysis in Markovian state models for protein folding, *J. Chem. Phys.* **123**, 204909.
- J. Sirignano and K. Spiliopoulos (2018), DGM: A deep learning algorithm for solving partial differential equations, *J. Comput. Phys.* **375**, 1339–1364.
- I. Steinwart and A. Christmann (2008), *Support Vector Machines*, first edition, Springer.
- E. Suarez, R. P. Wiewiora, C. Wehmeyer, F. Noé, J. D. Chodera and D. M. Zuckerman (2021), What Markov state models can and cannot do: Correlation versus path-based observables in protein-folding models, *J. Chem. Theory Comput.* **17**, 3119–3133.
- R. H. Swendsen and J.-S. Wang (1986), Replica Monte Carlo simulation of spin-glasses, *Phys. Rev. Lett.* **57**, 2607–2609.
- D. W. H. Swenson and P. G. Bolhuis (2014), A replica exchange transition interface sampling method with multiple interface sets for investigating networks of rare events, *J. Chem. Phys.* **141**, 044101.
- W. Tian and H. Wu (2021), Kernel embedding based variational approach for low-dimensional approximation of dynamical systems, *Comput. Methods Appl. Math.* **21**, 635–659.
- A. Tsourtis, Y. Pantazis, M. A. Katsoulakis and V. Harmandaris (2015), Parametric sensitivity analysis for stochastic molecular systems using information theoretic metrics, *J. Chem. Phys.* **143**, 014116.
- L. R. Tucker (1964), The extension of factor analysis to three-dimensional matrices, in *Contributions to Mathematical Psychology* (H. Gulliksen and N. Frederiksen, eds), Holt, Rinehart & Winston, pp. 110–127.
- P. Turkedjiev (2013), Numerical methods for backward stochastic differential equations of quadratic and locally Lipschitz type. Doctoral thesis, Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät II.
- O. Valsson and M. Parrinello (2014), Variational approach to enhanced sampling and free energy calculations, *Phys. Rev. Lett.* **113**, 090601.
- T. S. Van Erp, D. Moroni and P. G. Bolhuis (2003), A novel path sampling method for the calculation of rate constants, *J. Chem. Phys.* **118**, 7762–7774.
- E. Vanden-Eijnden and J. Weare (2012), Rare event simulation of small noise diffusions, *Commun. Pure Appl. Math.* **65**, 1770–1803.
- M. Villén-Altamirano and J. Villén-Altamirano (1994), Restart: A straightforward method for fast simulation of rare events, in *Proceedings of the 26th Conference on Winter Simulation (WSC '94)*, Society for Computer Simulation International, pp. 282–289.

- P. R. Vlachas, J. Zavadlav, M. Praprotnik and P. Koumoutsakos (2022), Accelerated simulations of molecular systems through learning of effective dynamics, *J. Chem. Theory Comput.* **18**, 538–549.
- A. F. Voter (1998), Parallel replica method for dynamics of infrequent events, *Phys. Rev. B* **57**, R13985–R13988.
- D. J. Wales (2003), *Energy Landscapes*, Cambridge University Press.
- D. J. Wales (2005), Energy landscapes and properties of biomolecules, *Phys. Biol.* **2**, S86–S93.
- Y. Wang, J. M. L. Ribeiro and P. Tiwary (2019), Past–future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics, *Nature Commun.* **10**, 3573.
- M. Weber and N. Ernst (2017), A fuzzy-set theoretical framework for computing exit rates of rare events in potential-driven diffusion processes. Available at [arXiv:1708.00679](https://arxiv.org/abs/1708.00679).
- M. Weber and K. Fackeldey (2014), Computing the minimal rebinding effect included in a given kinetics, *Multiscale Model. Simul.* **12**, 318–334.
- C. Wehmeyer and F. Noé (2018), Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics, *J. Chem. Phys.* **148**, 241703.
- A. West, R. Elber and D. Shalloway (2007), Extending molecular dynamics time scales with milestoneing: Example of complex kinetics in a solvated peptide, *J. Chem. Phys.* **126**, 145104.
- E. Wigner (1938), The transition state method, *Trans. Faraday Soc.* **34**, 29–41.
- M. O. Williams, I. G. Kevrekidis and C. W. Rowley (2015a), A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition, *J. Nonlinear Sci.* **25**, 1307–1346.
- M. O. Williams, C. W. Rowley and I. G. Kevrekidis (2015b), A kernel-based method for data-driven Koopman spectral analysis, *J. Comput. Dyn.* **2**, 247–265.
- H. Wu and F. Noé (2020), Variational approach for learning Markov processes from time series data, *J. Nonlinear Sci.* **30**, 23–66.
- H. Wu, F. Nüske, F. Paul, S. Klus, P. Koltai and F. Noé (2017), Variational Koopman models: Slow collective variables and molecular kinetics from short off-equilibrium simulations, *J. Chem. Phys.* **146**, 154104.
- Z. Yang, Y. Zang, H. Wang, Y. Kang, J. Zhang, X. Li, L. Zhang and S. Zhang (2022), Recognition between CD147 and cyclophilin A deciphered by accelerated molecular dynamics simulations, *Phys. Chem. Chem. Phys.* **24**, 18905–18914.
- E. Yeung, S. Kundu and N. Hodas (2019), Learning deep neural network representations for Koopman operators of nonlinear dynamical systems, in *2019 American Control Conference (ACC)*, IEEE, pp. 4832–4839.
- Y. Yu, T. Wang and R. J. Samworth (2015), A useful variant of the Davis–Kahan theorem for statisticians, *Biometrika* **102**, 315–323.
- W. Zhang and C. Schütte (2017), Reliable approximation of long relaxation timescales in molecular dynamics, *Entropy* **19**, 367.
- W. Zhang, C. Hartmann and C. Schütte (2016), Effective dynamics along given reaction coordinates, and reaction rate theory, *Faraday Discuss.* **195**, 365–394.
- W. Zhang, T. Li and C. Schütte (2022), Solving eigenvalue PDEs of metastable diffusion processes using artificial neural networks, *J. Comput. Phys.* **465**, 111377.
- W. Zhang, H. Wang, C. Hartmann, M. Weber and C. Schütte (2014), Applications of the cross-entropy method to importance sampling and optimal control of diffusions, *SIAM J. Sci. Comput.* **36**, A2654–A2672.

- D.-X. Zhou (2008), Derivative reproducing properties for kernel methods in learning theory, *J. Comput. Appl. Math.* **220**, 456–463.
- Y. Zhuang, H. R. Bureau, S. Quirk and R. Hernandez (2021), Adaptive steered molecular dynamics of biomolecules, *Molecular Simul.* **47**, 408–419.
- R. Zwanzig (1973), Nonlinear generalized Langevin equations, *J. Statist. Phys.* **9**, 215–220.