

Optimal Reaction Coordinates: Variational Characterization and Sparse Computation

Andreas Bittracher¹, Mattes Mollenhauer¹, Péter Koltai¹, and
Christof Schütte^{1,2}

¹Department of Mathematics and Computer Science, Freie Universität Berlin, Germany

²Zuse Institute Berlin, Germany

Abstract

Reaction Coordinates (RCs) are indicators of hidden, low-dimensional mechanisms that govern the long-term behavior of high-dimensional stochastic processes. We present a novel and general variational characterization of optimal RCs and provide conditions for their existence. Optimal RCs are minimizers of a certain loss function and reduced models based on them guarantee very good approximation of the long-term dynamics of the original high-dimensional process. We show that, for slow-fast systems, metastable systems, and other systems with known good RCs, the novel theory reproduces previous insight. Remarkably, the numerical effort required to evaluate the loss function scales only with the complexity of the underlying, low-dimensional mechanism, and not with that of the full system. The theory provided lays the foundation for an efficient and data-sparse computation of RCs via modern machine learning techniques.

1. Introduction

Complex high-dimensional dynamical processes, as they occur in molecular dynamics, statistical mechanics or finance, are often modelled by time- and space-continuous Markov processes on the microscopic level. The reason researchers are interested in such processes is however often not so much the microscopic dynamics itself, but rather the phenomenon that, over long time scales, the system often exhibits much more regularity and much less complexity than the sheer number of degrees of freedom would actually allow for. The desire to understand the emergence of this macro-scale behaviour, and to simulate it efficiently, motivates the derivation of reduced, in general non-Markovian models. By projecting the full dynamics onto some selected observables, using techniques such as the Mori-Zwanzig formalism [61], the reduced model describes the long-term dynamical behaviour of these observables, while discarding the short-term microscopic detail. For most such observables the resulting reduced dynamics is too complex (long-term memory effects, complicated noise process) to be useful. However, particular observables allow for a rather direct representation of the reduced dynamics. In the context of molecular dynamics, such particular observables are known as *reaction coordinates* (RCs) or *collective variables*; we will adopt the former designation herein.

This article illuminates the first, and arguably most crucial step in the model reduction pipeline, which is the selection of the “correct” or “optimal” RCs. We present a novel, comprehensive theory of optimal RCs based on a variational principle that is accessible by modern machine learning methods. Its key features are described below. Its differentiation from previous approaches can be found in Section 2. The novel theory presented herein has some crucial advantages:

Variational characterization with dynamical meaning. Our theory is the first to explicitly describe a dynamically interpretable variational principle for optimal RCs that goes beyond alternative approaches based on spectral decomposition or autocorrelation [45, 60] by avoiding any linearization steps. We derive a loss function that measures how well a given RC and the associated reduced dynamics preserve the long-term behaviour of the full system. To quantify the discrepancy in the long-term behaviour, a special metric based on two specific, dynamically interpretable and intuitive properties is used. Optimal RCs, i.e., the global minimizers of the loss function, then inherit this interpretability naturally.

Consistency with established theory. The new theory covers system classes for which good RCs are predefined by established theory or are intuitively clear. Applied to slow-fast dynamical systems [44], or systems with a timescale gap separating slowly and quickly equilibrating sub-processes [55], our variational principle correctly characterizes the respective slow components as optimal RCs. Moreover, our theory extends the transition manifold framework [5], which has been used to successfully identify RCs in several high-dimensional biomolecular systems [3, 6]. By utilizing transition manifold theory, it can be shown how to control the approximation error between the long-term behavior of the original system and the reduced model for optimal RCs.

Sparse sampling of low-rank dynamics. Key parts of our variational problem are insensitive to the curse of dimensionality. By interpreting the existence of a latent low-dimensional reduced model as a low-rank like property, and using arguments similar to low-rank matrix approximation [23] and compressive sampling [18] techniques, we show that approximating our loss function requires only short simulations started from sparsely sampled starting points. More precisely, the number of required starting points does not scale with the dimension of the full system, but only with the dimension and complexity of the reduced system.

Leverage point for modern machine learning. The variational problem can be used directly as a loss function in deep learning methods. We expect that a numerical scheme based on a combination of massively parallel dynamical sampling, automatic differentiation and stochastic gradient descent minimization of our loss function will be effective in identifying interpretable RCs in large-scale real-world systems, such as biomolecular complexes.

This article is structured as follows: Section 2 compares our approach to related work. Section 3 introduces our characterization of good and optimal RCs, and confirms the compatibility of that characterization with established concepts. Section 4 derives the variational principle in form of a loss function for optimal RCs. Section 5 shows how this loss function can be efficiently approximated via a Monte Carlo method with sparse samples. It also discusses the applicability of modern machine learning methods. Finally, section 6 demonstrates the variational principle and the sparse loss function approximation by two synthetic examples. Section 7 contains the conclusions, and an outlook on future work.

2. Related work

Characterization of dynamically meaningful RCs. Existing definitions of “good” RCs were mostly motivated by applications in computational physics and chemistry. However, these field to this date still predominantly use *heuristic* RCs in their model reduction efforts. This practice of manually selecting and combining coordinates from a pool of candidate physical observables such as bond angles or residue distances comes with obvious shortcomings with respect to dynamical meaningfulness, optimality and scalability. Mostly, one is mainly interested in reproducing statistical quantities like the free energy correctly without any strict characterization of the dynamical meaning. We now discuss the most popular of the few approaches to the identification of dynamically meaningful RCs.

The *committor function* is a one-dimensional reaction coordinate, which for some reactant and some product state indicates the probability to hit the latter before hitting the former [17, 35]. However, the committor can only be expected to describe the system’s rate-limiting processes well for sensible choices of the reactant and product (which obviously requires a priori macro-scale knowledge of the system).

The dominant eigenfunctions of the system’s transfer operator (or equivalently its Fokker Planck operator) are mathematically related to committor functions [52]. These eigenfunctions linearly decompose the system into independent sub-processes, which equilibrate with a rate determined by the associated eigenvalue [55]. Hence, the dominant eigenfunctions have been used as RCs [37]. While this corresponds well with our definition of a good RC, it has been demonstrated that the dominant eigenfunctions themselves can be reduced further, if the associated sub-processes are in some way “nonlinearly dependent” on each other [5]. Our second example system (Section 6.2) demonstrates this situation. A RC composed of dominant eigenfunctions is therefore in general not optimal in terms of its dimensionality.

The TICA (time-lagged independent component analysis) method constructs RCs as those linear combinations of the original degrees of freedom with the highest autocorrelation [39, 45]. Assuming the system is reversible, the TICA coordinates are the aforementioned eigenfunctions of the transfer operator projected onto the linear basis functions $\psi_i(x) = x_i$ [30]. They therefore suffer from the same non-optimality as the eigenfunction RCs, and in addition from non-optimality due to the overhead of linear approximation.

The most frequently analyzed special case of timescale-separated systems are *slow-fast systems* (see [44] for a text book introduction). They are characterized by the existence of a coordinate transformation such that the new coordinates can be subdivided into one quickly and one slowly moving part, and the two parts are approximately decoupled. The slow coordinates, which form a parametrization of the system’s slow manifold [59], can be considered, and can be used as a good RC of the system [57, 19]. We will see later that our characterization encompasses that of slow variables.

The RCs conceptually most alike to the definition presented in this article are the *transition manifold RCs*, proposed in [5] and further refined in [7]. Transition manifold theory characterizes good RCs as a parametrization of a low-dimensional manifold in a certain function space. This *transition manifold* emerges from the system’s transition densities with progressing equilibration. The characterization presented in the present article is a significant generalization of [7].

Computational strategies. Most of the aforementioned approaches to dynamically meaningful RCs come with a proposed numerical scheme for their computation. The committor function, which satisfies a backward Kolmogorov equation [17], can be computed using numerical PDE solvers (although this was never proposed as a practical scheme and a vastly more efficient scheme was proposed soon-after [35]). RCs based on transfer operator eigenfunctions (including TICA) can be computed by an eigendecomposition of a suitable discretization of that operator [12, 54, 15, 45]. Approaches that characterize RCs as parametrization of some manifold use unsupervised manifold learning methods such as diffusion maps to learn the variables in an equation-free manner [57, 19, 5, 4].

Over time, deep-rooted relationships between the different RCs, as well as extensions and generalizations were discovered (see [28] for a partial overview), which led to alternative and more efficient schemes for their computation. While a comprehensive listing would go beyond the scope of this article, we want to point out an emerging trend in these efforts: the formulation of a variational principle for the respective RC. There now exist variational approaches for the committor function [27], the TICA coordinate [60], the transfer operator eigenfunctions [36], and related dominant subspaces [46]. Alternative approaches try to maximize timescale separation between slow and fast processes via maximum caliber-based frameworks [43]. The driving force behind this trend is of course the desire to profit from the impressive performance that modern deep learning and neural network-based methods have demonstrated with regard to their generalization power, robustness to overfitting and seeming immunity to the curse of dimensionality [2].

Notably missing from the above list is however a variational principle for manifold-based RCs. As mentioned before, the characterization presented in this article generalizes the transition manifold, which in turn generalizes the slow manifold, so it can be seen as a completion in that regard. The variational approach then offers the additional advantage of yielding a closed form of the RC (in some finite-dimensional ansatz space), unlike the aforementioned geometric manifold learning algorithms, which output only discrete point-evaluations of the RC.

Sparsity. Our strategy to approximate the loss function from sparse samples of the dynamics shows parallels to other computational techniques that implicitly exploit some form of hidden regularity of the problem.

In a recent publication [8], a discrete version of the sparse sampling strategy to time- and space-*discrete* Markov chains was proposed. There it was assumed that the long-term behavior of a large (i.e., many-state) Markov chain is essentially determined by transitions between certain *aggregates* of these states. It was shown that the aggregates and the transition probabilities between them could be discovered from a vastly undersampled version of the transition matrix of the original chain (obtainable through simulations starting from random states). Some of the key ideas of [8] will be exploited herein by putting them into the continuous setting.

The fundamental idea behind both the discrete and continuous sparse sampling strategies is heavily inspired by the field of compressive sensing, see [18] for an introduction. The impressive feat of compressive sensing is its ability to re-construct a signal (a high-dimensional vector) from far less samples than the Nyquist–Shannon sampling theorem would actually demand by means of solving vastly underdetermined linear equations. The necessary assumption is that the system is *sparse* in some basis (for example the frequency domain), though the location or precise number of the sparse entries does not need to be known. In a way, the dominance of a Markov process by a single low-dimensional mechanism can be interpreted as a sort of “non-linear dynamical sparsity”.

Finally we would like to point out that the apparent similarity of our sampling strategy to randomized matrix low-rank approximation techniques [23] like the Nyström method [16] or randomized feature approximation [47] is rather superficial. While for these techniques a (nearly) low-rank structure of the target matrix is necessary to achieve low approximation error, they do not interpret this low rank as an underlying structure “generating” the matrix. Indeed, in [58] it has been argued that an (approximate) low rank is a generic property of large data matrices, and that the attribution of that rank to some “physical reason” is in general not possible. Consequently, this rank and with it the sampling requirement still scales with the size of the matrix (typically logarithmically [58]). In contrast, our theory assumes and exploits the existence of an underlying mechanism of fixed low dimension that governs the micro-scale model, which effectively decouples the sampling requirement from the size of the full model.

3. Characterization of good RCs

This section contains the central definition of good RCs and corresponding conditions on the dynamics for their existence. First, however, we introduce the notation and the dynamical setting.

3.1. Definition of the dynamics and fundamental assumptions

Let $\mathbb{X} \subset \mathbb{R}^n$ be a Lebesgue-measurable set (the *state space*) and $(X_t)_{t \in \mathbb{R}^+}$, or short (X_t) , be a time- and space-continuous Markov process on \mathbb{X} . Let $P^t : \mathbb{X} \times \mathcal{B} \rightarrow [0, 1]$ denote the *transition probability function* of (X_t) , where \mathcal{B} is the Borel σ -algebra on \mathbb{X} , i.e.,

$$P^t[x, B] = \text{Prob}[X_{t_0+t} \in B \mid X_{t_0} = x] \quad \text{for all } t_0 \geq 0.$$

For any $t > 0$ and $x \in \mathbb{X}$, $P^t(x, \cdot)$ is a probability measure on \mathcal{B} , and $P^t(\cdot, B)$ is a \mathcal{B} -measurable function for any $B \in \mathcal{B}$ [50]. Moreover, let the process be ergodic, such that a unique stationary

measure $\mu : \mathcal{B} \rightarrow \mathbb{R}_0^+$ exists. We require μ to be absolutely continuous with respect to the Lebesgue measure, i.e., there exists a density $\pi : \mathbb{X} \rightarrow \mathbb{R}^+$ such that

$$\mu(B) = \int_B \pi(x) dx.$$

Moreover, we require that π is continuous and strictly positive.

We also require the $P^t(x, \cdot)$ to be absolutely continuous with respect to the Lebesgue measure. Thus, we may assume that there exists a family of functions $p^t : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}^+$ such that

$$P^t(x, B) = \int_B p^t(x, y) dy \quad \text{for all } \tau > 0, x \in \mathbb{X}, B \in \mathcal{B}. \quad (1)$$

Many classes of Markov processes are absolutely continuous, including Itô diffusions with smooth coefficients [32]. Also, we assume that the system is reversible with respect to π , i.e., the *detailed balance equation* holds:

$$p^t(x, y)\pi(x) = p^t(y, x)\pi(y) \quad \text{for all } x, y \in \mathbb{X}, t \in \mathbb{R}^+. \quad (2)$$

The existence of good RCs will be determined by specific properties of the function p^t , hence we now examine it more closely. As a function of the second argument, $p^t(x, \cdot) \in L^1$ is the time- t *transition density function* of (X_t) , i.e.

$$p^t(x, \cdot) = \text{Law}(X_{t_0+t} \mid X_{t_0} = x) \quad \text{for all } t_0 \geq 0.$$

On the other hand, $p^t(\cdot, y)$ as a function of the first argument is harder to interpret, and discussed less in the literature of stochastic processes. Let $L_\mu^1(\mathbb{X})$ be the space equipped with the norm

$$\|f\|_{L_\mu^1} := \int_{\mathbb{X}} f(x) d\mu(x).$$

We then have $p^t(\cdot, y) \in L_\mu^1$, since, by reversibility of X_t ,

$$\begin{aligned} \int_{\mathbb{X}} p^t(x, y) d\mu(x) &= \int_{\mathbb{X}} p^t(x, y)\pi(x) dx \\ &= \int_{\mathbb{X}} p^t(y, x)\pi(y) dx \leq \|\pi\|_\infty \|p^t(y, \cdot)\|_{L^1} < \infty. \end{aligned}$$

To distinguish it from the transition density, we call $p^t(\cdot, y)$ the time- t *transition observable* of y .

As a function of *two* arguments, we call $p^t : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ the time- t *transition kernel* of (X_t) . It can be interpreted as an element of the space of functions $L_{\mu \times \lambda}^1(\mathbb{X}^2)$, where $\mu \times \lambda$ is the product measure on the space \mathbb{X}^2 given by the invariant measure μ and the Lebesgue measure λ on \mathbb{X} . For simplicity, we use the shorthand notation

$$\mathbb{K} := L_{\mu \times \lambda}^1(\mathbb{X}^2).$$

Note that by Fubini–Tonelli, we have

$$\|p(*, \cdot)\|_{\mathbb{K}} := \left\| \|p(*, \cdot)\|_{L^1(\mathbb{X})} \right\|_{L_\mu^1(\mathbb{X})} \quad (3)$$

as the norm on \mathbb{K} , where in (3) the inner norm applies to the argument “ \cdot ”, and the outer norm applies to the argument “ $*$ ” (this will be a convention from now on).

3.2. Lumpability and decomposability

We will now introduce two seemingly different conditions for a system/RC pair. Each condition may individually be taken as a definition for what a good RC is. It will turn out, however, that the two conditions are equivalent to each other for reversible systems, so a good RC with respect to one condition is a good RC with respect to the other.

Let $r < n$, and $\mathbb{Z} \subset \mathbb{R}^r$ be a domain. We call any function $\xi \in C(\mathbb{X}, \mathbb{Z})$ an *r-dimensional RC*. Later, in particular in Section 4, we will require ξ to be “smooth”:

Definition 3.1 (Smooth RCs). *Let $r < n$, and $\mathbb{Z} \subset \mathbb{R}^r$ be a domain. If $\xi : \mathbb{X} \rightarrow \mathbb{Z}$ fulfils the two assumptions*

1. ξ is continuously differentiable,
2. for all $z \in \mathbb{Z}$, the level sets

$$\Sigma_\xi(z) := \{x \in \mathbb{X} \mid \xi(x) = z\}$$

are $n - r$ -dimensional topological submanifolds of \mathbb{X} ,

we call ξ a smooth r -dimensional RC. We denote the function space of smooth RCs from \mathbb{X} to \mathbb{Z} by $S(\mathbb{X}, \mathbb{Z})$.

In case where $\xi \in S(\mathbb{X}, \mathbb{Z})$, the marginal stationary measure μ_z on $\Sigma_\xi(z)$ can be defined by

$$\int_{\Sigma_\xi(z)} f(x) d\mu_z(x) = \int_{\Sigma_\xi(z)} f(x) \pi(x) \det(\nabla \xi(x)^\top \nabla \xi(x))^{-1/2} dH_{n-r}(x), \quad (4)$$

where H_d denotes the d -dimensional Hausdorff measure. By $|\mathbb{Z}|$, we denote the Lebesgue-measure of \mathbb{Z} .

We emphasize, however, that we do not require smoothness for the following characterizations of good RCs.

Lumpability

The first condition on the process (X_t) is as follows:

Definition 3.2 (Lumpability). *If there exists a domain $\mathbb{Z} \subset \mathbb{R}^r$, a function $\xi \in C(\mathbb{X}, \mathbb{Z})$, a family of time-parametrized functions $p_L^t : \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}^+$ and a lag time $\tau > 0$ such that*

$$\frac{1}{|\mathbb{Z}|} \|p^t(*, \cdot) - p_L^t(\xi(*), \cdot)\|_{\mathbb{K}} \leq \varepsilon \quad (\text{L})$$

is fulfilled for all $t \geq \tau$, we say the system is ε -lumpable with respect to ξ .

In words, lumpability means that for sufficiently large t , the transition densities $p^t(x, \cdot)$, i.e., the probabilities of transitions starting from x , depend essentially only on the value $\xi(x)$ of the RC at x , and not on the precise location of x on the $\xi(x)$ -level set of ξ . Condition (L) is therefore a sensible definition for a RC that is supposed to describe the effective long-term dynamics of X_t .

Remark 3.3. Our definition of lumpability can be seen as a continuous version of the lumpability condition for discrete Markov chains, originally formulated by Kemeny and Snell [26]. There it has been shown that the discrete version of lumpability is a necessary and sufficient condition for the Markov chain to be “compressible” into a Markov chain between certain *aggregates* of the original chain. In general, this compression however leads to a loss of information, i.e., restoration of the original chain is in general not possible under the lumpability assumption alone.

Decomposability

The second condition on (X_t) we will discuss is the following:

Definition 3.4 (Decomposability). *If there exists a domain $\mathbb{Z} \subset \mathbb{R}^r$, a function $\xi \in C(\mathbb{X}, \mathbb{Z})$, a family of time-parametrized functions $p_D^t : \mathbb{X} \times \mathbb{Z} \rightarrow \mathbb{R}^+$ and a lag time $\tau > 0$ such that*

$$\frac{1}{|\mathbb{Z}|} \|p^t(*, \cdot) - p_D^t(*, \xi(\cdot))\pi(\cdot)\|_{\mathbb{K}} \leq \varepsilon \quad (\text{D})$$

is fulfilled for $t \geq \tau$, we say the system is ε -decomposable with respect to ξ .

The value of $p^t(x, y)$ can be interpreted as the transition probability from an infinitesimal neighborhood around x to an infinitesimal neighborhood around y . Decomposability now holds, among others, for systems for which this transition can be decomposed into two consecutive stages:

- 1) a slow transition from x to anywhere on the level set $\Sigma_\xi(\xi(y))$, for which the (infinitesimal) transition probability is given by $p_D^t(x, \xi(y))$, and
- 2) an instantaneous equilibration on that level set with respect to the invariant density $\pi(y)$, which does no longer depend on the starting point.

Hence, for systems for which this decomposition exists, condition (D) again describes a sensible definition of good RCs. Note however that not only such “physically decomposable” systems fulfil the above definition, but that (D) describes a more generic “conceptual decomposability”, as we will see in Section 3.3.

Remark 3.5. The above notion of decomposability also has a time- and space-discrete equivalent for discrete Markov chains, which was defined by some of the authors in [8] (however the condition was called “deflatability” there). There it was shown that for Markov chains fulfilling both the lumpability and deflatability condition, a “compressed” Markov chain between certain aggregates of the original states exists, and that the full Markov chain can be restored from the compressed chain.

Remark 3.6. It should be noted that, for large enough lag time, every uniformly ergodic system is trivially lumpable and decomposable since

$$\sup_{x \in \mathbb{X}} \|P^t(x, \cdot) - \mu(\cdot)\|_{TV} = \sup_{x \in \mathbb{X}} \frac{1}{2} \|p^t(x, \cdot) - \pi(\cdot)\|_{L^1} \rightarrow 0 \quad (5)$$

as $t \rightarrow \infty$ (see for example [49, Section 3.3 together with Proposition 3(f)]. Hence, choosing $p_L^t(z, \cdot) = \pi$ and $p_D^t(z, \cdot) = 1$ in the above definitions will give lumpability and decomposability with respect to any constant RC since then p_L^t and p_D^t are independent of x . Likewise, every system is trivially lumpable and decomposable, for any tolerance and lag time, with respect to the trivial n -dimensional RC $\xi(x) = x$, since choosing $p_L^t(x, \cdot) = p^t(x, \cdot)$ and $p_D^t(\cdot, y) = p^t(\cdot, y)/\pi(y)$ fulfils (L) and (D) with $\varepsilon = 0$. We emphasize that in this paper we specifically care about systems which are lumpable/decomposable with respect to intermediate lag times τ that are much smaller than the equilibration time scale of the system, as well as small dimensions r and $\varepsilon > 0$.

Moreover, a system may be lumpable/decomposable with respect to more than one non-trivial combination of ε, τ and r . In cases where no clear time scale separation exists in the full system, a balance has to be struck between the achievable approximation error of a reduced model built using ξ (acceptable ε), the time scale above which the reduced model is valid (choice of τ), and the dimension of the reduced model (choice of r).

Although an interesting objective especially in systems with cascades of time scales, optimizing the choice of r and τ are not subject of this paper. We will later consider r and τ to be fixed, and search for corresponding “optimal” RCs, i.e. an r -dimensional ξ for which (L) and/or (D) are fulfilled for the smallest possible ε (more on that in Section 4).

Connection to Reversibility

As mentioned above, the two conditions (L) and (D) are equivalent in reversible systems:

Proposition 3.7. *Let the system be reversible, i.e., let (2) hold. Then the system is ε -lumpable if and only if it is ε -decomposable.*

Proof. Let (L) hold for some family of functions $p_L^t : \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}$. Define the family of functions $p_D^t : \mathbb{X} \times \mathbb{Z} \rightarrow \mathbb{R}$ by

$$p_D^t(x, z) := \frac{p_L^t(z, x)}{\pi(x)}.$$

Then

$$\begin{aligned}
\|p_D^\tau(x, \xi(\cdot))\pi(\cdot) - p^\tau(x, \cdot)\|_{L^1} &= \|p_L^\tau(\xi(\cdot), x) \frac{\pi(\cdot)}{\pi(x)} - p^\tau(x, \cdot)\|_{L^1} \\
&\stackrel{(2)}{=} \|p_L^\tau(\xi(\cdot), x) \frac{\pi(\cdot)}{\pi(x)} - p^\tau(\cdot, x) \frac{\pi(\cdot)}{\pi(x)}\|_{L^1} \\
&= \|p_L^\tau(\xi(\cdot), x) - p^\tau(\cdot, x)\|_{L_\mu^1} \pi(x)^{-1}.
\end{aligned}$$

Hence, with this p_D^t , it holds

$$\begin{aligned}
\frac{1}{|\mathbb{Z}|} \|p^\tau(*, \cdot) - p_D^\tau(*, \xi(\cdot))\pi(\cdot)\|_{\mathbb{K}} &= \frac{1}{|\mathbb{Z}|} \int_{\mathbb{X}} \|p^\tau(x, \cdot) - p_D^\tau(x, \xi(\cdot))\pi(\cdot)\|_{L^1} d\mu(x) \\
&= \frac{1}{|\mathbb{Z}|} \int_{\mathbb{X}} \|p_L^\tau(\xi(x), \cdot) - p^\tau(x, \cdot)\|_{L_\mu^1} \pi(x)^{-1} d\mu(x) \\
&= \frac{1}{|\mathbb{Z}|} \|p_L^\tau(\xi(*), \cdot) - p^\tau(*, \cdot)\|_{\mathbb{K}} \leq \varepsilon.
\end{aligned}$$

For the reverse direction, let (D) hold for some family of functions $p_D^t : \mathbb{X} \times \mathbb{Z} \rightarrow \mathbb{R}$ and define $p_L^t : \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}$ by

$$p_L^t(z, y) := p_D^t(y, z)\pi(y).$$

We then obtain ε -lumpability with respect to ξ and this p_L^t by performing the above transformations in reverse. \square

Remark 3.8. Proposition 3.7 implies that, whenever a system is ε -lumpable or ε -decomposable, there exists a reduced transition kernel $\tilde{p}^t : \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{R}^+$, such that

$$\|p^t(*, \cdot) - \tilde{p}^t(\xi(*), \xi(\cdot))\pi(\cdot)\|_{\mathbb{K}} \leq \varepsilon$$

for $t \geq \tau$. Under this condition, knowing the triple (ξ, \tilde{p}^t, π) allows us to approximately re-construct the effective long-term dynamics of the full system. This represents the theoretical justification of our definition of a good RC.

3.3. Examples of lumpable and decomposable systems

An obvious question is whether the conditions (L) and (D) are consistent with established concepts of reducible systems. We hence now present several systems with known good RCs and show that they are indeed either lumpable or decomposable.

3.3.1. Existence of a transition manifold

Lumpability is strongly connected to the concept of the so-called *transition manifold*, which was introduced a few years ago by some of the authors in order to formulate a geometrical approach to the computation of optimal RCs [5].

Definition 3.9 (Reducibility and Transition Manifold). *For $\varepsilon > 0, r \leq n, \tau \in \mathbb{R}_0^+$, we call the system (ε, r, τ) -reducible if there exists an r -dimensionally parametrizable manifold $\mathbb{M} \subset \{p^\tau(x, \cdot), x \in \mathbb{X}\}$ so that for all $x \in \mathbb{X}$*

$$\|Q(x) - p^\tau(x, \cdot)\|_{L_{1/\pi}^2} \leq \varepsilon, \tag{6}$$

where $Q : \mathbb{X} \rightarrow \mathbb{M}$ is the nearest point projection onto \mathbb{M} ,

$$Q(x) := \arg \min_{p \in \mathbb{M}} \|p^\tau(x, \cdot) - p\|_{L_{1/\pi}^2}. \tag{7}$$

We call any \mathbb{M} that fulfills (6) a transition manifold of the system.

The intuition behind (6) is that the set of all transition densities $\{p^t(x, \cdot), x \in \mathbb{X}\}$ clusters ε -closely around an r -dimensional manifold \mathbb{M} with respect to the $L^2_{1/\pi}$ norm. The density $\mathcal{Q}(x)$ can be understood as the reduced density $p_L^\tau(\xi(x), \cdot)$ in (L). Let $\mathcal{E} : \mathbb{M} \rightarrow \mathbb{R}^r$ be any parametrization of \mathbb{M} . It can be shown that the *transition manifold RC*

$$\xi(x) := \mathcal{E}(\mathcal{Q}(x)) \quad (8)$$

is a good RC, in the sense that the projection error of the leading transfer operator eigenfunctions onto ξ is at most ε [5].

The computational strategy behind the transition manifold approach now is to sample the set $\{p^t(x, \cdot), x \in \mathbb{X}\}$ (for example by randomly selecting starting points $x_m \in \mathbb{X}$, $m = 1, 2, \dots$ and estimating the $p^\tau(x_m, \cdot)$ by parallel simulation), and applying an *unsupervised manifold learning method* (such as diffusion maps [11]) to the samples. This strategy has been successfully applied to multiple high-dimensional molecular systems and was confirmed to produce physically interpretable RCs [3, 4].

The concept of the transition manifold was recently re-visited and extended to a broader class of dynamical systems [7] (the central object now being called *weak transition manifold*), by requiring the closeness to the manifold now only averaged over the level sets of \mathcal{Q} . It however requires the restriction to the following class of “smooth” transition manifolds:

Definition 3.10 (Smooth manifold). *We call an r -dimensionally parametrizable manifold $\mathbb{M} \subset \{p^\tau(x, \cdot), x \in \mathbb{X}\}$ smooth, if for its nearest point projection \mathcal{Q} holds*

1. \mathcal{Q} is continuously differentiable,
2. for each $p \in \mathbb{M}$, the level set

$$\Sigma_{\mathcal{Q}}(p) = \{x \in \mathbb{X} \mid \mathcal{Q}(x) = p\}$$

is an $n - r$ -dimensional topological submanifold of \mathbb{X} .

With that we can introduce the weak transition manifold:

Definition 3.11 (Weak reducibility and weak transition manifold). *For $\varepsilon > 0, r \leq n, \tau \in \mathbb{R}_0^+$, we call the system weakly (ε, r, τ) -reducible if there exists an r -dimensionally parametrizable smooth manifold $\mathbb{M} \subset \{p^\tau(x, \cdot), x \in \mathbb{X}\}$ so that for all $x \in \mathbb{X}$*

$$\int_{\Sigma_{\mathcal{Q}}(\mathcal{Q}(x))} \|\mathcal{Q}(x) - p^\tau(x', \cdot)\|_{L^2_{1/\pi}} d\mu_{\mathcal{Q}(x)}(x') \leq \varepsilon, \quad (9)$$

where $\mu_{\mathcal{Q}}$ is the marginal invariant measure on $\Sigma_{\mathcal{Q}}(p)$, defined analogously to (4). We call any smooth manifold \mathbb{M} that fulfils (9) a weak transition manifold.

It is easy to confirm that every reducible system with a smooth transition manifold is weakly reducible.

The closeness condition for the weak transition manifold (9) is similar in nature to the lumpability condition (L), in that the original transition density functions, $p^\tau(x, \cdot)$, are close so some projected or reduced transition density function, $\mathcal{Q}(p^\tau(x, \cdot))$ and $p_L(\xi(x), \cdot)$, respectively. Indeed, we will now show that a system which is weakly (ε, r, τ) -reducible is ε -lumpable with respect to the transition manifold RC ξ . Lumpability, and equivalently decomposability, are therefore generalizations of weak reducibility.

Proposition 3.12. *Let the system be weakly (ε, r, τ) -reducible with transition manifold \mathbb{M} . Assume that there exists a parametrization $\mathcal{E} : \mathbb{M} \rightarrow \mathbb{Z} \subset \mathbb{R}^r$ such that ξ defined by (8) is a smooth RC (see Definition 3.1). Then there exists a family of functions $p_L^\xi : \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}^+$ such that (L) is fulfilled with respect to ξ .*

Proof. Let $\mathcal{E} : \mathbb{M} \rightarrow \mathbb{R}^r$ be any parametrization of the transition manifold \mathbb{M} and let $\mathbb{Z} := \mathcal{E}(\mathbb{M})$. In particular, $\mathcal{E} : \mathbb{M} \rightarrow \mathbb{Z}$ is one-to-one. Define the RC ξ by (8) and the reduced density p_L^τ by

$$p_L^\tau(z, \cdot) := \mathcal{E}^{-1}(z).$$

Then for any $z \in \mathbb{Z}$ there exists an $x \in \mathbb{X}$ such that $z = \mathcal{E}(\mathcal{Q}(x))$ and hence, due to \mathcal{E} being one-to-one, $\Sigma_\xi(z) = \Sigma_{\mathcal{Q}}(\mathcal{Q}(x))$. For this x it holds

$$\begin{aligned} \int_{\Sigma_\xi(z)} \|p_L^\tau(z, \cdot) - p^\tau(x', \cdot)\|_{L^1} d\mu_z(x') &= \int_{\Sigma_{\mathcal{Q}}(\mathcal{Q}(x))} \|p_L^\tau(\mathcal{E}(\mathcal{Q}(x)), \cdot) - p^\tau(x', \cdot)\|_{L^1} d\mu_{\mathcal{Q}(x)}(x') \\ &= \int_{\Sigma_{\mathcal{Q}}(\mathcal{Q}(x))} \|\mathcal{E}^{-1}(\mathcal{E}(\mathcal{Q}(x))) - p^\tau(x', \cdot)\|_{L^1} d\mu_{\mathcal{Q}(x)}(x') \\ &= \int_{\Sigma_{\mathcal{Q}}(\mathcal{Q}(x))} \|\mathcal{Q}(x) - p^\tau(x', \cdot)\|_{L^1} d\mu_{\mathcal{Q}(x)}(x'). \end{aligned}$$

By application of the co-area formula, and the definition (4) of μ_z we get

$$\frac{1}{|\mathbb{Z}|} \|p^\tau(*, \cdot) - p_L^\tau(\xi(*), \cdot)\|_{\mathbb{K}} = \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\xi(z)} \|p_L^\tau(z, \cdot) - p^\tau(x', \cdot)\|_{L^1} d\mu_z(x') dz,$$

which finally, with $\|f\|_{L^1} = \|f/\pi\|_{L_\mu^1} \leq \|f/\pi\|_{L_\mu^2} = \|f\|_{L_{1/\pi}^2}$, becomes

$$\begin{aligned} &\leq \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\mathcal{Q}}(\mathcal{Q}(x))} \|\mathcal{Q}(x) - p^\tau(x', \cdot)\|_{L_{1/\pi}^2} d\mu_{\mathcal{Q}(x)}(x') dz \\ &\leq \sup_{z \in \mathbb{Z}} \int_{\Sigma_{\mathcal{Q}}(\mathcal{Q}(x))} \|\mathcal{Q}(x) - p^\tau(x', \cdot)\|_{L_{1/\pi}^2} d\mu_{\mathcal{Q}(x)}(x') \\ &\stackrel{(9)}{\leq} \varepsilon. \end{aligned}$$

□

3.3.2. Slow- and fast components

Next, we show that a process defined by an SDE with slow and fast components is decomposable with respect to the slow component. We utilize a multiscale decomposition of the corresponding infinitesimal generator, together with a multiscale ansatz for the transition densities $p^t(x, \cdot)$. In that we utilize well-known averaging techniques from [44].

It will prove advantageous to consider the transition densities as densities with respect to the stationary density π . That is, we define for each $x \in \mathbb{X}$ the density $q^t(x, \cdot)$ with respect to π by

$$p^t(x, \cdot) = q^t(x, \cdot)\pi(\cdot).$$

Let the components of (X_t) be dividable into two processes (Y_t) on \mathbb{Y} , (Z_t) on \mathbb{Z} , such that $\mathbb{X} = \mathbb{Y} \oplus \mathbb{Z}$,

$$(X_t) = \begin{pmatrix} Y_t \\ Z_t \end{pmatrix},$$

and the components fulfil the system of SDEs

$$\begin{aligned} \varepsilon dY_t &= -\nabla_y V(Y_t, Z_t) dt + \sqrt{\varepsilon} \sigma dW_t^{\mathbb{Y}} \\ dZ_t &= -\nabla_z V(Y_t, Z_t) dt + \sigma dW_t^{\mathbb{Z}}, \end{aligned} \tag{10}$$

with potential $V : \mathbb{X} \rightarrow \mathbb{R}$, scalar diffusion parameter $\sigma > 0$ and $(W_t^{\mathbb{Y}})$, $(W_t^{\mathbb{Z}})$ standard Wiener processes on \mathbb{Y} and \mathbb{Z} , respectively. A timescale separation parameter $0 < \varepsilon \ll 1$ ensures that (Y_t) evolves “fast” compared to (Z_t) .

The evolution of q^t under (X_t) is now governed by the Fokker Planck equation

$$\partial_t q^t(x, \cdot) = \mathcal{L}_\varepsilon q^t(x, \cdot), \quad q^0(x, \cdot) = \delta_x(\cdot),$$

where the infinitesimal generator \mathcal{L}_ε has the multiscale structure

$$\mathcal{L}_\varepsilon = \frac{1}{\varepsilon} \mathcal{L}_y + \mathcal{L}_z, \quad (11)$$

with the two components

$$\begin{aligned} \mathcal{L}_y &= \frac{\sigma^2}{2} \Delta_y - \nabla_y V \cdot \nabla_y, \\ \mathcal{L}_z &= \frac{\sigma^2}{2} \Delta_z - \nabla_z V \cdot \nabla_z. \end{aligned}$$

We want to investigate to what extent $q^t(x, \cdot)$, and by extension $p^t(x, \cdot)$, can be approximated by an ‘‘essential transition density’’ that depends only on the slow variable z in the setting $\varepsilon \ll 1$, $t = \mathcal{O}(1)$. To this end, we make the multiscale ansatz for q^t

$$q^t(x, \cdot) = q_0^t(x, \cdot) + \varepsilon q_1^t(x, \cdot) + \mathcal{O}(\varepsilon^2). \quad (12)$$

Inserting (12) into (11) gives

$$\partial_t q_0^t(x, \cdot) + \varepsilon \partial_t q_1^t(x, \cdot) + \mathcal{O}(\varepsilon^2) = \frac{1}{\varepsilon} \mathcal{L}_y q_0^t(x, \cdot) + \mathcal{L}_y q_1^t(x, \cdot) + \mathcal{L}_z q_0^t(x, \cdot) + \mathcal{O}(\varepsilon). \quad (13)$$

Comparing the terms of order ε^{-1} gives

$$\mathcal{L}_y q_0^t(x, \cdot) = 0.$$

By [44, Sec. 10.2], \mathcal{L}_y has a one-dimensional null space consisting of functions constant in y . Hence, $q_0^t(x, \cdot)$ is independent of y in its second argument, and so

$$q^t((y, z), (y', z')) = q_0^t((y, z), z') + \mathcal{O}(\varepsilon).$$

Therefore, for $t = \mathcal{O}(1)$, the transition density $p^t(x, \cdot)$ takes the form

$$p^t((y, z), (y', z')) = q_0^t((y, z), z') \pi((y', z')) + g((y, z), (y', z')) \pi((y', z')).$$

for some function $g \in \mathcal{O}(\varepsilon)$. Applying the $\|\cdot\|_{\mathbb{K}}$ -norm, it follows that the system is decomposable with respect to the RC $\xi(x) = z$, i.e.,

$$\frac{1}{|\mathbb{Z}|} \|q_0^t(*, \xi(\cdot)) \pi(\cdot) - p^t(*, \cdot)\|_{\mathbb{K}} = \mathcal{O}(\varepsilon).$$

Remark 3.13. The above result, or to be precise, the multiscale extension (12), holds only for times t up to $\mathcal{O}(1)$, as one can show that the second term q_1 is bounded uniformly only in t [31], hence may grow linearly in t . However, we are only interested in lumpability on moderate lag times anyway, as every system becomes trivially lumpable for $t \rightarrow \infty$ (see Remark 3.6).

Remark 3.14. While not necessarily in the scope of this paper, we can continue the multiscale analysis in order to derive an evolution equation for q_0^t . See Appendix A for details.

3.3.3. Generator spectral gap

Finally, we show that systems whose infinitesimal generator exhibits a spectral gap of sufficient size, such as metastable systems [12, 14], are decomposable with respect to some non-trivial RC.

Consider again the transition densities q^t with respect to the stationary measure, and its Fokker Planck equation

$$\partial_t q^t(x, \cdot) = \mathcal{L}q^t(x, \cdot), \quad q^0(x, \cdot) = \delta_x(\cdot).$$

We assume that the spectrum of the generator \mathcal{L} is real, non-positive and discrete, and denote the eigenvalues of \mathcal{L} in descending order:

$$0 = \theta_0 > \theta_1 \geq \theta_2 \geq \dots,$$

repeated by geometric multiplicity. Let furthermore φ_i denote the eigenfunction belonging to θ_i . The φ_i then form an orthonormal basis of $L_\mu^2(\mathbb{X})$ [42]. As $q^t(x, \cdot) \in L_\mu^1(\mathbb{X}) \cap L_\mu^\infty(\mathbb{X})$ for any x , $q^t(x, \cdot) \in L_\mu^2(\mathbb{X})$. We can then describe the evolution of the density $q^t(x, \cdot)$ by

$$q^t(x, \cdot) = \sum_{i=0}^{\infty} e^{\theta_i t} c_i(x) \varphi_i(\cdot) \quad (14)$$

where $c_i : \mathbb{X} \rightarrow \mathbb{R}$ are some functions that do not depend on the θ_i .

Now, we additionally assume that the eigenvalues can be separated into a dominant and a non-dominant part. Specifically, we assume there exists an index $K > 0$, so that the ratio

$$\frac{\theta_K}{\theta_{K+1}} > 0 \quad (15)$$

is small. This situation for example occurs if the system is *metastable*, i.e., there exists a partition $\mathbb{X} = \mathbb{X}_1 \cup \dots \cup \mathbb{X}_K$ of state space into disjoint regions, and the system is almost invariant on each \mathbb{X}_i on relatively long time scales. For a precise introduction of metastability and its connection to the dominant spectrum see [55].

Now suppose that there exists an integer $r \leq K$ and a RC $\xi : \mathbb{X} \rightarrow \mathbb{R}^r$ such that ξ parametrizes the dominant φ_i , i.e., there exist some functions $\tilde{\varphi}_i : \mathbb{R}^r \rightarrow \mathbb{R}$ such that

$$\varphi_i = \tilde{\varphi}_i \circ \xi \quad i = 1, \dots, K. \quad (16)$$

Such a ξ always exists, as one can always choose

$$r := K, \quad \xi_i := \varphi_i \quad \text{and} \quad \tilde{\varphi}_i(z) := z_i.$$

Often, however, the dominant eigenfunctions possess some common lower-dimensional, non-linear parametrization. For metastable systems, this is for example the case if the metastable sets $\mathbb{X}_1, \dots, \mathbb{X}_K$ are connected by just a small number of transition pathways. An example system with five metastable sets, but one common transition pathway, hence a one-dimensional RC, can be found in Section 6.2.

Let $\varepsilon > 0$ be some small constant. We now show that, if $t = t(\varepsilon)$ is large enough, and the spectral ratio (17) is small enough, then the system is ε -decomposable with respect to any ξ fulfilling (16). To see this, split the right hand side of (14) into the dominant and the non-dominant part:

$$q^t(x, \cdot) = \sum_{i=0}^K e^{\theta_i t} c_i(x) \varphi_i(\cdot) + \sum_{i=K+1}^{\infty} e^{\theta_i t} c_i(x) \varphi_i(\cdot).$$

Due to (16), the first summand depends only on ξ :

$$\sum_{i=0}^K e^{\theta_i t} c_i(x) \varphi_i(\cdot) = \underbrace{\sum_{i=0}^K e^{\theta_i t} c_i(x) \tilde{\varphi}_i(\xi(\cdot))}_{=: p_D^t(x, \xi(\cdot))}.$$

To show that the system is ε -decomposable with respect to ξ and p_D^t , we thus have to ensure that

$$\frac{1}{|\mathbb{Z}|} \left\| \sum_{i=K+1}^{\infty} e^{\theta_i t} c_i(\cdot) \varphi_i(\cdot) \pi(\cdot) \right\|_{\mathbb{K}} \leq \varepsilon. \quad (17)$$

Since the θ_i are decreasing, it holds

$$\left\| \sum_{i=K+1}^{\infty} e^{\theta_i t} c_i(*) \varphi_i(\cdot) \pi(\cdot) \right\|_{\mathbb{K}} \leq e^{\theta_{K+1} t} \left\| \sum_{i=K+1}^{\infty} c_i(*) \varphi_i(\cdot) \pi(\cdot) \right\|_{\mathbb{K}},$$

and further

$$\leq e^{\theta_{K+1} t} \left(\underbrace{\left\| \sum_{i=0}^K c_i(*) \varphi_i(\cdot) \pi(\cdot) \right\|_{\mathbb{K}}}_{=: \tilde{C}} + \left\| \sum_{i=0}^{\infty} c_i(*) \varphi_i(\cdot) \pi(\cdot) \right\|_{\mathbb{K}} \right).$$

The first summand, denoted \tilde{C} , is finite as a finite sum. As

$$p^t(*, \cdot) = q^t(*, \cdot) \pi(\cdot) = \sum_{i=0}^{\infty} e^{\theta_i t} c_i(*) \varphi_i(\cdot) \pi(\cdot),$$

and $1 = e^{\theta_i 0}$, the second summand can be interpreted as the \mathbb{K} -norm of $p^t|_{t=0}$:

$$\left\| \sum_{i=0}^{\infty} c_i(*) \varphi_i(\cdot) \pi(\cdot) \right\|_{\mathbb{K}} = \lim_{t \rightarrow 0} \|p^t(*, \cdot)\|_{\mathbb{K}}.$$

By writing out the \mathbb{K} -norm, it can easily be seen that $\|p^t(*, \cdot)\|_{\mathbb{K}} = 1$ for all $t > 0$, and thus

$$\lim_{t \rightarrow 0} \|p^t(*, \cdot)\|_{\mathbb{K}} = 1.$$

With $C := \frac{\tilde{C}+1}{|\mathbb{Z}|}$, we therefore get

$$\frac{1}{|\mathbb{Z}|} \left\| \sum_{i=K+1}^{\infty} e^{\theta_i t} c_i(*) \varphi_i(\cdot) \pi(\cdot) \right\|_{\mathbb{K}} \leq C e^{\theta_{K+1} t}.$$

Hence, if we choose

$$t \geq t(\varepsilon) := \frac{1}{\theta_{K+1}} \log\left(\frac{\varepsilon}{C}\right), \quad (18)$$

then (17) is fulfilled, and the system is ε -decomposable.

Now, of course, for a decreasing decomposability tolerance ($\varepsilon \rightarrow 0$), we will need to increase the lag time ($t(\varepsilon) \rightarrow \infty$). In this situation, every system becomes trivially decomposable, see Remark 3.6. In order to claim *non-trivial* decomposability for the lag time $t(\varepsilon)$, we have to ensure that $p_D^{t(\varepsilon)}(x, \cdot)$ remains “expressive”, i.e., is not close to the identity. For this we need to show that the factors $e^{\theta_i t(\varepsilon)}$, $i = 1, \dots, K$, have not yet decayed substantially. Note that

$$\begin{aligned} e^{\theta_i t(\varepsilon)} &= e^{\frac{\theta_i}{\theta_{K+1}} \log\left(\frac{\varepsilon}{C}\right)} \\ &\geq e^{\frac{\theta_K}{\theta_{K+1}} \log\left(\frac{\varepsilon}{C}\right)} \quad \text{for } i = 1, \dots, K. \end{aligned}$$

We need to ensure that the spectral ratio $\frac{\theta_K}{\theta_{K+1}}$ decreases sufficiently quickly for $\varepsilon \rightarrow 0$. Assuming

$$\frac{\theta_K}{\theta_{K+1}} = \mathcal{O}(\log(\varepsilon)^{-1}) \quad (\varepsilon \rightarrow 0), \quad (19)$$

we have

$$e^{\theta_i t(\varepsilon)} = \mathcal{O}(1) \quad (\varepsilon \rightarrow 0),$$

and thus also, for all $x \in \mathbb{X}$,

$$\left| p_D^{t(\varepsilon)}(x, \cdot) - 1 \right| = \mathcal{O}(1) \quad (\varepsilon \rightarrow 0).$$

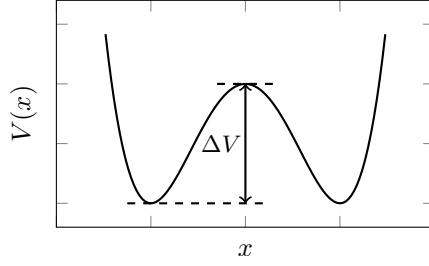


Figure 1.: One-dimensional double well potential. At low temperature, transitions across the central energy barrier are rare. The barrier height ΔV , along with the system temperature, determines the transition rate.

Example 3.15. The assumption (19) on the spectral ratio is readily fulfilled in systems with a natural spectral gap, such as systems with energy barriers. Consider the overdamped Langevin equation

$$dX_t = -\nabla V(X_t) + \sqrt{2\varepsilon} dW_t, \quad (20)$$

describing the motion of a particle in the potential V at temperature ε . Now assume that V possesses multiple local minima separated by saddle points, such as the double well potential depicted in Figure 1. Transitions out of the regions around the local minima (called *metastable sets*) are induced by the stochastic forcing in (20), with a rate governed by ε . To be specific, the expected exit rate ρ from one minimum to another across a barrier of height ΔV is given by the van't Hoff-Arrhenius law

$$\rho \sim e^{-\frac{\Delta V}{\varepsilon}},$$

where “ \sim ” indicates asymptotic equality in the sense of $\rho = C \exp(-\Delta V/\varepsilon)(1 + \mathcal{O}(\sqrt{\varepsilon} \log \varepsilon))$, where C is an ε -independent prefactor [21].

Recall that the i -th eigenvalue θ_i is the negative of the equilibration rate of the i -th slowest sub-process. At low enough temperature, the metastable transitions are hence associated with the dominant eigenvalues. The equilibration within the wells on the other hand resembles an Ornstein-Uhlenbeck process, whose eigenvalues are independent of the temperature [29]. There hence exists an index $K \geq 1$ (which depends on the number of local minima and barriers between them) such that [9]

$$\begin{aligned} \theta_i &\sim e^{-\frac{\Delta V_i}{\varepsilon}}, & i = 1, \dots, K, \\ \theta_i &\sim 1, & i \geq K + 1, \end{aligned}$$

where ΔV_i denotes the barrier height between the i -th pair of local minima. For the spectral ratio then holds

$$\frac{\theta_K}{\theta_{K+1}} \sim e^{-\frac{\Delta V_K}{\varepsilon}}.$$

For a given small enough temperature ε , the system is now non-trivially ε -decomposable on the timescale $t(\varepsilon)$ predicted by (18). On one hand, by construction of $t(\varepsilon)$, the fast processes have already decayed. On the other hand, for the slow timescales holds, for some C, \tilde{C} ,

$$e^{\theta_i t(\varepsilon)} = e^{\tilde{C} e^{-\frac{\Delta V_K}{\varepsilon}} \log(\frac{\varepsilon}{\tilde{C}})}.$$

As one can easily check,

$$\lim_{\varepsilon \rightarrow 0} e^{\tilde{C} e^{-\frac{\Delta V_K}{\varepsilon}} \log(\frac{\varepsilon}{\tilde{C}})} = 1,$$

so for small enough ε , the slow processes have not decayed yet on timescale $t(\varepsilon)$. Hence, the “natural” spectral ratio $e^{-\frac{\Delta V}{\varepsilon}}$ has a faster decay rate in ε than what is required by assumption (19).

4. A variational principle for optimal RCs

From here on, we consider the reduced dimension $r \leq n$ as well the lag time τ to be predetermined and fixed. For simplicity of notation, we will omit the lag time as parameter for the transition kernel, i.e., define $p(\cdot, \cdot) := p^\tau(\cdot, \cdot)$, $p_L(\cdot, \cdot) := p_L^\tau(\cdot, \cdot)$, $p_D(\cdot, \cdot) := p_D^\tau(\cdot, \cdot)$.

Our goal is now to find an optimal RC, i.e., a function $\xi : \mathbb{X} \rightarrow \mathbb{Z} \subset \mathbb{R}^r$ that fulfills (L) or equivalently (D) for the smallest possible $\varepsilon \geq 0$. Hence formally, we seek the minimizers of the following loss functions:

Definition 4.1 (Lumpability and decomposability loss function). *The nonlinear functional $\mathcal{L}_L : C(\mathbb{X}, \mathbb{Z}) \rightarrow \mathbb{R}^+$, defined by*

$$\mathcal{L}_L(\vartheta) := \frac{1}{|\mathbb{Z}|} \min_{p_L : \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}^+} \|p(\cdot, \cdot) - p_L(\vartheta(\cdot), \cdot)\|_{\mathbb{K}} \quad (21)$$

is called the lumpability loss function of the system.

The nonlinear functional $\mathcal{L}_D : C(\mathbb{X}, \mathbb{Z}) \rightarrow \mathbb{R}^+$, defined by

$$\mathcal{L}_D(\vartheta) := \frac{1}{|\mathbb{Z}|} \min_{p_D : \mathbb{X} \times \mathbb{Z} \rightarrow \mathbb{R}^+} \|p(\cdot, \cdot) - p_D(\cdot, \vartheta(\cdot))\pi(\cdot)\|_{\mathbb{K}} \quad (22)$$

is called the decomposability loss function of the system.

From the equivalence of lumpability and decomposability (Proposition 3.7) it follows that for every $\vartheta \in C(\mathbb{X}, \mathbb{Z})$ holds

$$\mathcal{L}_L(\vartheta) = \mathcal{L}_D(\vartheta), \quad (23)$$

hence we can find the optimal RC with respect to both (L) and (D) by solving

$$\xi := \arg \min_{\vartheta \in C(\mathbb{X}, \mathbb{Z})} \mathcal{L}_L(\vartheta), \quad (24)$$

at least in theory. Evaluating, let alone minimizing \mathcal{L}_L (or \mathcal{L}_D) would however prove difficult, due to the minimization over an infinite-dimensional function space involved in its definition (the search for the functions p_L or p_D in every step). In the following section, we therefore derive *essentially* equivalent reformulations of \mathcal{L}_L and \mathcal{L}_D that do not involve this minimization.

4.1. Differential formulation of lumpability and decomposability

The condition (L) can be interpreted as the closeness of the transition densities $p(x, \cdot)$ to some reduced reference density $p_L(\xi(x), \cdot)$. This implies that all densities $p(x, \cdot)$ whose starting points x lie on one level set of ξ are close to each other. Likewise, condition (D) can be seen as the closeness of the transition observables $p(\cdot, y)$ to some reduced reference observable $p_D(\cdot, \xi(y))\pi(y)$, which implies that all observables $p(\cdot, y)$ whose end points y lie on one level set of ξ are close to each other. This observation motivates the following “differential” characterization of lumpability and decomposability, for which we now require the smoothness assumptions on ξ and its level sets from Definition 3.1.

Definition 4.2 (Differential lumpability). *If there exists a domain $\mathbb{Z} \subset \mathbb{R}^r$ and a smooth RC $\xi \in S(\mathbb{X}, \mathbb{Z})$ such that*

$$\frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\xi(z)} \int_{\Sigma_\xi(z)} \|p(x^{(1)}, \cdot) - p(x^{(2)}, \cdot)\|_{L^1} d\mu_z(x^{(1)}) d\mu_z(x^{(2)}) dz \leq \varepsilon \quad (L')$$

is fulfilled, we say the system is differentially ε -lumpable with respect to ξ .

Definition 4.3 (Differential decomposability). *If there exists a domain $\mathbb{Z} \subset \mathbb{R}^r$ and a smooth RC $\xi \in S(\mathbb{X}, \mathbb{Z})$ such that*

$$\frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\xi(z)} \int_{\Sigma_\xi(z)} \|p(\cdot, y^{(1)})/\pi(y^{(1)}) - p(\cdot, y^{(2)})/\pi(y^{(2)})\|_{L^1_\mu} d\mu_z(x^{(1)}) d\mu_z(x^{(2)}) dz \leq \varepsilon. \quad (D')$$

is fulfilled, we say the system is differentially ε decomposable with respect to ξ .

As one easily sees using a triangle inequality argument, the differential conditions imply the original conditions, and the original conditions *almost* imply the differential conditions:

Lemma 4.4. *If the system is ε -lumpable with respect to ξ , then the system is differentially 2ε -lumpable with respect to ξ .*

Conversely, if the system is differentially ε -lumpable with respect to ξ , then the system is ε -lumpable with respect to ξ .

Proof. Assume that the system is ε -lumpable, i.e., (L) holds for some function p_L . Then we have

$$\begin{aligned} & \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \int_{\Sigma_{\xi}(z)} \left\| p(x^{(1)}, \cdot) - p(x^{(2)}, \cdot) \right\|_{L^1} d\mu_z(x^{(1)}) d\mu_z(x^{(2)}) dz \\ &= \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \int_{\Sigma_{\xi}(z)} \left\| p(x^{(1)}, \cdot) - \underbrace{p_L(z, \cdot)}_{=p_L(\xi(x^{(1)}), \cdot)} + \underbrace{p_L(z, \cdot)}_{=p_L(\xi(x^{(2)}), \cdot)} - p(x^{(2)}, \cdot) \right\|_{L^1} d\mu_z(x^{(1)}) d\mu_z(x^{(2)}) dz \\ &\leq \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \left\| p(x^{(1)}, \cdot) - p_L(\xi(x^{(1)}), \cdot) \right\|_{L^1} d\mu_z(x^{(1)}) dz + \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \left\| p(x^{(2)}, \cdot) - p_L(\xi(x^{(2)}), \cdot) \right\|_{L^1} d\mu_z(x^{(2)}) dz \\ &\leq 2\varepsilon. \end{aligned}$$

For the reverse statement, assume that (L') holds. Define

$$p_L(z, \cdot) := \int_{\Sigma_{\xi}(z)} p(x', \cdot) d\mu_z(x').$$

This p_L exists because $p(\cdot, y) \in L^1_{\mu}$ for all $y \in \mathbb{X}$.

Recall that all μ_z are probability measures on the respective $\Sigma_{\xi}(z)$. Then

$$\begin{aligned} & \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \left\| p_L(z, \cdot) - p(x, \cdot) \right\|_{L^1} d\mu_z(x) dz \\ &= \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \left\| \int_{\Sigma_{\xi}(z)} p(x', y) - p(x, y) d\mu_z(x') \right\|_{L^1} d\mu_z(x) dz \\ &\leq \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \int_{\Sigma_{\xi}(z)} \left\| p(x', \cdot) - p(x, \cdot) \right\|_{L^1} d\mu_z(x') d\mu_z(x) dz \\ &\stackrel{(L')}{\leq} \varepsilon. \end{aligned}$$

□

Lemma 4.5. *If the system is ε -decomposable with respect to ξ , then the system is differentially 2ε -decomposable with respect to ξ .*

Conversely, if the system is differentially 2ε -decomposable with respect to ξ , then the system is ε -decomposable with respect to ξ .

Proof. Assume that the system is ε -decomposable, i.e., (D) holds for some function p_D . Then we

have

$$\begin{aligned}
& \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \int_{\Sigma_{\xi}(z)} \left\| p(\cdot, y^{(1)})/\pi(y^{(1)}) - p(\cdot, y^{(2)})/\pi(y^{(2)}) \right\|_{L^1_{\mu}} d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz \\
&= \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \int_{\Sigma_{\xi}(z)} \left\| p(\cdot, y^{(1)})/\pi(y^{(1)}) - p_D(\cdot, z) + p_D(\cdot, z) - p(\cdot, y^{(2)})/\pi(y^{(2)}) \right\|_{L^1_{\mu}} d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz \\
&\leq \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \left\| p(\cdot, y^{(1)})/\pi(y^{(1)}) - p_D(\cdot, \xi(y^{(1)})) \right\|_{L^1_{\mu}} d\mu_z(y^{(1)}) dz \\
&\quad + \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \left\| p(\cdot, y^{(2)})/\pi(y^{(2)}) - p_D(\cdot, \xi(y^{(2)})) \right\|_{L^1_{\mu}} d\mu_z(y^{(2)}) dz \\
&= \frac{2}{|\mathbb{Z}|} \int_{\mathbb{X}} \|p(\cdot, y) - p_D(\cdot, \xi(y))\pi(y)\|_{L^1_{\mu}} dy \\
&= \frac{2}{|\mathbb{Z}|} \left\| \|p(*, \cdot) - p_D(*, \cdot)\pi(\cdot)\|_{L^1} \right\|_{L^1_{\mu}} \stackrel{\text{(D)}}{\leq} 2\varepsilon.
\end{aligned}$$

For the reverse statement, define

$$p_D(\cdot, z) := \int_{\Sigma_{\xi}(z)} p(\cdot, y')/\pi(y') d\mu_z(y').$$

This p_D exists because $p(x, \cdot) \in L^1$ for all $x \in \mathbb{X}$. Then

$$\begin{aligned}
\frac{1}{|\mathbb{Z}|} \left\| \|p(*, \cdot) - p_D(*, \xi(\cdot))\pi(\cdot)\|_{\mathbb{K}} \right\|_{\mathbb{K}} &= \frac{1}{|\mathbb{Z}|} \left\| \|p(*, \cdot) - p_D(*, \xi(\cdot))\pi(\cdot)\|_{L^1} \right\|_{L^1_{\mu}} \\
&= \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \left\| p_D(\cdot, z) - p(\cdot, y)/\pi(y) \right\|_{L^1_{\mu}} d\mu_z(y) dz \\
&= \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \int_{\mathbb{X}} \left| \int_{\Sigma_{\xi}(z)} p(x, y')/\pi(y') - p(x, y)/\pi(y) d\mu_z(y') \right| d\mu(x) d\mu_z(y) dz \\
&\leq \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\xi}(z)} \int_{\Sigma_{\xi}(z)} \left\| p(\cdot, y')/\pi(y') - p(\cdot, y)/\pi(y) \right\|_{L^1_{\mu}} d\mu_z(y') d\mu_z(y) dz \\
&\stackrel{\text{(D)'}}{\leq} \varepsilon.
\end{aligned}$$

□

4.2. Differential loss functions

We can now define *differential* versions of the loss functions \mathcal{L}_L and \mathcal{L}_D in which the hard-to-identify terms p_L and p_D no longer appear:

Definition 4.6 (differential lumpability and differential decomposability loss function). *The nonlinear functional $\widetilde{\mathcal{L}}_L : S(\mathbb{X}, \mathbb{Z}) \rightarrow \mathbb{R}^+$, defined by*

$$\widetilde{\mathcal{L}}_L(\vartheta) := \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\vartheta}(z)} \int_{\Sigma_{\vartheta}(z)} \left\| p(x^{(1)}, \cdot) - p(x^{(2)}, \cdot) \right\|_{L^1} d\mu_z(x^{(1)}) d\mu_z(x^{(2)}) dz \quad (25)$$

is called the differential lumpability loss function of the system.

The nonlinear functional $\widetilde{\mathcal{L}}_D : S(\mathbb{X}, \mathbb{Z}) \rightarrow \mathbb{R}^+$, defined by

$$\widetilde{\mathcal{L}}_D(\vartheta) := \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\vartheta}(z)} \int_{\Sigma_{\vartheta}(z)} \left\| p(\cdot, y^{(1)})/\pi(y^{(1)}) - p(\cdot, y^{(2)})/\pi(y^{(2)}) \right\|_{L^1_{\mu}} d\mu_z(x^{(1)}) d\mu_z(x^{(2)}) dz \quad (26)$$

is called the differential decomposability loss function of the system.

Note that, unlike \mathcal{L}_L and \mathcal{L}_D , $\widetilde{\mathcal{L}}_L$ and $\widetilde{\mathcal{L}}_D$ are in general neither identical to each other, nor to \mathcal{L}_L . The following result characterizes the relationship between the different loss functions:

Corollary 4.7. *For any $\vartheta \in S(\mathbb{X}, \mathbb{Z})$ holds*

$$\mathcal{L}_L(\vartheta) \leq \widetilde{\mathcal{L}}_L(\vartheta) \leq 2\mathcal{L}_L(\vartheta),$$

and

$$\mathcal{L}_L(\vartheta) \leq \widetilde{\mathcal{L}}_D(\vartheta) \leq 2\mathcal{L}_L(\vartheta).$$

Proof. The first pair of inequalities follows directly from Lemma 4.4. The second pair of inequalities follows from Lemma 4.5 and (23). \square

Hence, for arbitrary (non-optimal) RCs ϑ , we cannot expect $\widetilde{\mathcal{L}}_L(\vartheta)$ or $\widetilde{\mathcal{L}}_D(\vartheta)$ to be similar to $\mathcal{L}_L(\vartheta)$. However, under the assumption that the system is indeed ε -lumpable for small ε , we can expect their *minima* to be similar:

Corollary 4.8. *Let ξ be an optimal RC, defined by (24), and set $\varepsilon := \mathcal{L}_L(\xi)$. Let ξ_L and ξ_D be minimizers of $\widetilde{\mathcal{L}}_L$ and $\widetilde{\mathcal{L}}_D$, respectively. Then*

$$\widetilde{\mathcal{L}}_L(\xi) \leq 2\varepsilon \quad \text{and} \quad \widetilde{\mathcal{L}}_D(\xi) \leq 2\varepsilon.$$

Moreover,

$$\mathcal{L}_L(\xi_L) \leq 2\varepsilon \quad \text{and} \quad \mathcal{L}_L(\xi_D) \leq 2\varepsilon.$$

Proof. We show the assertions for $\widetilde{\mathcal{L}}_L$. For $\widetilde{\mathcal{L}}_D$, the proof is identical.

The first inequality follows from applying Corollary 4.7 to $\vartheta = \xi$. The second inequality then follows from

$$\mathcal{L}_L(\xi_L) \stackrel{\text{Cor. 4.7}}{\leq} \widetilde{\mathcal{L}}_L(\xi_L) \stackrel{\text{Def. } \xi_L}{\leq} \widetilde{\mathcal{L}}_L(\xi) \leq 2\varepsilon.$$

\square

Remark 4.9. The above variational principle implies that a minimizer η to $\widetilde{\mathcal{L}}_L$ or $\widetilde{\mathcal{L}}_D$ is not necessarily a strict minimizer of \mathcal{L}_L . However, the difference between the minima will be at most ε . In other words, the system will be 2ε -lumpable and -decomposable with respect to η . Thus, for practical purposes, we can expect

$$\xi_L := \arg \min_{\vartheta \in S(\mathbb{X}, \mathbb{Z})} \widetilde{\mathcal{L}}_L(\vartheta) \tag{27}$$

and

$$\xi_D := \arg \min_{\vartheta \in S(\mathbb{X}, \mathbb{Z})} \widetilde{\mathcal{L}}_D(\vartheta) \tag{28}$$

to be “quasi-optimal” RCs.

5. Numerical approximation of the loss function

In order to solve the above optimization problems, the loss functions $\widetilde{\mathcal{L}}_L$ and $\widetilde{\mathcal{L}}_D$, and, depending on the optimization scheme, also their gradients, need to be evaluated numerically for candidate RCs ϑ . The presumed high dimension n of the state space however presents several challenges, which are discussed in the following.

5.1. Sampling requirements of the decomposability loss function

Note that at the core of both $\widetilde{\mathcal{L}}_L$ and $\widetilde{\mathcal{L}}_D$ lies the evaluation of the transition density p^t at certain start and end points. The function p^t is however not known analytically in practice, and must be estimated empirically by simulations of the dynamics. A critical question is thus how many of these estimates are required in order to approximate $\widetilde{\mathcal{L}}_L$ or $\widetilde{\mathcal{L}}_D$ up to a given tolerance. In particular, an exponential dependency of that number on the dimension n would be fatal. We will now show that, if the system is indeed highly lumpable with respect to some RC ξ (which in general is different from the candidate RC ϑ), the differential decomposability loss function $\widetilde{\mathcal{L}}_D$ can be approximated very efficiently. Crucially, ξ does not have to be known, its implied existence is sufficient to guarantee the efficiency.

By expanding the $\|\cdot\|_{L^1_\mu}$ norm in (26), we can write $\widetilde{\mathcal{L}}_D$ as

$$\widetilde{\mathcal{L}}_D(\vartheta) = \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\vartheta}(z)} \int_{\Sigma_{\vartheta}(z)} \int_{\mathbb{X}} \left| \frac{p(x, y^{(1)})}{\pi(y^{(1)})} - \frac{p(x, y^{(2)})}{\pi(y^{(2)})} \right| \pi(x) dx d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz. \quad (29)$$

or, by changing the integration order, as

$$= \frac{1}{|\mathbb{Z}|} \int_{\mathbb{X}} \left(\int_{\mathbb{Z}} \int_{\Sigma_{\vartheta}(z)} \int_{\Sigma_{\vartheta}(z)} \left| \frac{p(x, y^{(1)})}{\pi(y^{(1)})} - \frac{p(x, y^{(2)})}{\pi(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz \right) \pi(x) dx. \quad (30)$$

Exact evaluation of the outermost integral in (30) would require analytical knowledge of the transition kernel p . Specifically, it would require knowledge of *all* transition densities, i.e., of $p(x, \cdot)$ for *all* starting points $x \in \mathbb{X}$, which we cannot assume in practice. We can however assume that a sufficiently precise approximation of a certain *limited* number of transition densities $p(x^{(k)}, \cdot)$, $k = 1, \dots, M$ can be obtained. This approximation would typically be realized by parallel simulation of the stochastic dynamics starting from $x^{(k)}$, thus creating a point cloud that is distributed according to $p(x^{(k)}, \cdot)$. One can then apply density estimation techniques such as kernel- or spectral density estimation to obtain an analytic expression of $p(x^{(k)}, \cdot)$. In any realistic scenario, we would however not be able to approximate $p(x^{(k)}, \cdot)$ for a uniform covering of \mathbb{X} by a sufficient amount of points, due to the typically high dimension of \mathbb{X} . For example, a regular grid-based covering of the unit cube $[0, 1]^n$ with spacing $\delta > 0$ would require $M = (\frac{1}{\delta})^n$ grid points. This exponential dependence of M on the dimension renders classical numerical quadrature schemes infeasible, leaving Monte Carlo (MC) quadrature as the only viable option for solving the integral over \mathbb{X} .

To quantify the MC error in M for this integral, define

$$f(x) := \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_{\vartheta}(z)} \int_{\Sigma_{\vartheta}(z)} \left| \frac{p(x, y^{(1)})}{\pi(y^{(1)})} - \frac{p(x, y^{(2)})}{\pi(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz, \quad (31)$$

$$I(f) := \int_{\mathbb{X}} f(x) d\mu(x), \quad (32)$$

$$I_M(f) := \frac{1}{M} \sum_{i=1}^M f(x^{(i)}), \quad x^{(i)} \sim \mu \text{ i.i.d.} \quad (33)$$

For the expected approximation error then holds [25, p. 77]

$$\mathbb{E}[|I(f) - I_M(f)|] = \frac{\text{Var}_\mu(f)}{\sqrt{M}}, \quad (34)$$

where $\text{Var}_\mu(f)$ denotes the variance of f with respect to μ , defined by

$$\text{Var}_\mu(f) = \mathbb{E}_\mu[f^2] - (\mathbb{E}_\mu[f])^2.$$

The independence of the convergence rate $1/\sqrt{M}$ from the dimension is what gives MC methods an edge above conventional methods, at least in theory. However, the *effective* convergence speed, and hence the required number of start points $x^{(k)}$, is highly influenced by the prefactor $\text{Var}_\mu(f)$.

We will show now that for highly lumpable systems, $\text{Var}_\mu(f)$ tends to be substantially smaller than for non-lumpable systems. The intuitive explanation is that, for systems that are lumpable with respect to some RC ξ , f is essentially constant along every level set of ξ , hence only the variation of f along ξ contributes to $\text{Var}_\mu(f)$. This holds even in the candidate RC ϑ that appears in the definition of f is not close to ξ . The intuition is formalized by the following theorem:

Theorem 5.1. *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*

$$f_L(z) := \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\vartheta(z')} \int_{\Sigma_\vartheta(z')} \left| \frac{p_L(z, y^{(1)})}{\pi(y^{(1)})} - \frac{p_L(z, y^{(2)})}{\pi(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz'. \quad (35)$$

Furthermore, define the effective invariant density by

$$\bar{\pi}(z) := \int_{\Sigma_\xi(z)} \pi(x) d\sigma_z(x) \quad (36)$$

where σ_z denotes the surface measure on $\Sigma_\xi(z)$. Then there exists a constant $C > 0$ such that

$$|\text{Var}_\mu(f) - \text{Var}_\mu(f_L \circ \xi)| \leq 2(1 + C^2) \|f_L \circ \xi\|_{L^1_\mu} \|\bar{\pi}\|_\infty \varepsilon + \mathcal{O}(\varepsilon^2).$$

In words, the variance of f is ε -close to the variance of $f_L \circ \xi$. Also, the variance of $f_L \circ \xi$ is equal to the variance of f_L :

Lemma 5.2. *Let $\bar{\mu}$ be defined as in Lemma 5.1 and $h \in L^1_{\bar{\mu}}(\mathbb{Z})$. Then*

$$\text{Var}_\mu(h \circ \xi) = \text{Var}_{\bar{\mu}}(h).$$

The proof of Theorem 5.1 and Lemma 5.2 can be found in Appendix B.

Remark 5.3. To summarize, 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 n through the dimension of its argument.

Looking at the definition of f_L , one sees that the variance of the effective transition density $p_L(z, y)$ in its first argument influences $\text{Var}_{\bar{\mu}}(f_L)$. Note, however, that the dimension n also indirectly appears in the definition of f_L , through the integrations over the $(n - r)$ -dimensional level sets $\Sigma_\vartheta(z)$. As the candidate RC ϑ can be arbitrarily complex, it is very hard to give any general estimates on the influence of its level sets and n on $\text{Var}_{\bar{\mu}}(f_L)$. However, we found no argument why $\text{Var}_{\bar{\mu}}(f_L)$ should *increase* with increasing n in general. Indeed, in Section 6.1, we present an example system for which $\text{Var}_{\bar{\mu}}(f_L)$ is inversely correlated to n . There, we also numerically confirm the predicted dependence of the Monte Carlo error on $\text{Var}_\mu(f)$, and the predicted dependence of $\text{Var}_\mu(f)$ on ε .

5.2. Further numerical challenges

In order to solve the optimization problem (28) in practice, several more numerical challenges will need to be overcome. These challenges do however not involve the collection of dynamical data, i.e., the expensive realization of the dynamical system, and hence can be considered peripheral. We will therefore only briefly sketch possible solutions to these challenges, and leave the exact elaboration to future work.

Sampling the invariant measure

In high dimensions, generating the μ -distributed samples $x^{(k)}$ in (33) by classical techniques such as importance sampling is notoriously inefficient [1]. As an alternative, Markov Chain Monte Carlo techniques [48], such as Metropolis-Hastings or Gibbs Sampling can be used. Due to their Markov chain fulfilling the detailed balance condition w.r.t. π , their generated samples are precisely μ -distributed, so the variance estimates derived in Section 5.1 hold. The same is true for Boltzmann generators [40], which in addition overcome the problem of sampling long-lived states.

Discretization of the solution space

The solution space $S(\mathbb{X}, \mathbb{Z})$ for ξ_D in (27) needs to be replaced by a finite-dimensional parametric ansatz space. While in principle classical approaches like Galerkin or grid-based finite element discretization methods could be explored, these methods suffer heavily from the curse of dimensionality, that is, the exponential dependence of the number of parameters on the system dimension [41].

Due to its prevalence in scientific computing, many approaches have been suggested to tackle the curse of dimensionality, including sparse grids [22], mesh-free methods [20], and non-parametric approaches [53]. All of these methods are in principle compatible with the task of evaluating $\widetilde{\mathcal{L}}_D(\vartheta)$, provided the level sets $\Sigma_\vartheta(z)$ can be sampled for the discretized ϑ .

However, one particular discretization method practically suggests itself for our overall task of finding the minimizer of \mathcal{L}_D , and that is the representation of ξ_D via a multilayer neural network [34]. These models have demonstrated impressive performance for problems of high input dimensions, such as image recognition [33], protein structure prediction [56] or, quite relevant to our task, Markov model construction via the approximation of transfer operator eigenfunctions [36]. Minimization of the loss function $\widetilde{\mathcal{L}}_D$ can then be accomplished by gradient descent methods, of which there exist numerous efficient variants and implementations [51].

Level set integrals

Besides the integral over \mathbb{X} , computation of $\widetilde{\mathcal{L}}_D(\vartheta)$ requires the numerical solution of an integral over \mathbb{Z} and, for each $z \in \mathbb{Z}$, two surface integrals over the z -level set $\Sigma_\vartheta(z)$ of ϑ , see (30). As the dimension r of \mathbb{Z} was assumed to be small, classical grid-based quadrature methods are sufficient for the former. The level sets $\Sigma_\vartheta(z)$, however, are $n - r$ dimensional nonlinear sub-manifolds of \mathbb{X} , and thus too high-dimensional for grid based methods.

The canonical approach would be again to use Monte Carlo quadrature, like for the integral over \mathbb{X} . This approach was used for the low-dimensional examples presented in Section 6. Sampling of the surface measure μ_z can be accomplished by restricted simulation of the dynamics, for which numerical schemes were suggested in [10]. Moreover, if ϑ is modeled by a neural network, sampling $\Sigma_\vartheta(z)$ could be alternated with the optimization step, leading to an optimization scheme similar to stochastic gradient descent.

6. Numerical examples

6.1. Timescale-separated Brownian motion

We consider a process $(X_t) = (X_t^{(1)}, \dots, X_t^{(n)})$ on the n -dimensional torus $\mathbb{T}^n := [-\pi, \pi]^n$ for $n \geq 2$ with slow diffusion in the first coordinate direction (standard Brownian motion) and instantaneous and pairwise independent equilibrations to the uniform distribution in the $n - 1$ remaining coordinate components. We can write this system in SDE form as

$$\begin{aligned} dX_t^{(1)} &= dW_t, \\ X_t^{(i)} &\sim U([-\pi, \pi]), \quad 2 \leq i \leq n, \end{aligned}$$

such that

$$\begin{aligned} X_{t_1}^{(i)} &\perp X_{t_2}^{(i)}, \quad 1 \leq i \leq n, t_1 \neq t_2, \\ X_{t_1}^{(i)} &\perp X_{t_2}^{(j)} \quad 1 \leq i \neq j \leq n, \forall t_1, t_2, \end{aligned}$$

where W_t denotes a standard Brownian motion, U denotes the uniform distribution, and \perp marks independence. We introduce a second process (\widetilde{X}_t) on \mathbb{T}^n which is parametrized by a variance

$\sigma^2 > 0$ by the global diffusion

$$\begin{aligned} d\tilde{X}_t^{(1)} &= d\tilde{W}_t^{(1)} \\ d\tilde{X}_t^{(i)} &= \sigma d\tilde{W}_t^{(i)}, \quad 2 \leq i \leq n, \end{aligned}$$

where the $\tilde{W}_t^{(i)}, 1 \leq i \leq n$ are pairwise independent standard Brownian motions.

6.1.1. Transition densities

The transition densities of (X_t) and (\tilde{X}_t) can be conveniently described in terms of the one-dimensional *wrapped normal distribution* [24] on the circle $[-\pi, \pi]$ with variance $\sigma^2 \geq 0$, which is defined by the density $g^\sigma : [-\pi, \pi] \times [-\pi, \pi] \rightarrow \mathbb{R}_+$ given by

$$g^\sigma(x, y) := \frac{1}{2\pi} \left(1 + 2 \sum_{k=1}^{\infty} \rho^{k^2} \cos(k(y-x)) \right), \text{ where } \rho = \exp(-\sigma^2/2).$$

In particular, for $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n) \in \mathbb{T}^n$ we straightforwardly obtain the transition density of (X_t) as

$$p^{\tau, \infty}(x, y) := \frac{1}{(2\pi)^{n-1}} g^\tau(x_1, y_1)$$

as well as the transition density of (\tilde{X}_t) as

$$p^{\tau, \sigma}(x, y) := g^\tau(x^{(1)}, y^{(1)}) \prod_{i=2}^n g^{\tau\sigma}(x_i, y_i). \quad (37)$$

The transition densities are illustrated in Figure 2 for the case $n = 2$.

6.1.2. Lumpability

As σ increases, the transition density $p^{\tau, \sigma}(x, y)$ uniformly approximates $p^{\tau, \infty}(x, y)$, we may therefore understand (X_t) as the limiting process of (\tilde{X}_t) for large variances. Because the transition density of $p^{\tau, \infty}(x, y)$ only depends on the first coordinate components x_1, y_1 , we can verify lumpability of (\tilde{X}_t) for large enough σ , as we will briefly illustrate now.

Consider the RC $\xi : [-\pi, \pi]^n \rightarrow [-\pi, \pi]$ given by the orthogonal projection onto the first coordinate component $\xi((x_1, x_2, \dots, x_n) = x_1$. We call ξ the *optimal* RC. With the effective transition density

$$p_L^\tau(z, y) := \frac{1}{(2\pi)^{n-1}} g^\tau(z, y_1),$$

one can then show that

$$\|p^{\tau, \sigma}(*, \cdot) - p_L^\tau(\xi(*), \cdot)\|_{\mathbb{K}} = \mathcal{O}(\exp(-(\tau\sigma)^2/2)) \quad (38)$$

for both $\sigma \rightarrow \infty$ or $\tau \rightarrow \infty$. The derivation of (38) can be found in Appendix C. In particular, for any lag time $\tau > 0$, we will find a σ for which the system becomes arbitrarily lumpable, while in the limit $\sigma \rightarrow \infty$, the system is “perfectly” lumpable (i.e., $\|p^{\tau, \sigma}(*, \cdot) - p_L^\tau(\xi(*), \cdot)\|_{\mathbb{K}} = 0$) for every $\tau > 0$. Note that ε -lumpability of (\tilde{X}_t) also implies ε -decomposability due to Proposition 3.7.

As our estimates of the distance $\|p^{\tau, \sigma}(*, \cdot) - p_L^\tau(\xi(*), \cdot)\|_{\mathbb{K}}$ in Appendix C are asymptotic in nature, we can give no precise formula for the dependence of the minimal ε in (L) on σ . In particular, we cannot predict the dependence of ε on the system dimension n . Nevertheless, when computing the distance numerically, its dependence on n appears to be moderate, and seems to diminish with growing n (Figure 3). Hence, we can expect a high degree of lumpability for moderately large σ and moderately high dimensions. Also, the computations confirm the convergence rate predicted in (38).

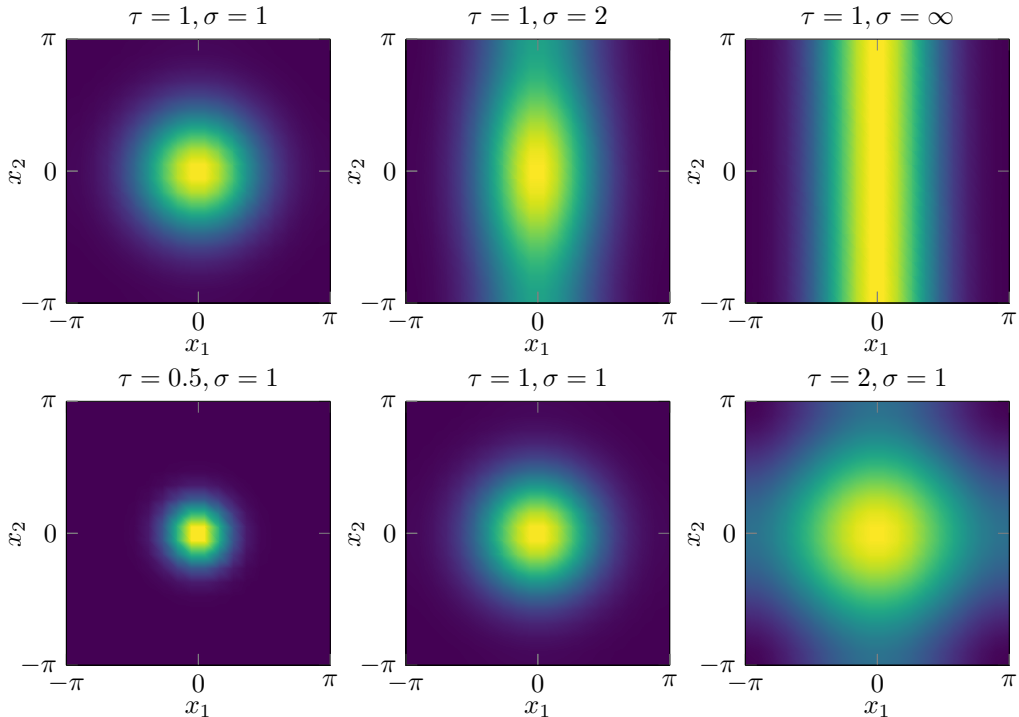


Figure 2.: Illustration of transition densities $p^{\tau, \sigma}(0, \cdot)$ for different lag times τ and standard deviations σ in the case $n = 2$. The top row illustrates how the transition densities $p^{\tau, \sigma}(x, \cdot)$ approximate $p^{\tau, \infty}(x, \cdot)$ for an increasing standard deviation σ when τ is fixed. Note that for fixed x_1 coordinates, the limiting density $p^{\tau, \infty}(x, \cdot)$ is constant along the x_2 -coordinate. The bottom row illustrates the larger degree of global dispersion for increasing lag times when σ is fixed.

6.1.3. Loss function illustration

Next, we investigate the dependence of the loss function $\widetilde{\mathcal{L}}_{\text{D}}$ on the RC, for the two-dimensional process (one slow and one fast component). To be precise, we investigate the dependence of $\widetilde{\mathcal{L}}_{\text{D}}(\vartheta_{\alpha})$ on the parameter α of the family of “test” reaction coordinates

$$\vartheta_{\alpha}(x) := \frac{1}{\pi(\cos |\alpha| + \sin |\alpha|)} (\cos \alpha \quad \sin \alpha) \cdot \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \quad \alpha \in (-\pi, \pi). \quad (39)$$

for different values of the diffusion constant σ . The level sets of ϑ_{α} are one-dimensional hyperplanes that intersect the x_1 -axis with angle $\pi/2 - \alpha$. The prefactor in (39) ensures that $\mathbb{Z} = \text{range}(\vartheta_{\alpha}) = [-1, 1]$ for all α . Figure 4 illustrates ϑ_{α} and its level sets.

We use a combination of symbolic computation and numerical quadrature techniques for computing $\widetilde{\mathcal{L}}_{\text{D}}(\vartheta_{\alpha})$, using Mathematica and Matlab. In particular, for sampling points on the level sets $\Sigma_{\vartheta_{\alpha}}(z)$, required for numerically computing the two innermost integrals in (30), we utilized the Matlab function `randFixedLinearCombination`, written by John D’Errico and provided through MatlabCentral [13]. The scripts containing our computations are provided in the SI. We do however *not* suggest to use these scripts for the computation and minimization of $\widetilde{\mathcal{L}}_{\text{D}}$ in real-world systems, as in particular the integration over the $(n - r)$ -dimensional level sets quickly become infeasible. See the discussion in Section 5.2 for a first outlook on how we plan to solve the optimization problem in practice.

The differential decomposability loss function $\widetilde{\mathcal{L}}_{\text{D}}(\vartheta_{\alpha})$ in dependence of α is shown in Figure 5. We observe that, for $\sigma = 2$ and $\sigma = \infty$, $\widetilde{\mathcal{L}}_{\text{D}}(\vartheta_{\alpha})$ indeed takes its unique global minimum for $\alpha = 0$, i.e., the optimal RC $\vartheta_0(x) = \xi(x) = x_1$. We also observe that for $\sigma = 2$, the minimum value is

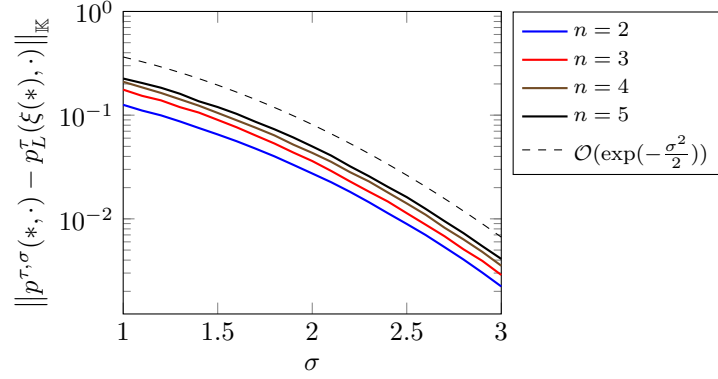


Figure 3.: \mathbb{K} -distance of the transition density $p^{\tau, \sigma}$ to the effective transition density p_L^τ , indicating the degree of lumpability. We observe the predicted decay in the diffusion constant σ , as well as a relative insensitivity to the system dimension n .

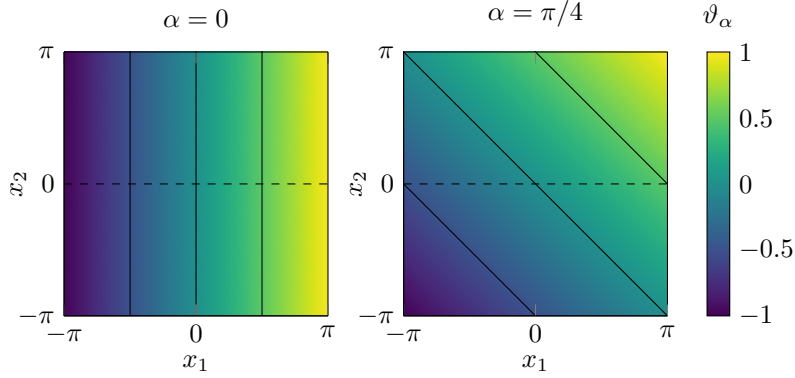


Figure 4.: The linear RC ϑ_α for two different values of α .

not zero, as the system is not “perfectly lumpable” with respect to ϑ_0 , whereas for $\sigma = \infty$ it is. Further, the “worst” RCs $\vartheta_{\pm\pi/2}(x) = \pm x_2$, which project onto the system’s fast instead of its slow coordinate, correctly get assigned the global maximum value. For $\alpha = 1$, on the other hand, the diffusion is isotropic, and hence the “slow directions” coincide with the directions of largest extent, which are the diagonals of the domain $[-\pi, \pi]^2$. Consequently, $\widetilde{\mathcal{L}}_D(\vartheta_\alpha)$ becomes minimal for $\alpha = \pm\pi/4$, for which ϑ_α projects onto the diagonals $\{x_1 = \pm x_2\}$.

The behavior of $\widetilde{\mathcal{L}}_D$ is therefore perfectly consistent with our intuition, and we can expect to identify the optimal RC by minimizing $\widetilde{\mathcal{L}}_D$.

6.1.4. Loss function Monte Carlo error

Finally we investigate the Monte Carlo quadrature error for the integral over \mathbb{X} in $\widetilde{\mathcal{L}}_D$. As discussed in Section 5.1, this error is, besides the number of sample points M , determined by the variance of the integrand f , which for the current system takes the form

$$f(x) = \frac{(2\pi)^n}{|\vartheta_{\max} - \vartheta_{\min}|} \int_{\vartheta_{\min}}^{\vartheta_{\max}} \int_{\Sigma_{\vartheta}(z)} \int_{\Sigma_{\vartheta}(z)} \left| p^{\tau, \sigma}(x, y^{(1)}) - p^{\tau, \sigma}(x, y^{(2)}) \right| d\sigma_z(y^{(1)}) d\sigma_z(y^{(2)}) dz,$$

where ϑ_{\max} and ϑ_{\min} describe the maximum and minimum of ϑ . We point out in particular that f depends on the diffusion coefficient σ and, through the RC ϑ , on the system dimension n . The influence of these two parameters on the Monte Carlo error is the primary subject of investigation for this section.

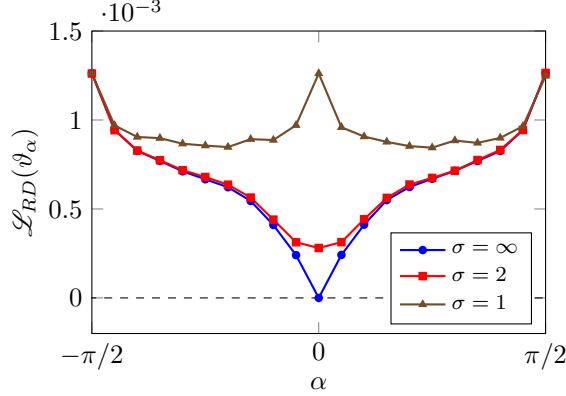


Figure 5.: The differential decomposability loss function of the RC ϑ_α for $\alpha \in (-\pi/2, \pi/2)$ and different values of σ .

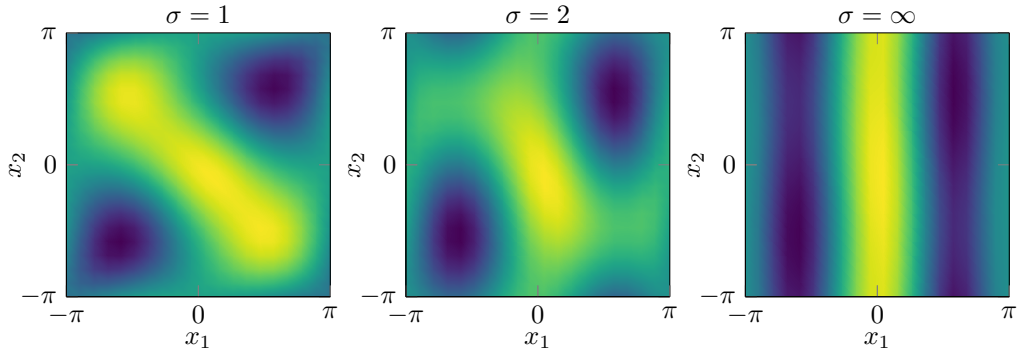


Figure 6.: The function f , i.e., the integrand of the integral over \mathbb{X} in $\widetilde{\mathcal{L}}_D(\vartheta)$ for $\vartheta(x) = x_1 + x_2$ and different values of σ .

We consider in dimension 2 and 3, respectively, the test RCs

$$\vartheta^2(x) = x_1 + x_2, \quad \vartheta^3(x) = x_1 + x_2 + x_3.$$

Figure 6 shows f for the RC ϑ^2 and different values of σ . We observe that with increasing σ , f indeed becomes increasingly constant on the level sets $\{x_1 = z\}$ of the optimal RC ξ . This behavior was predicted in Section 5.1.

Figure 7 illustrates the variance of f . We first note that the variance indeed decreases for increasing σ , and converges towards the variance of the process with instantaneously equilibrating components x_2, \dots, x_n (equivalent to choosing “ $\sigma = \infty$ ”). Moreover, we observe that the variance of the three-dimensional process is substantially smaller compared to the two-dimensional process. This observation still holds when considering the relative variance $\text{Var}[f]/\mathbb{E}[f]$. This demonstrates that higher-dimensional function do not per se possess higher variance. In fact, $\text{Var}[f]$ appears to be more dependent on the choice of the particular RC ϑ than on n , although, to avoid digression, we will refrain from detailed analysis.

Finally, we estimate the relative expected MC error

$$\frac{\mathbb{E}[|I(f) - I_M(f)|]}{I(f)}$$

where $I(f)$ and $I_M(f)$ are the exact integral and Monte Carlo integral with M samples of f defined in (32) and (33). In practice, this error indicates how many “dynamical samples” $p^{\tau, \sigma}(x^{(k)}, \cdot)$ need

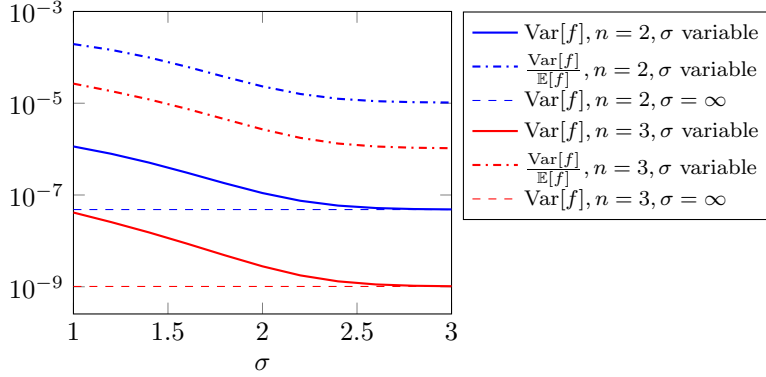


Figure 7.: (Relative) variance of f associated in dimensions 2 and 3 for various σ . We observe convergence of $\text{Var}[f]$ with variable σ towards $\text{Var}[f]$ associated with the σ -independent limit process X_t .

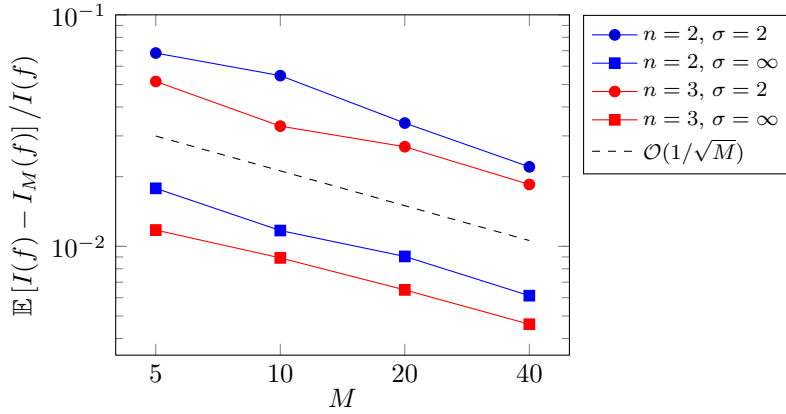


Figure 8.: The relative Monte Carlo quadrature error for the integral over \mathbb{X} in $\widetilde{\mathcal{L}}_D(\vartheta_w)$ for different dimensions n and different diffusion coefficients σ .

to be created by numerical simulation in order to approximate $\widetilde{\mathcal{L}}_D(\vartheta)$ up to a given accuracy. In the present example, however, $p^{\tau \cdot \sigma}$ is known analytically by (37).

Figure 8 shows the relative error in dependence on the sample size M for ϑ^2 and ϑ^2 , each for a finite value of σ as well as $\sigma = \infty$. In all four cases we observe the expected Monte Carlo convergence rate of $\mathcal{O}(1/\sqrt{M})$. Moreover, the error is significantly smaller for higher σ , i.e., the more lumpable system. Finally, the expected error for the higher dimensional system is slightly smaller. All these phenomena are perfectly consistent with the preceding analysis of $\text{Var}[f]$.

We conclude from this example that for lumpable systems, we can expect to find the optimal RC by minimising $\widetilde{\mathcal{L}}_D$, that the Monte Carlo approximation to $\widetilde{\mathcal{L}}_D(\vartheta)$ requires only few dynamical samples for highly lumpable systems, and that a high dimension of the base system has no negative impact on the performance.

6.2. Metastable circular system

To demonstrate the behavior of the loss function for nonlinear optimal RCs, we consider as a second example a two-dimensional system governed by the overdamped Langevin dynamics

$$dX_t = -\nabla V(X_t) + \sqrt{2\beta^{-1}}dW_t, \quad (40)$$

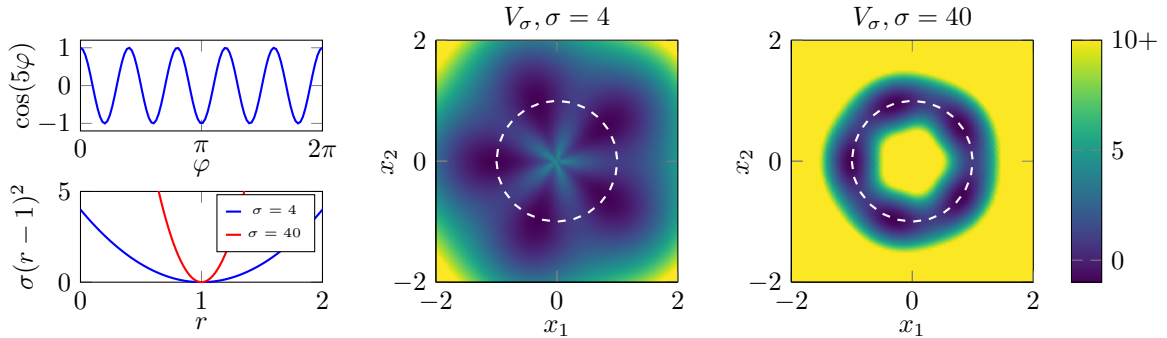


Figure 9.: Left: The two components of the potential, which depend on the angular and the radial part of the polar coordinates, respectively. Center & right: The circular potential and its components for different values of σ . The dashed white line indicates the unit circle, i.e., the minimal energy pathway connecting the local minima.

where $V : \mathbb{R}^2 \rightarrow \mathbb{R}$ is the potential energy surface and $\beta > 0$ is the inverse temperature determining the strength of the Brownian motion W_t . Informally, movement of this system can be described as a random walk within the energy landscape, aiming in the direction of steepest descent of V but being disturbed by temperature-scaled white noise.

In particular, we here consider the family of potentials

$$V_\sigma(x) = \cos(5\varphi(x)) + \sigma(r(x) - 1)^2, \quad \sigma > 0,$$

where $(\varphi(x), r(x))$ describe the polar coordinates of x , i.e.,

$$\varphi(x) = \text{atan2}(x), \quad r(x) = \|x\|_2.$$

The potential consists of two components: a cosine term in the angular coordinate with five local minima of equal depth, and a quadratic term in the radial coordinate with a single minimum at $r = 1$. The full potential, shown in Figure 9, therefore possesses five local minima arranged along the unit circle. The two components φ and r as functions on \mathbb{X} are shown in Figure 11 (left).

6.2.1. Metastability analysis

At moderate temperatures, the five local minima of V induce *metastability* in the system. This means that a typical trajectory will vibrate around a minimum for a long time, until suddenly, induced by sufficiently strong stochastic excitation, undergo a rapid transition to another local minimum. If additionally the equilibration in the radial direction is sufficiently fast, i.e., the parameter σ is sufficiently large, then these rare transitions will be the slowest sub-processes of the system. To confirm this, we compute the leading eigenvalues and associated eigenfunctions of the system's transfer operator $\mathcal{P}^t : L^2(\mathbb{X}) \rightarrow L^2(\mathbb{X})$, which is defined by

$$\mathcal{P}^t f(x) = \int_{\mathbb{X}} f(y) p^t(y, x) dy.$$

Going back to the late 90s, spectral analysis of the transfer operator, or its adjoint, the Koopman operator, forms the basis of many model reduction methods that aim to preserve the system's long timescales, for an overview see [28].

The eigenvalues of \mathcal{P}^t for $t = 0.1$ and various values of σ are shown in Figure 10. We see that for $\sigma = 10$ and $\sigma = 100$, there is a significant *spectral gap* after λ_5 , indicating a significant time scale separation between the associated sub-processes. Analysing the sign structure of the associated eigenfunctions confirms that the slowest processes are indeed associated with the transitions between the metastable sets (see the SI).

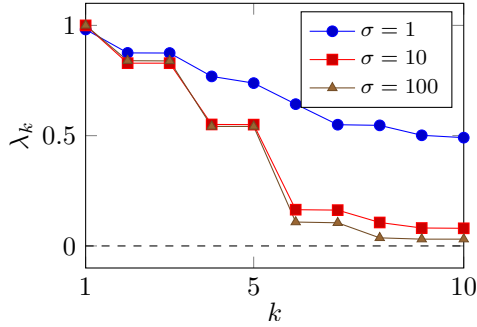


Figure 10.: Leading eigenvalues of the circular system’s transfer operator for different values of σ . For larger σ , we observe a spectral gap, which indicates that the equilibration times of the metastable transitions are much slower than that of other sub-processes in the system.

For small σ , on the other hand, no spectral gap can be observed, hence the metastable transitions are not considerably slower than the remaining processes. Indeed, sign analysis of the sixth eigenfunction for $\sigma = 1$ shows that the corresponding sub-process describes the equilibration in the radial direction (see the SI). As λ_6 is not significantly separated from λ_5 , the radial equilibration occurs on roughly the same timescale as the metastable transitions.

6.2.2. Lossfunction computation

The preceding analysis confirms that, for high enough σ , a good RC should resolve the transitions between the metastable sets. As these transitions occur predominantly within certain “transition channels” that surround the minimum energy pathways (MEPs) [38], and as in our system the MEPs between two neighbouring minima are segments of the unit circle, we would expect the angular component $\varphi(x)$ of the polar coordinates of x to be a good RC. Conversely, we would expect the radial component $r(x)$ to be a bad RC, especially for high values of σ .

To test this hypothesis, we now compute the differential lumpability loss function $\widetilde{\mathcal{L}}_{\mathcal{D}}$ for both φ and r . For the lag time parameter in $\widetilde{\mathcal{L}}_{\mathcal{D}}$, we use $\tau = 0.1$, as the observed spectral gap for this time indicates that the time scale separation between the slow and fast processes has already manifested itself. We again solve the various integrals in (30) by Monte Carlo quadrature. One complication over the simple slow-fast system is that the transition density functions $p^{\tau}(x, \cdot)$ are not known analytically, so they have to be approximated empirically. For that, we first draw an empirical sample of $p^t(x, \cdot)$ by simulating many trajectories with starting point x , and then apply the kernel density estimation algorithm to the end points. For details on the numerical implementation see the SI.

Figure 11 (right) shows the loss function for the two RCs in dependence of σ . We see that, indeed, $\widetilde{\mathcal{L}}_{\mathcal{D}}(\varphi) < \widetilde{\mathcal{L}}_{\mathcal{D}}(r)$ for all values of σ , hence φ is the better RC. Moreover, for increasing σ , $\widetilde{\mathcal{L}}_{\mathcal{D}}(\varphi)$ continually decreases¹, whereas $\widetilde{\mathcal{L}}_{\mathcal{D}}(r)$ increases. This agrees well with our intuitive understanding of the role of σ : with increasing σ , the radial component r equilibrates more quickly, so the long-term future evolution of X_t depends more and more only on $\varphi(X_t)$ (i.e., the degree of lumpability with respect to φ increases), and less and less on $r(X_t)$ (i.e., the degree of lumpability with respect to r decreases).

¹Note that the flattening of $\widetilde{\mathcal{L}}_{\mathcal{D}}(\varphi)$ towards the right side of the plot is a numerical artefact. For high degrees of lumpability, miniscule differences between transition densities need to be quantified, which presents a challenge to our fixed-bandwidth kernel density estimator.

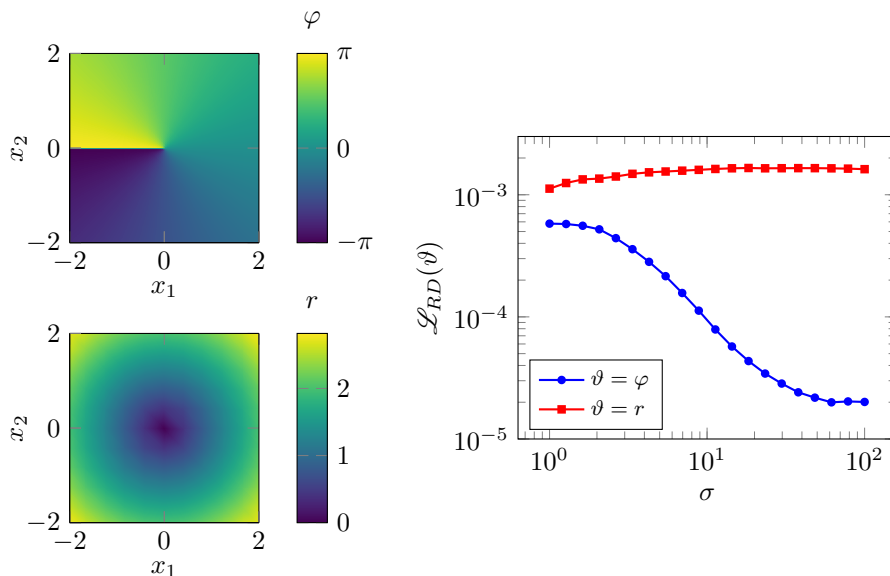


Figure 11.: Left: the two components φ and r as functions on \mathbb{X} . Right: Loss function \mathcal{L}_D for φ and r interpreted as RCs.

7. Discussion and outlook

In this paper, we formally characterized optimal reaction coordinates (RCs) in continuous Markovian systems, that is, observables that optimally describe latent long-term mechanisms of the dynamics. We have seen, by both theoretical analysis (Section 3.3) and numerical examples (Section 6), that our definition is applicable to several common types of multiscale systems, such as slow-fast systems and metastable systems. To add further interpretability to our definition, it would be desirable to draw a formal connection to the well-established transition path theory [17, 38], which characterizes RCs in terms of committor functions and minimum energy pathways on the potential energy surface. In [3], a connection between the transition manifold approach (a predecessor to the present characterization) and transition path theory has been discussed rather informally, but a rigorous investigation is still outstanding.

We then presented a variational formulation of this characterization as a computational strategy for the discovery of optimal RCs. In particular, that variational formulation provides a leverage point for modern machine learning methods, such as deep learning and stochastic gradient descent. The implementation of these methods and their demonstration of efficiency are subject of ongoing work. The reduced data requirement for the optimization problem that was demonstrated in this paper raises confidence that we will be able to apply our method to high dimensional problems such as the identification of RCs in large molecular systems.

Acknowledgements

This research has been funded by Deutsche Forschungsgemeinschaft (DFG) through grant CRC 1114 “Scaling Cascades in Complex Systems”, Project Number 235221301, Project B03 “Multilevel coarse graining of multiscale problems” and Project A01 “Coupling a multiscale stochastic precipitation model to large scale atmospheric flow dynamics”, and by Deutsche Forschungsgemeinschaft (DFG) through grant EXC 2046 “MATH+”, Project Number 390685689, Project AA1-2 “Learning Transition Manifolds and Effective Dynamics of Biomolecules”.

References

- [1] S. K. Au and J. L. Beck. Important sampling in high dimensions. *Structural Safety*, 25(2):139–163, Apr. 2003.
- [2] J. Berner, P. Grohs, G. Kutyniok, and P. Petersen. The Modern Mathematics of Deep Learning. *arXiv:2105.04026 [cs, stat]*, May 2021.
- [3] A. Bittracher, R. Banisch, and C. Schütte. Data-driven computation of molecular reaction coordinates. *The Journal of Chemical Physics*, 149(15):154103, 2018.
- [4] A. Bittracher, S. Klus, B. Hamzi, P. Koltai, and C. Schütte. Dimensionality Reduction of Complex Metastable Systems via Kernel Embeddings of Transition Manifolds. *J Nonlinear Sci*, 31(1):3, Dec. 2020.
- [5] A. Bittracher, P. Koltai, S. Klus, R. Banisch, M. Dellnitz, and C. Schütte. Transition Manifolds of Complex Metastable Systems: Theory and Data-driven Computation of Effective Dynamics. *J. Nonlinear Sci.*, 28(2):471–512, 2017.
- [6] A. Bittracher, J. Moschner, B. Koksich, R. R. Netz, and C. Schütte. Exploring the Locking Stage of NFGAILS Amyloid Fibrillation via Transition Manifold Analysis. *submitted to EPJ B*, 2021.
- [7] A. Bittracher and C. Schütte. A Weak Characterization of Slow Variables in Stochastic Dynamical Systems. In O. Junge, O. Schütze, G. Froyland, S. Ober-Blöbaum, and K. Padberg-Gehle, editors, *Advances in Dynamics, Optimization and Computation*, Studies in Systems, Decision and Control, pages 132–150, Cham, 2020. Springer International Publishing.
- [8] A. Bittracher and C. Schütte. A probabilistic algorithm for aggregating vastly undersampled large Markov chains. *Physica D: Nonlinear Phenomena*, 416:132799, Feb. 2021.
- [9] A. Bovier, V. Gayraud, and M. Klein. Metastability in reversible diffusion processes II: Precise asymptotics for small eigenvalues. 7:69-99, 2002.
- [10] G. Ciccotti, T. Lelievre, and E. Vanden-Eijnden. Projection of Diffusions on Submanifolds: Application to Mean Force Computation. *Communications on Pure and Applied Mathematics*, 61(3):371–408, 2007.
- [11] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, July 2006.
- [12] M. Dellnitz and O. Junge. On the approximation of complicated dynamical behaviour. *SIAM J. Num. Anal.*, 36(2), 1999.
- [13] J. D’Errico. randFixedLinearCombination. <https://de.mathworks.com/matlabcentral/fileexchange/49795-randfixedlinearcombination>, Retrieved June 16, 2021.
- [14] P. Deuffhard, W. Huisinga, A. Fischer, and C. Schütte. Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains. *Linear Algebra Appl.*, 315(13):39–59, 2000.
- [15] P. Deuffhard and M. Weber. Robust Perron cluster analysis in conformation dynamics. *Linear algebra and its applications*, 398:161–184, 2005.
- [16] P. Drineas and M. W. Mahoney. On the Nyström Method for Approximating a Gram Matrix for Improved Kernel-Based Learning. *J. Mach. Learn. Res.*, 6:2153–2175, Dec. 2005.
- [17] W. E and E. Vanden-Eijnden. Towards a Theory of Transition Paths. *J. Stat. Phys.*, 123(3):503–523, 2006.

- [18] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Springer New York, New York, NY, 2013.
- [19] G. Froyland, G. A. Gottwald, and A. Hammerlindl. A Computational Method to Extract Macroscopic Variables and Their Dynamics in Multiscale Systems. *SIAM Journal on Applied Dynamical Systems*, 13(4):1816–1846, Jan. 2014.
- [20] S. Garg and M. Pant. Meshfree Methods: A Comprehensive Review of Applications. *Int. J. Comput. Methods*, 15(04):1830001, June 2018.
- [21] V. Gayraud, A. Bovier, M. Eckhoff, and M. Klein. Metastability in reversible diffusion processes I: Sharp asymptotics for capacities and exit times. 2004.
- [22] M. Griebel, P. Oswald, and T. Schiekofer. Sparse grids for boundary integral equations. *Numerische Mathematik*, 83(2):279–312, 1999.
- [23] N. Halko, P. Martinsson, and J. Tropp. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Review*, 53(2):217–288, 2011.
- [24] S. R. Jammalamadaka and A. SenGupta. *Topics in Circular Statistics*, volume 5 of *Series on Multivariate Analysis*. WORLD SCIENTIFIC, Apr. 2001.
- [25] M. H. Kalos and P. A. Whitlock. *Monte Carlo Methods*. John Wiley & Sons, June 2009.
- [26] J. Kemeny and J. L. Snell. *Finite Markov Chains*. Springer-Verlag New York, 1976.
- [27] Y. Khoo, J. Lu, and L. Ying. Solving for high-dimensional committor functions using artificial neural networks. *Res Math Sci*, 6(1):1, Oct. 2018.
- [28] S. Klus, F. Nüske, P. Koltai, H. Wu, I. Kevrekidis, C. Schütte, and F. Noé. Data-Driven Model Reduction and Transfer Operator Approximation. *J Nonlinear Sci*, 28(3):985–1010, June 2018.
- [29] S. Klus, F. Nüske, S. Peitz, J.-H. Niemann, C. Clementi, and C. Schütte. Data-driven approximation of the Koopman generator: Model reduction, system identification, and control. *Physica D: Nonlinear Phenomena*, 406:132416, May 2020.
- [30] S. Klus, I. Schuster, and K. Muandet. Eigendecompositions of Transfer Operators in Reproducing Kernel Hilbert Spaces. *Journal of Nonlinear Science*, 30(1):283–315, Feb. 2020.
- [31] P. Koltai and C. Schütte. A Multiscale Perturbation Expansion Approach for Markov State Modeling of Nonstationary Molecular Dynamics. *Multiscale Model. Simul.*, 16(4):1455–1485, Jan. 2018.
- [32] U. Krengel and A. Brunel. *Ergodic Theorems*. Number 6 in De Gruyter Studies in Mathematics. W. de Gruyter, Berlin ; New York, 1985.
- [33] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. *Commun. ACM*, 60(6):84–90, May 2017.
- [34] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015.
- [35] L. Maragliano, A. Fischer, E. Vanden-Eijnden, and G. Ciccotti. String method in collective variables: Minimum free energy paths and isocommittor surfaces. *J. Chem. Phys.*, 125(2):024106, 2006.
- [36] A. Mardt, L. Pasquali, H. Wu, and F. Noé. VAMPnets for deep learning of molecular kinetics. *Nat. Commun.*, 9(1):5, 2018.

- [37] R. T. McGibbon, B. E. Husic, and V. S. Pande. Identification of simple reaction coordinates from complex dynamics. *The Journal of Chemical Physics*, 146(4):044109, 2017.
- [38] P. Metzner, C. Schütte, and E. Vanden-Eijnden. Illustration of transition path theory on a collection of simple examples. *The Journal of Chemical Physics*, 125(8):084110, Aug. 2006.
- [39] L. Molgedey and H. G. Schuster. Separation of a mixture of independent signals using time delayed correlations. *Physical Review Letters*, 72(23):3634–3637, 1994.
- [40] F. Noé, S. Olsson, J. Köhler, and H. Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, Sept. 2019.
- [41] E. Novak and H. Woźniakowski. Approximation of infinitely differentiable multivariate functions is intractable. *Journal of Complexity*, 25(4):398–404, Aug. 2009.
- [42] F. Nüske, P. Koltai, L. Boninsegna, and C. Clementi. Spectral Properties of Effective Dynamics from Conditional Expectations. *Entropy*, 23(2):134, Feb. 2021.
- [43] S. Pant, Z. Smith, Y. Wang, E. Tajkhorshid, and P. Tiwary. Confronting pitfalls of ai-augmented molecular dynamics using statistical physics. *The Journal of Chemical Physics*, 153:234118, 12 2020.
- [44] G. A. Pavliotis and A. M. Stuart. *Multiscale Methods: Averaging and Homogenization*. Springer Science & Business Media, 2008.
- [45] G. Perez-Hernandez, F. Paul, T. Giorgino, G. de Fabritiis, and F. Noé. Identification of slow molecular order parameters for Markov model construction. *The Journal of Chemical Physics*, 139(1):015102, 2013.
- [46] R. J. Rabben, S. Ray, and M. Weber. ISOKANN: Invariant subspaces of koopman operators learned by a neural network. *The Journal of Chemical Physics*, 153(11):114109, 2020.
- [47] A. Rahimi and B. Recht. Random Features for Large-Scale Kernel Machines. page 8.
- [48] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer New York, New York, NY, 1999.
- [49] G. O. Roberts and J. S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1(none):20–71, Jan. 2004.
- [50] M. Rosenblatt. Transition probability operators. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Contributions to Probability Theory, Part 2*, pages 473–483, Berkeley, Calif., 1967. University of California Press.
- [51] S. Ruder. An overview of gradient descent optimization algorithms. *arXiv:1609.04747 [cs]*, June 2017.
- [52] M. Sarich, R. Banisch, C. Hartmann, and C. Schütte. Markov State Models for Rare Events in Molecular Dynamics. *Entropy*, 16(1):258–286, Jan. 2014.
- [53] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization and Beyond*. MIT press, Cambridge, USA, 2001.
- [54] C. Schütte, A. Fischer, W. Huisinga, and P. Deuffhard. A Direct Approach to Conformational Dynamics Based on Hybrid Monte Carlo. *J. Comput. Phys.*, 151(1):146–168, 1999.
- [55] C. Schütte and M. Sarich. *Metastability and Markov State Models in Molecular Dynamics: Modeling, Analysis, Algorithmic Approaches*. Courant Lecture Notes in Mathematics. American Mathematical Society, 2014.

- [56] A. W. Senior, R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, C. Qin, A. Židek, A. W. R. Nelson, A. Bridgland, H. Penedones, S. Petersen, K. Simonyan, S. Crossan, P. Kohli, D. T. Jones, D. Silver, K. Kavukcuoglu, and D. Hassabis. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, Jan. 2020.
- [57] A. Singer, R. Erban, I. G. Kevrekidis, and R. R. Coifman. Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps. *Proceedings of the National Academy of Sciences of the United States of America*, 106(38):16090–5, 2009.
- [58] M. Udell and A. Townsend. Why Are Big Data Matrices Approximately Low Rank? *SIAM Journal on Mathematics of Data Science*, 1(1):144–160, Jan. 2019.
- [59] W. Wang and A. Roberts. Slow manifold and averaging for slow–fast stochastic differential system. *Journal of Mathematical Analysis and Applications*, 398(2):822–839, Feb. 2013.
- [60] C. Wehmeyer and F. Noé. Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics. *J. Chem. Phys.*, 148(24):241703, June 2018.
- [61] R. Zwanzig. *Nonequilibrium Statistical Mechanics*. Oxford University Press, 2001.

A. Extended multiscale expansion of slow-fast systems

In this section, we continue the multiscale analysis for the slow-fast system from Section 3.3.2. That is, we derive an evolution equation for the dominant component q_0^t of the transition density. Comparing the terms of order ε^0 in (13) yields

$$\partial_t q_0^t(x, \cdot) = \mathcal{L}_y q_1^t(x, \cdot) + \mathcal{L}_z q_0^t(x, \cdot). \quad (41)$$

The middle term is zero, which can be seen as follows: Let the averaging operator $\Pi : L_\mu^1(\mathbb{X}) \rightarrow L_{\bar{\mu}}^1(\mathbb{Z})$ be defined by

$$\Pi g(z) = \frac{\int_{\mathbb{Y}} g(y, z) \pi(y, z) dy}{\bar{\pi}(z)} \quad \text{where} \quad \bar{\pi}(z) = \int_{\mathbb{Y}} \pi(y, z) dy \quad (42)$$

and $\bar{\mu}$ is the measure induced by $\bar{\pi}$. As $q_0^t(x, \cdot)$ is independent of y in its second argument, we have

$$\partial_t q_0^t(x, \cdot) = \frac{1}{\bar{\pi}(z)} \Pi \partial_t q_0^t(x, \cdot) \quad \text{and} \quad \mathcal{L}_z q_0^t(x, \cdot) = \frac{1}{\bar{\pi}(z)} \Pi \mathcal{L}_z q_0^t(x, \cdot).$$

Furthermore, $\Pi \mathcal{L}_y = 0$. Hence, applying $\frac{1}{\bar{\pi}(z)} \Pi$ to (41) leaves the first and third term invariant, while eliminating the second term, from which follows that $\mathcal{L}_y q_1^t(x, \cdot) = 0$.

Therefore, applying Π to (41) gives

$$\partial_t \Pi q_0^t(x, \cdot) = \Pi \mathcal{L}_z q_0^t(x, \cdot).$$

As both $q_0^t(x, \cdot)$ and $\mathcal{L}_y g$ (for all g) are independent of y , we finally get for the evolution equation for q_0^t

$$\partial_t q_0^t(x, \cdot) = \mathcal{L}_y q_0^t.$$

Hence, with $\bar{\mathcal{L}} := \Pi \mathcal{L}_z \Pi$, the evolution equation of $q_0^t(x, \cdot)$ up to order ε becomes

$$\partial_t q_0^t(x, \cdot) = \bar{\mathcal{L}} q_0^t(x, \cdot) + \mathcal{O}(\varepsilon) \quad \text{for all } x \in \mathbb{X}.$$

B. Proof of Theorem 5.1

The proof consists of simple integral and norm estimations. The main argument, used multiple times in the final proof, is formulated in the following lemma:

Lemma B.1. *Assume that the system is ε -lumpable with respect to $\xi : \mathbb{X} \rightarrow \mathbb{Z}$ and the effective density $p_L : \mathbb{Z} \times \mathbb{X} \rightarrow \mathbb{R}$. Let f_L and $\text{Var } \pi$ be defined as in Theorem 5.1. Then*

$$\|f - f_L \circ \xi\|_{L^1_\mu} \leq \|\bar{\pi}\|_\infty \varepsilon. \quad (43)$$

Proof. Writing out the left hand side of (43), we get

$$\begin{aligned} \|f - f_L \circ \xi\|_{L^1_\mu} &= \int_{\mathbb{X}} \left| \left(\frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\vartheta(z)} \int_{\Sigma_\vartheta(z)} \left| \frac{p(x, y^{(1)})}{\pi(y^{(1)})} - \frac{p(x, y^{(2)})}{\pi(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz \right) \right. \\ &\quad \left. - \left(\frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\vartheta(z)} \int_{\Sigma_\vartheta(z)} \left| \frac{p_L(\xi(x), y^{(1)})}{\pi(y^{(1)})} - \frac{p_L(\xi(x), y^{(2)})}{\pi(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz \right) \right| d\mu(x). \end{aligned}$$

Applying the reverse triangle inequality, this becomes

$$\begin{aligned} \|f - f_L \circ \xi\|_{L^1_\mu} &\leq \int_{\mathbb{X}} \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\vartheta(z)} \int_{\Sigma_\vartheta(z)} \left| \frac{p(x, y^{(1)})}{\pi(y^{(1)})} - \frac{p(x, y^{(2)})}{\pi(y^{(2)})} \right. \\ &\quad \left. - \frac{p_L(\xi(x), y^{(1)})}{\pi(y^{(1)})} + \frac{p_L(\xi(x), y^{(2)})}{\pi(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) d\mu(x) dz \\ &\leq \int_{\mathbb{X}} \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\vartheta(z)} \int_{\Sigma_\vartheta(z)} \left| \frac{p(x, y^{(1)})}{\pi(y^{(1)})} - \frac{p_L(\xi(x), y^{(1)})}{\pi(y^{(1)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz d\mu(x) \\ &\quad + \int_{\mathbb{X}} \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\vartheta(z)} \int_{\Sigma_\vartheta(z)} \left| \frac{p(x, y^{(2)})}{\pi(y^{(2)})} - \frac{p_L(\xi(x), y^{(2)})}{\pi(y^{(2)})} \right| d\mu_z(y^{(1)}) d\mu_z(y^{(2)}) dz d\mu(x). \end{aligned}$$

In each of the two summands, the integrand is independent of $y^{(2)}$ and $y^{(1)}$, respectively, hence one integral over $\Sigma_\vartheta(z)$ becomes the factor $\bar{\pi}(z)$:

$$\begin{aligned} \|f - f_L \circ \xi\|_{L^1_\mu} &\leq 2 \int_{\mathbb{X}} \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\vartheta(z)} |p(x, y) - p_L(\xi(x), y)| \bar{\pi}(z) d\sigma_z(y) dz d\mu(x) \\ &\leq 2 \|\bar{\pi}\|_\infty \int_{\mathbb{X}} \frac{1}{|\mathbb{Z}|} \int_{\mathbb{Z}} \int_{\Sigma_\vartheta(z)} |p(x, y) - p_L(\xi(x), y)| d\sigma_z(y) dz d\mu(x) \end{aligned}$$

By the coarea formula, the inner two integrals simply describe the integration over \mathbb{X} with respect to the Lebesgue measure, i.e.,

$$\|f - f_L \circ \xi\|_{L^1_\mu} \leq 2 \|\bar{\pi}\|_\infty \frac{1}{|\mathbb{Z}|} \int_{\mathbb{X}} \|p(x, \cdot) - p_L(\xi(x), \cdot)\|_{L^1} d\mu(x).$$

The integral over \mathbb{X} is the L^1_μ norm, so the overall expression is the \mathbb{K} -norm (see (3)):

$$\|f - f_L \circ \xi\|_{L^1_\mu} \leq 2 \|\bar{\pi}\|_\infty \frac{1}{|\mathbb{Z}|} \|p(\cdot, \cdot) - p_L(\xi(\cdot), \cdot)\|_{\mathbb{K}}$$

which by the ε -lumpability assumption (L) is

$$\leq 2 \|\bar{\pi}\|_\infty \varepsilon.$$

□

The proof of the main result now consists only of reducing the difference of the variances to the expression $\|f - f_L \circ \xi\|_{L^1_\mu}$.

Proof of Theorem 5.1. We have

$$\begin{aligned} |\mathrm{Var}_\mu(f) - \mathrm{Var}_\mu(f_L \circ \xi)| &= |\mathbb{E}_\mu[f^2 - (f_L \circ \xi)^2] - (\mathbb{E}_\mu[f]^2 - \mathbb{E}_\mu[f_L \circ \xi]^2)| \\ &\leq \underbrace{|\mathbb{E}_\mu[f^2 - (f_L \circ \xi)^2]|}_{=:(\star)} + \underbrace{|\mathbb{E}_\mu[f]^2 - \mathbb{E}_\mu[f_L \circ \xi]^2|}_{=:(\star\star)}. \end{aligned}$$

For the first summand, we get

$$\begin{aligned} (\star) &= \|f^2 - (f_L \circ \xi)^2\|_{L_\mu^1} = \|(f - f_L \circ \xi)(f + f_L \circ \xi)\|_{L_\mu^1} \\ &\leq \|f - f_L \circ \xi\|_{L_\mu^2} \|f + f_L \circ \xi\|_{L_\mu^2}. \end{aligned}$$

As μ is a finite measure on \mathbb{X} , we have $L_\mu^2(\mathbb{X}) \subset L_\mu^1(\mathbb{X})$, hence there exists a $C > 0$ such that

$$(\star) \leq C^2 \|f - f_L \circ \xi\|_{L_\mu^1} \|f + f_L \circ \xi\|_{L_\mu^1},$$

which by Lemma B.1 can be estimated as

$$\leq C^2 \|\bar{\pi}\|_{\infty} \varepsilon \|f + f_L \circ \xi\|_{L_\mu^1}.$$

Using the inverse triangle inequality, the remaining factor $\|f + f_L \circ \xi\|_{L_\mu^1}$ can be estimated as

$$\begin{aligned} \|f + f_L \circ \xi\|_{L_\mu^1} &\leq \|f - f_L \circ \xi\|_{L_\mu^1} + 2 \|f_L \circ \xi\|_{L_\mu^1} \\ &\leq \|\bar{\pi}\|_{\infty} \varepsilon + 2 \|f_L \circ \xi\|_{L_\mu^1}. \end{aligned}$$

Overall, we get for the first summand

$$(\star) \leq 2C^2 \|f_L \circ \xi\|_{L_\mu^1} \|\bar{\pi}\|_{\infty} \varepsilon + C^2 \|\bar{\pi}\|_{\infty}^2 \varepsilon^2$$

For the second summand, we get

$$\begin{aligned} (\star\star) &= |(\mathbb{E}_\mu[f] + \mathbb{E}_\mu[f_L \circ \xi]) (\mathbb{E}_\mu[f] - \mathbb{E}_\mu[f_L \circ \xi])| \\ &\leq \|f + f_L \circ \xi\|_{L_\mu^1} \|f - f_L \circ \xi\|_{L_\mu^1}. \end{aligned}$$

By using Lemma B.1, and the same estimation of $\|f + f_L \circ \xi\|_{L_\mu^1}$ as above, this becomes

$$(\star\star) \leq 2 \|f_L \circ \xi\|_{L_\mu^1} \|\bar{\pi}\|_{\infty} \varepsilon + \|\bar{\pi}\|_{\infty}^2 \varepsilon^2$$

Overall, we receive

$$|\mathrm{Var}_\mu(f) - \mathrm{Var}_\mu(f_L \circ \xi)| \leq 2(1 + C^2) \|f_L \circ \xi\|_{L_\mu^1} \|\bar{\pi}\|_{\infty} \varepsilon + (1 + C^2) \|\bar{\pi}\|_{\infty}^2 \varepsilon^2.$$

□

Finally, we show Lemma 5.2:

Proof of Lemma 5.2. We have

$$\begin{aligned} \mathrm{Var}_\mu(h \circ \xi) &= \mathbb{E}_\mu[(h \circ \xi)^2] - (\mathbb{E}_\mu[h \circ \xi])^2 \\ &= \int_{\mathbb{X}} (h(\xi(x)))^2 \pi(x) dx - \left(\int_{\mathbb{X}} h(\xi(x)) \pi(x) dx \right)^2. \end{aligned}$$

Using the coarea formula, this becomes

$$\begin{aligned} &= \int_{\mathbb{Z}} h^2(z) \int_{\Sigma_\xi(z)} \pi(x) d\sigma_z(x) dz + \left(\int_{\mathbb{Z}} h(z) \int_{\Sigma_\xi(z)} \pi(x) d\sigma_z(x) dz \right)^2 \\ &= \int_{\mathbb{Z}} h^2(z) \bar{\pi}(z) dz + \left(\int_{\mathbb{Z}} h(z) \bar{\pi}(z) dz \right)^2 \\ &= \mathrm{Var}_{\bar{\mu}}(h). \end{aligned}$$

□

C. Lumpability of Example 6.1

For $x, y \in \mathbb{T}$, we define the shorthand notation $h^\sigma(x, y) := \left(1 + 2 \sum_{k=1}^{\infty} \rho^{k^2} \cos(k(y-x))\right)$, where $\rho = \exp(-\sigma^2/2)$. Note that since $g^\sigma(x, y) = h^\sigma(x, y)/(2\pi)$ is a density, the function h is clearly nonnegative. For all $x, y \in \mathbb{T}^n$, we have

$$\begin{aligned} |p^{\tau, \infty}(x, y) - p^{\tau, \sigma}(x, y)| &= \left| \frac{h^\tau(x^{(1)}, y^{(1)})}{(2\pi)^n} \left(1 - \prod_{i=2}^n h^{\tau\sigma}(x^{(i)}, y^{(i)})\right) \right| \\ &\leq \frac{C(\tau)}{(2\pi)^d} \max \left\{ \max_{x, y \in \mathbb{T}^n} \left| 1 - \prod_{i=2}^n h^{\tau\sigma}(x^{(i)}, y^{(i)}) \right|, \min_{x, y \in \mathbb{T}^n} \left| 1 - \prod_{i=2}^n h^{\tau\sigma}(x^{(i)}, y^{(i)}) \right| \right\} \\ &= \frac{C(\tau)}{(2\pi)^n} \max \left\{ \left| 1 - \left(1 + 2 \sum_{k=1}^{\infty} (\pm 1)^k \rho^{k^2}\right)^{n-1} \right| \right\} \end{aligned} \quad (44)$$

with the constant $C(\tau) := \max_{x^{(1)}, y^{(1)} \in \mathbb{T}} h^\tau(x^{(1)}, y^{(1)})$. Here, we use the fact that for all $\tau, \sigma > 0$, the maximum of $\prod_{i=2}^n h^{\tau\sigma}(x^{(i)}, y^{(i)})$ is attained at $x^{(i)} = y^{(i)}$ and its minimum at $|x^{(i)} - y^{(i)}| = \pi$. When we apply the binomial theorem to the right hand side of (44), we obtain

$$\begin{aligned} \left| 1 - \left(1 + 2 \sum_{k=1}^{\infty} (\pm 1)^k \rho^{k^2}\right)^{n-1} \right| &= \left| \sum_{i=1}^{n-1} \binom{n-1}{i} \left(2 \sum_{k=1}^{\infty} (\pm 1)^k \rho^{k^2}\right)^i \right| \\ &= 2(n-1) \left(\sum_{k=1}^{\infty} \rho^{k^2} \right) + \mathcal{O} \left(\left(\sum_{k=1}^{\infty} \rho^{k^2} \right)^2 \right). \end{aligned} \quad (45)$$

We now neglect the higher order terms in (45). Whenever σ is large enough such that $\rho < 1$, we have by the limit of the geometric series

$$\left(\sum_{k=1}^{\infty} \rho^{k^2} \right) \leq \left(\sum_{k=1}^{\infty} \rho^k \right) = \frac{1}{1-\rho} - 1 = \mathcal{O}(\rho) = \mathcal{O}(\exp(-(\tau\sigma)^2/2)).$$

Hence, we have shown

$$\max_{x, y \in \mathbb{T}^n} |p^{\tau, \infty}(x, y) - p^{\tau, \sigma}(x, y)| = (\exp(-(\tau\sigma)^2/2)).$$

As such, bounding both integrals in the norms in the lumpability condition (L) by the maximum above, we obtain the assertion (38).

Note that by neglecting the higher order terms in (45), we essentially ignore the impact of the prefactor depending on n in terms of the binomial coefficient, constituting the unknown dependence of the error for flexible n as mentioned in Section 6.1.2.