

A COMPUTATIONAL METHOD TO EXTRACT MACROSCOPIC VARIABLES AND THEIR DYNAMICS IN MULTISCALE SYSTEMS

GARY FROYLAND ^{*}, GEORG A. GOTTWALD [†], AND ANDY HAMMERLINDL [‡]

Abstract. This paper introduces coordinate-independent methods for analysing multiscale dynamical systems using numerical techniques based on the transfer operator and its adjoint. In particular, we present a method for testing whether an arbitrary dynamical system exhibits multiscale behaviour and for estimating the time-scale separation. For systems with such behaviour, we establish techniques for analysing the fast dynamics in isolation, extracting slow variables for the system, and accurately simulating these slow variables at a large time step. We illustrate our method with numerical examples and show how the reduced slow dynamics faithfully represents statistical features of the full dynamics which are not coordinate dependent.

Key words. multiscale systems, slow-fast systems, transfer operator, Koopman operator

AMS subject classifications. 37xx, 37Mxx, 65Pxx

1. Introduction. Devising efficient computational methods to simulate complex systems is of paramount importance to a wide range of scientific fields, including biomolecular dynamics, material science and climate science. The sheer dimensionality, often paired with the presence of a rich hierarchy of temporal scales and several metastable states, makes direct computational modelling over the whole range of temporal scales intractable.

Scientists therefore seek reduced models for some designated variables which carry the relevant information. There are two separate scenarios when such a reduction is possible: scale separation and weak coupling [35]. Here we concentrate on the large class of time-scale separated systems which include models from molecular dynamics, chemical kinetics, and climate. In scale-separated systems a reduction to slow degrees of freedom has two computational advantages. In addition to the obvious computational advantages of decreasing the dimension, model reductions offer the advantage that the time-step to be used in simulations can be orders of magnitudes larger since the reduced model only involves slow variables.

How to extract the set of relevant variables and the associated relevant dynamics from a dynamical system is one of the most challenging problems in computational modelling. In this work we address this problem on several levels: We first devise a method which tests whether the complex system under consideration is actually a scale-separated system which allows for a decomposition into slow and fast degrees of freedom. In a second step we then identify possible slow variables. We accurately describe the statistics of the fast dynamics by using transfer operator techniques and then construct the reduced dynamics for the slow variables.

For model reduction techniques to be successful, a time-scale separation must exist in the first place. Many treatments simply assume that such behaviour exists. They further assume the directions of the slow and fast dynamics are known in advance and that these directions align with the coordinates in which the system is defined.

^{*}School of Mathematics and Statistics, University of New South Wales, Australia (G.Froyland@unsw.edu.au).

[†]School of Mathematics and Statistics, University of Sydney, Australia (georg.gottwald@sydney.edu.au).

[‡]School of Mathematics and Statistics, University of Sydney, Australia and School of Mathematics and Statistics, University of New South Wales, Australia (andy@maths.usyd.edu.au).

In order to analyse general dynamical systems without making such assumptions, we develop methods based on the transfer operator and its dual, the so-called Koopman operator. These two operators represent a global description of the dynamical system’s action on ensembles and observables, respectively, and provide powerful tools for identifying global slow dynamical modes also known as strange eigenmodes, persistent patterns, metastable sets and almost-invariant sets [18, 19, 20, 51, 58, 56, 57, 26]. Transfer operator methodologies have been used to detect and approximate slowly decaying modes in a number of settings, including molecular dynamics, ocean dynamics, atmospheric dynamics, and general fluid flow [67, 30, 31].

The spectrum of the transfer operator is closely related to the ergodic properties of the system, the rate of the decay of correlations, and the overall “speed” of the dynamics [25, 5, 27]. To exploit these properties, we develop techniques for isolating the fast dynamics of the system, computing its transfer operator, and using this to estimate the scale of the temporal separation. This study of the fast dynamics in isolation is a topic which has received comparatively little attention.

Once a clear time-scale separation has been established, one needs to identify the slow variables. We use the eigenfunctions of the Koopman operator for the full system to define a projection to a lower dimensional space. This, in effect, identifies the slow variables of the system. Our approach is similar to the diffusion map approach [13, 12, 33]. Therein slow variables are defined as functions of the physical variables, and the reduction then consists of determining the temporal evolution of those functions. In contrast, we propose a method that allows for a direct representation of the slow variables, rather than describing the dynamics in function space.

The next step is to construct reduced dynamics for slow variables which reliably approximate the slow dynamics of the full system, with the computational advantages mentioned above. There exists a plethora of methods to construct reduced dynamics for slow variables [35]. Averaging of deterministic and stochastic systems over the measure induced by the fast process has been widely used [1, 66]. In the case when the averaged dynamics turns out to be trivial, diffusive effects become important and singular stochastic perturbation theory (homogenization) can be employed [43, 46, 63, 53, 40, 64] to derive reduced equations for the slow variables, even if the underlying dynamics is deterministic (but sufficiently chaotic) [55, 37]. On short time scales, deterministic and stochastic center manifold theory [9, 7, 61, 62, 4, 65] is a well-known method to describe the dynamics locally close to a fixed point, in which the fast variables are slaved to the dynamics of the slow variables. A different approach [11, 10, 69, 40] uses the Mori-Zwanzig projection formalism [60, 72]. Some of these methods have a rigorous mathematical footing with clearly stated assumptions for their validity [43, 46, 63, 9, 55, 37, 52]. However, from a practical view point these methods are limited because they require the underlying equations to be of a simple enough form to allow for the required analytical manipulations.

In more complicated situations, i.e., in realistic applications, one has to resort to numerical simulations. We mention here the equation-free projection algorithms [34, 42] and the heterogeneous multiscale method [21, 23]. These methods have been applied to a wide range of problems, including modelling of water in nanotubes, micelle formation, chemical kinetics and climate modelling and data assimilation [53, 41, 59, 36]. The underlying assumption of these methods is that the fast dynamics quickly relaxes to its equilibrium value (conditioned on the slow variables). The vector field of the reduced slow equation is then estimated by averaging over short bursts of the full dynamics. The presence of metastable states and the rare transitions between

them severely impede the relaxation towards equilibrium [71], rendering short bursts ineffective to estimate the slow dynamics. Instead we employ the transfer operator to estimate the averages. The transfer operator has been used to devise reduced slow dynamics in [15, 14] for estimating stochastic models from discrete time series. The analytical results in [14] provide support for the approach we choose here. Our reduction method differs from previous methods in that we use the invariant measure computed by the transfer operator to populate the fast fibers conditioned on the slow states. This cloud of initial conditions is then propagated for an intermediate time scale and, by using the eigenfunctions of the Koopman operator, projected back onto the slow variables. This allows us to estimate the drift and diffusion coefficients via ensemble averages. The existence of the intermediate time scale is assured by homogenisation theory and numerically estimated by requiring that the increments of the slow dynamics are nearly Gaussian.

The paper is organised as follows. In Section 2 we present an abstract framework for slow-fast systems. The basic concepts of the transfer and the Koopman operators are introduced in Section 3. Section 4 presents a novel algorithm to effectively detect the presence of a multiscale structure in a given dynamical system as well as identifying the possible slow degrees of freedom. The algorithm is illustrated with a suite of test problems. In Section 5 we present an algorithm to extract the reduced dynamics of the slow variables which were determined in Section 4. Here we use results from homogenization theory to fit parameters of a reduced stochastic differential equation (SDE) for the slow variables. Details on the computational issues involved in calculating the transfer operator and the Koopman operator are discussed in Section 6. We conclude with a discussion and outlook in Section 7.

2. An abstract framework for multiscale systems. In model reduction it is common to consider a projection $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$, where \mathcal{Z} represents the phase space of the full system, and \mathcal{X} is a lower-dimensional space where a reduced model of the dynamics captures its slowest scale behaviour; see for example [73, 35]. The lower-dimensional space \mathcal{X} does not necessarily have to be a subspace of \mathcal{Z} . In practice, however, one may want to reduce onto a slow space $\mathcal{X} \subset \mathcal{Z}$ to identify slow coordinates as a subset of the original (full) coordinate system. In this paper, we assume that we have no *a priori* knowledge about the map $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$, only that it exists. In particular, we do not assume that we know what variables are slow or fast, respectively. Therefore, none of the techniques developed in this paper explicitly use the map \mathcal{P} , but rather exploit the structure it gives to the overall system.

For this, we assume that there is a dynamical system $T : \mathcal{Z} \rightarrow \mathcal{Z}$ on the full space and a “reduced dynamics” $S : \mathcal{X} \rightarrow \mathcal{X}$ such that the following properties hold:

1. S models the long-term behaviour of T ; that is, there is a large iterate $N \gg 1$ such that $\mathcal{P} \circ T^N \approx S \circ \mathcal{P}$.
2. The short term behaviour of T can be approximated by dynamics on a fiber of \mathcal{P} . That is, if

$$z, T(z), \dots, T^n(z)$$

is an orbit of finite length $n \ll N$, then $\mathcal{P}(T^i(z))$ is nearly constant for $i = 0, \dots, n$ and there is a dynamical system $\hat{T} : F \rightarrow F$ on the fiber $F = \mathcal{P}^{-1}(\mathcal{P}(z))$ and the orbit

$$z, \hat{T}(z), \dots, \hat{T}^n(z)$$

of length n closely approximates that of T .




FIG. 2.1. An illustration of the abstract slow-fast framework. On short time scales, orbits of T stay close to fibers of \mathcal{P} . On long time scales, the orbits of T , after projecting by \mathcal{P} , can be approximated by orbits of S .

Figure 2.1 gives an illustration of this abstraction. Note that there is no unique choice for \mathcal{P} . Indeed, if $h : \mathcal{X} \rightarrow \mathcal{X}$ is any continuous invertible function, then $h \circ \mathcal{P}$ would also define a map from \mathcal{Z} to \mathcal{X} which satisfies the condition in the slow-fast framework. However, from a numerical standpoint, one projection may be better than another, and a change of projection is used for such a purpose in Section 5.

For ease of exposition we restricted the presentation here to the case of deterministic systems with discrete time. In the general case, the full dynamics and/or the reduced dynamics could be continuous-time systems, and could have stochastic components, with the obvious modifications.

To turn this abstraction into a fully rigorous definition, one could consider both the dynamical system T_ϵ and the number N_ϵ as depending on a parameter $\epsilon > 0$, such that, as ϵ tends to zero,

$$N_\epsilon \rightarrow \infty, \quad \mathcal{P} \circ T_\epsilon^{N_\epsilon} \rightarrow S \circ \mathcal{P}, \quad \text{and} \quad T_\epsilon \rightarrow \hat{T}$$

under the appropriate notions of convergence. Indeed, such an approach is used in homogenization theory [35]. In this paper, we study numerical examples where $\epsilon > 0$ is fixed, and so do not consider these convergence results in detail. However, for the techniques we develop, a clear time-scale separation is needed. The existence of such a time-scale separation is explicitly tested for by our algorithm.

3. The transfer and Koopman operator. Information of the long-term behaviour of a dynamical system can often be obtained from spectral properties of linear operators associated with the dynamics. If $T : \mathcal{Z} \rightarrow \mathcal{Z}$ represents a deterministic, discrete dynamical system, the transfer operator (or Perron-Frobenius operator) \mathcal{L} on $L^2(\mathcal{Z})$ is defined by

$$(3.1) \quad \int g \cdot \mathcal{L}f \, dm = \int (g \circ T) \cdot f \, dm$$

for all $f, g \in L^2(\mathcal{Z})$. Unless otherwise noted, we take m to be Lebesgue measure on $\mathcal{Z} \subset \mathbb{R}^d$. If μ is a measure defined by a density f (that is $d\mu = f dm$), then $\mathcal{L}f$ gives the resulting density after pushing the measure forward by T . A density f is invariant exactly when $\mathcal{L}f = f$. In other words, f is an eigenfunction with eigenvalue $\lambda = 1$. If this eigenvalue is simple, then f is the unique invariant density.

In certain settings [3, 5, 6], the transfer operator can be decomposed as $\mathcal{L} = F + V$ where F is a finite operator and V is contractive, that is, $\|V\| < 1$. The spectrum of \mathcal{L} then contains a finite number of isolated eigenvalues λ_i satisfying $|\lambda_i| > \|V\|$. In particular, there is a gap in the spectrum between $\lambda_1 = 1$ and λ_2 , the next largest eigenvalue in magnitude. The magnitude of λ_2 can be interpreted as the geometric rate at which non-equilibrium densities converge to the invariant density, often called the *rate of mixing*. In general, the leading eigenvalues λ_i with $|\lambda_i|$ close to one correspond to the slowest decaying transients of the system. Rigorous results for computing this “outer spectrum” of eigenvalues can be found in [25, 5, 27].

A similar theory holds for the adjoint of the transfer operator [58, 56, 57]. This is the Koopman operator \mathcal{K} defined simply by $\mathcal{K}g = g \circ T$.

In practice in this paper, we apply numerical schemes, which approximate \mathcal{L} and \mathcal{K} by finite dimensional operators. These approximations can be viewed as the transfer and Koopman operators $\tilde{\mathcal{L}}$ and $\tilde{\mathcal{K}}$ now defined for a non-deterministic dynamical system \tilde{T} which adds local noise to the dynamics of T . The effect of this noise is to ensure that the images of the operators $\tilde{\mathcal{L}}$ and $\tilde{\mathcal{K}}$ are finite-dimensional, and so their eigenfunctions exist and can be computed. These operators can be defined in such a way that the dynamics of \tilde{T} closely resembles that of T , and the numerical implementation of these operators is described in detail in Section 6. Because of this, for the remainder of the paper we treat \mathcal{L} and \mathcal{K} as if they are compact operators on $L^2(\mathcal{Z})$ and assume their leading eigenvalues and associated eigenfunctions can be calculated numerically.

4. Testing for multiscale behaviour and identifying fast and slow coordinates. We now consider the Koopman operator for a slow-fast system, and argue that its eigenfunctions can be used to find the projection $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$ without any *a priori* knowledge of the system.

4.1. Defining fiber dynamics. In the framework for slow-fast systems of Section 2, we can consider Koopman operators for both the full dynamics T with

$$\mathcal{K}_T g = g \circ T,$$

and the reduced dynamics S with

$$\mathcal{K}_S g = g \circ S.$$

We now compare the spectrum of \mathcal{K}_S with that of \mathcal{K}_T .

Suppose that $\psi : \mathcal{X} \rightarrow \mathbb{R}$ is an eigenfunction for the Koopman operator of the reduced dynamics \mathcal{K}_S . That is, $\psi \circ S = \hat{\lambda}\psi$ with eigenvalue $\hat{\lambda}$. Then,

$$\psi \circ \mathcal{P} \circ T^N \approx \psi \circ S \circ \mathcal{P} = \hat{\lambda}(\psi \circ \mathcal{P}),$$

so that $\psi \circ \mathcal{P} : \mathcal{Z} \rightarrow \mathbb{R}$ approximately solves the eigenfunction equation $\phi \circ T^N = \hat{\lambda}\phi$ for the Koopman operator of T^N . This suggests that \mathcal{K}_T has an eigenfunction ϕ close to $\psi \circ \mathcal{P}$ with eigenvalue $\hat{\lambda}^{\frac{1}{N}}$. As $N \gg 1$, this eigenvalue is very close to one.

On the other hand, if $\phi : \mathcal{Z} \rightarrow \mathbb{R}$ is an eigenfunction of \mathcal{K}_T with eigenvalue λ such that $0 < |\lambda| < 1$, then as N is large, either λ^N is close to zero, or λ is very

close to 1. In the latter case, for a finite orbit $z, T(z), \dots, T^n(z)$ with $n \ll N$ the values $\phi \circ T^j(z) = \lambda^j \phi(z)$ are nearly constant. Since this orbit stays close to the fiber $\mathcal{P}^{-1}(\{v\})$ and N is large, it suggests that ϕ will be nearly constant along the fiber. In certain settings, the above arguments can be made rigorous with convergence results as the time-scale separation tends to infinity [14].

Suppose the Koopman operator associated with the full dynamics K_T has a number of associated eigenfunctions $\phi_i : \mathcal{Z} \rightarrow \mathbb{R}$ ($i = 1, \dots, m$), all with eigenvalues close to one. Then these have associated approximations $\psi_i \circ \mathcal{P} : \mathcal{X} \rightarrow \mathbb{R}$. If, further, \mathcal{X} can be identified with its image under the product map

$$\psi_1 \times \psi_2 \times \dots \times \psi_m : \mathcal{X} \rightarrow \mathbb{R}^m$$

then, the product

$$\phi_1 \times \phi_2 \times \dots \times \phi_m : \mathcal{Z} \rightarrow \mathbb{R}^m$$

gives an approximation of the mapping $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$.

For simplicity, we restrict our study here to examples where the reduced space \mathcal{X} is one-dimensional. In this case the mapping $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$ can be well approximated by a single eigenfunction $\phi : \mathcal{Z} \rightarrow \mathbb{R}$. Possible extensions to higher-dimensional slow subspaces are discussed in Section 7.

To see if an eigenfunction ϕ is actually approximating the (non-linear) projection \mathcal{P} of a multiscale system, we develop a test for multiscale behaviour. To do this, we first approximate the fast dynamics $\hat{T} : F \rightarrow F$ defined on a fiber $F := \mathcal{P}^{-1}(\{v\}) = \phi^{-1}(\{v\})$. Assuming ϕ is differentiable, its gradient $\nabla\phi$ is a vector field on \mathcal{Z} . If this vector field is non-zero on a neighbourhood U of the fast fiber F , it defines a map π from U to F simply by flowing either along $\nabla\phi$ or $-\nabla\phi$ depending on whether the starting point $z \in U$ satisfies $\phi(z) < v$ or $\phi(z) > v$. Moreover, one can show by basic calculus that

$$\|\pi(z) - z\| \leq \frac{1}{C} \|\phi(\pi(z)) - \phi(z)\| = \frac{1}{C} \|v - \phi(z)\|$$

where C is the minimum value of $\|\nabla\phi(z)\|$ on U . For a deterministic T , define $\hat{T} : F \rightarrow F$ by $\hat{T} = \pi \circ T$. Then,

$$(4.1) \quad \|\hat{T}(z) - T(z)\| < \frac{1}{C} |v - \phi(T(z))| = \frac{1}{C} |v - \lambda\phi(z)| = \frac{1}{C} |1 - \lambda| |v|.$$

For λ close to one, the term $|1 - \lambda|$ above will be close to zero, and \hat{T} will well approximate T and can be used to analyze the fast dynamics of the system. The closer to 1, the better the approximation for the dynamics \hat{T} on the fast fibers; therefore, in practice, we use the eigenfunction ϕ which is associated with the second leading eigenvalue λ_2 to define the fast fibers and the projection \mathcal{P} .

4.2. Algorithm to test for multiscale behaviour and identify slow variables. To test for multiscale behaviour, we look at the Koopman operator for the fiber dynamics \hat{T} . The spectrum of the operator is connected to the rate of decay of correlations for the system [16]. Therefore, we expect for the fast dynamics that the spectrum does not have values as close to one, as for the spectrum of the full system.

In fact, we propose the following algorithm to test for multiscale behaviour:

Algorithm 1

1. Compute a numerical approximation \mathcal{K} of the Koopman operator of the full dynamical system T .
2. Determine its leading eigenvalues

$$\lambda_1, \lambda_2, \lambda_3, \dots$$

such that $1 = |\lambda_1| \geq |\lambda_2| \geq |\lambda_3| \geq \dots$.

3. Using the eigenfunction ϕ associated to λ_2 , take a fiber $F = \phi^{-1}(\{v\})$ on which $\|\nabla\phi\|$ is bounded away from zero.
4. Define the “fiber dynamics” $\hat{T} : F \rightarrow F$ by $\hat{T} = \pi \circ T$ where π is determined by the gradient flow associated with ϕ .
5. Compute a numerical approximation $\hat{\mathcal{K}}$ of the Koopman operator of the fiber dynamics \hat{T} .
6. Determine the leading eigenvalues of $\hat{\mathcal{K}}$

$$\hat{\lambda}_1, \hat{\lambda}_2, \hat{\lambda}_3, \dots$$

7. Compare $\hat{\lambda}_i$ to λ_i . A large ratio implies multiscale behaviour.

In a multiscale system, we expect the ratios to be large, roughly on the same order as the time-scale separation between the slow and fast dynamics, and in this paper we give examples where this is the case.

If the same techniques are applied to a system without multiscale behaviour, then the simulated dynamics on a fiber will evolve at more-or-less the same rate as the full system. The resulting eigenvalues for the fiber dynamics may be smaller than for the full system, due to the fiber dynamics acting on a domain of smaller dimension, but the difference will be modest in nature. Specific examples of this are given in 4.3.1 and 4.3.3.

Once it is determined that a given dynamical system has multiscale behaviour and that $\phi : \mathcal{Z} \rightarrow \mathbb{R}$ reasonably approximates the map $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$, we know that \hat{T} defined on a fiber reasonably approximates the fast dynamics. The map ϕ in effect defines the slow variable. It remains to give an approximation for the reduced slow dynamics $S : \mathcal{X} \rightarrow \mathcal{X}$. We describe such a construction in Section 5 and how it can be used to simulate the slow dynamics of the system for a time step much larger than is possible for the full dynamics.

In Algorithm 1, one must compute numerical approximations of Koopman operator and its eigenfunctions. In the examples in this paper, we use Ulam’s method [70, 49] which approximates the transfer and Koopman operators by dividing the space into a finite number of boxes. Care must be taken when applying Ulam’s method to multiscale systems, and specific details about our implementation of the method are given in Section 6. Also, most of our example systems are defined by stochastic differential equations, and Section 6 also shows how the operators and their numerical approximations can be defined in this case.

One could alternatively consider the eigenfunctions of the transfer operator defined with respect to the invariant measure. However, as discussed in Section 6, there are significant numerical problems in computing them to sufficient accuracy. Therefore, in this paper, we only consider eigenfunctions of the Koopman operator when computing level sets.

4.3. Examples.

4.3.1. An analytic example. Before delving into numerics, we first analyze an overly simple example which has known analytic formulas for the eigendecomposition. Consider two independent Ornstein-Uhlenbeck processes given by the SDE

$$(4.2) \quad dx_t = -x dt + dV_t$$

$$(4.3) \quad dy_t = -\frac{1}{\epsilon^2}y dt + \frac{1}{\epsilon}dW_t$$

for some fixed $\epsilon > 0$ and where V_t and W_t are independent Wiener processes. If $0 < \epsilon \ll 1$, the y -variable evolves much faster than the x -variable and the SDE can be thought of as an example of a multiscale dynamical system where, in the abstract slow-fast framework of Section 2, $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$ is given by $\mathcal{P}(x, y) = x$.

For now, consider the x variable on its own. This is a classical one-dimensional Ornstein-Uhlenbeck system, and its spectral properties are known [68]. For each choice of flow time τ there is an associated Koopman operator \mathcal{K}_τ for the one-dimensional system given by (4.2). These form a semi-group of operators $\mathcal{K}_\tau = e^{-\tau\mathcal{A}}$, generated by an operator \mathcal{A} . Note the convention here that we include a minus sign when defining the generator, implying that eigenvalues of \mathcal{A} are non-negative. In the appropriate Hilbert space, the eigenfunctions of \mathcal{A} (and therefore of \mathcal{K}_τ for each τ) are the Hermite polynomials H_n . The eigenvalues of the infinitesimal generator \mathcal{A} are exactly the numbers $\{0, 1, 2, \dots\}$ and $\mathcal{A}H_n = nH_n$ for each n . To keep the analysis independent of the choice of flow-time τ , we express all eigenvalues in terms of the generator \mathcal{A} .

The system in (4.3) is the same as that of (4.2), but with a rescaling of time by a factor of ϵ^{-2} . Therefore, its generator has eigenvalues $\{0, \epsilon^{-2}, 2\epsilon^{-2}, \dots\}$.

For the two-dimensional system, as the two processes are independent, products of the form $H_m \times H_n : \mathbb{R}^2 \rightarrow \mathbb{R}$, $(x, y) \mapsto H_m(x) \cdot H_n(y)$ are eigenfunctions, and the eigenvalues for the full system are then of the form $m + \epsilon^{-2}n$ for $m, n \geq 0$. Assuming ϵ is sufficiently small, the leading eigenvalues, i.e., those with real part closest to zero, will be those with m small and n equal to zero. The function H_0 is identically equal to one, and therefore $(H_m \times H_0)(x, y) = H_m(x)$.

The level sets of such a function are lines where the x value is constant, and the gradient $\nabla(H_m \times H_0)$ points purely in the x direction. Therefore, one can verify that the fiber dynamics $\hat{T} = \pi \circ T$ described in the last section corresponds exactly to leaving x constant and evolving y according to (4.3). That is, in this simple case, the fiber dynamics exactly captures the fast dynamics.

The leading eigenvalues associated to the fiber dynamics are (in terms of the generator) $\{0, \epsilon^{-2}, 2\epsilon^{-2}, \dots\}$ which is exactly ϵ^{-2} times the leading eigenvalues of the full system. Recall that ϵ^{-2} is the difference in speed of the slow and fast systems.

In contrast, consider the same system, but with ϵ equal to one. Then, x and y evolve under identical independent processes and one can check that the eigenspace associated to an eigenvalue $\lambda = n$ is of dimension $n + 1$. In particular, for the second eigenvalue $\lambda = 1$, every eigenfunction $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ is in the span of $H_1 \times H_0$ and $H_0 \times H_1$ where $H_1(x) = x$ and therefore ϕ is of the form $\phi(x, y) \mapsto (ax + by)$. For $\epsilon = 1$, the system (4.2)–(4.3) is invariant under rotation about the origin, and so we can show that regardless of the exact values of a and b , the fiber dynamics are given by an Ornstein-Uhlenbeck process exactly as in (4.2). Hence, the eigenvalues associated to the fiber dynamics are $\{0, 1, 2, \dots\}$. The leading ten eigenvalues of the full system, are (when counted with multiplicity as they would be if calculated numerically)

$$0, 1, 1, 2, 2, 2, 3, 3, 3, 3$$

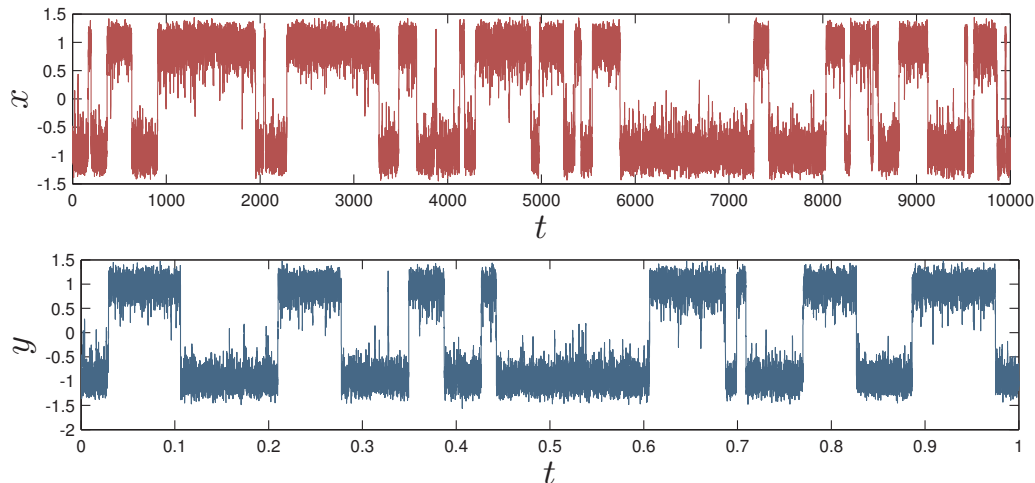


FIG. 4.1. The slow variable x and fast variable y computed for the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$. The x variable is plotted over an interval of 10^4 units of time, while the y variable is plotted over an interval 1 unit of time.

and the ratios are

$$\frac{0}{0}, \frac{1}{1}, \frac{2}{1}, \frac{3}{2}, \frac{4}{2}, \frac{5}{2}, \frac{6}{3}, \frac{7}{3}, \frac{8}{3}, \frac{9}{3}.$$

These ratios are growing linearly, but are not large as in the slow-fast case. Ratios computed for a system with $\epsilon \approx 1$ will be similar in magnitude.

4.3.2. A skew-product example. We now give an example system which exhibits metastable behaviour in both slow and fast variables, defined by the following SDE on \mathbb{R}^2

$$(4.4) \quad dx = (x - x^3 + \frac{a}{\epsilon}y) dt$$

$$(4.5) \quad dy = \frac{1}{\epsilon^2}(y - y^3) dt + \frac{\sigma}{\epsilon} dW_t.$$

Here, W_t represents the standard Wiener process. Homogenization results (see [35, 64]) show that as $\epsilon \rightarrow 0$, the evolution of the x variable approaches that of a drift-diffusion process with the contribution of $\frac{a}{\epsilon}y$ replaced by Gaussian noise. For numerical analysis, however, we only consider the system with the constants fixed at $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$. For these parameters the fast y variable switches between states near $y = -1$ and $y = +1$ on the order of about 0.1 units of time. The slow x variable switches between states near $x = -1$ and $x = +1$ on the order of hundreds of units of time, as illustrated in Figure 4.1. Because of the long time step needed to switch between them, we refer to these two states near $x = \pm 1$ as the *slow metastable states* of the system.

For the algorithms developed here, the dynamical system is treated as a black box. It is represented by a computer routine which, given initial data $z = (x, y) \in \mathbb{R}^2$, returns a point $T(z) \in \mathbb{R}^2$ given by the Euler-Maruyama integration method [44] after a single time step of $\Delta t = 2 \times 10^{-6}$. As the system is defined by a stochastic




FIG. 4.2. The invariant density of the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$, computed using Ulam’s method with a flow time of $\tau = 0.04$. (See Section 6 for numerical details.)

differential equation, calling the routine multiple times with the same input will yield different outputs.

The value Δt was chosen to be large enough to allow efficient computation of the system while being small enough to still simulate the fast dynamics of the system to an acceptable degree of accuracy.

To compute approximations of the eigenfunctions using Ulam’s method, we apply 2×10^4 iterates of this time step $\Delta t = 2 \times 10^{-6}$ for a total time of $\tau = 0.04$. Note that while this time interval corresponds to a large number of iterates, it is still much smaller than the time necessary for x to switch between its two metastable states. However, Ulam’s method computes accurate approximations of the eigenfunctions as shown in Figures 4.2 and 4.3.

We now choose a fiber and simulate dynamics restricted to that fiber as described in Algorithm 1. As will be explained in Section 6, a time interval of $\hat{\tau} = 4 \times 10^{-5}$ is sufficient to accurately compute eigenvalues of the fiber dynamics. Note that this $\hat{\tau}$ is much smaller than the time of $\tau = 0.04$ used for the full dynamics.

The leading ten eigenvalues of both the full dynamics and the fiber dynamics are given in Table 4.1. For this example, a range of several values of v was used to compute dynamics on the fibers $\phi^{-1}(\{v\})$ and gave similar eigenvalues. As these were computed using the different time steps τ and $\hat{\tau}$, we also give the eigenvalues χ_i for the infinitesimal generator, computed by taking $\chi_i = -\frac{1}{\tau} \log(\lambda_i)$. The final column gives the ratio $\hat{\chi}_i/\chi_i$ of the eigenvalues of the two generators. These ratios are large (over 10^3), and all within the same order of magnitude, which gives strong evidence the system has slow-fast behaviour with a time-scale separation of between 10^3 and 10^4 . Compare this to the scaling term $\epsilon^{-2} = 10^4$ in equation (4.5) and with the observed switching times for x and y in Figure 4.1.

As a further test that the fiber dynamics is a valid approximation of the fast




FIG. 4.3. The second eigenfunction of the Koopman operator for the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$, computed using Ulam’s method with a flow time of $\tau = 0.04$. (See Section 6 for numerical details.)

TABLE 4.1

Leading eigenvalues λ_i of the Koopman operator for the full two-dimension dynamics of the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$ with flow time $\tau = 0.04$. Also computed are the flow-independent values $\chi_i = -\frac{1}{\tau} \log |\lambda_i|$ corresponding to the infinitesimal generator of the Koopman operator. Next, are the leading eigenvalues of the Koopman operator of the fiber dynamics computed with flow time $\hat{\tau} = 4 \times 10^{-5}$ and the corresponding $\hat{\chi}_i = -\frac{1}{\hat{\tau}} \log |\hat{\lambda}_i|$. Here, the fiber $\phi^{-1}(\{v\})$ is for $v = 0.8$ where ϕ is the eigenfunction plotted in Figure 4.3. The final column shows the ratio $\hat{\chi}_i/\chi_i$.

Full dynamics		Fiber dynamics		Ratio
λ_i	χ_i	$\hat{\lambda}_i$	$\hat{\chi}_i$	$\hat{\chi}_i/\chi_i$
1.0000	0	1.0000	0	—
0.9995	1.25×10^{-2}	0.9979	5.26×10^1	4.20×10^3
0.9665	8.52×10^{-1}	0.7212	8.17×10^3	9.59×10^3
0.9429	1.47	0.5838	1.35×10^4	9.15×10^3
0.9153	2.21	0.4694	1.89×10^4	8.55×10^3
0.8801	3.19	0.3463	2.65×10^4	8.30×10^3
0.8397	4.37	0.2424	3.54×10^4	8.11×10^3
0.7954	5.72	0.1619	4.55×10^4	7.95×10^3
0.7477	7.27	0.1055	5.62×10^4	7.74×10^3
0.6969	9.02	0.0646	6.85×10^4	7.59×10^3

dynamics, we populated the fiber with 10^4 points $\{z_i\}$ distributed in proportion to the computed invariant density of the system, and then computed $\|T(z_i) - z_i\|$ and $\|\hat{T}(z_i) - T(z_i)\|$ for each point. For this example, the average value of $\|T(z_i) - z_i\|$ was 3.69×10^{-2} and the average value of $\|\hat{T}(z_i) - T(z_i)\|$ was 4.69×10^{-4} . Further, the maximum value of $\|\hat{T}(z_i) - T(z_i)\|$ over all of the points z_i was 1.26×10^{-2} . This shows that the fiber dynamics \hat{T} closely approximates T .

4.3.3. A “slow-slow” system. We now show that the same procedure when applied to a system without multiscale behaviour gives noticeably different results. As with the last example, consider the system defined by (4.4)–(4.5), but now set $\epsilon = 1$ and $a = 5$. The value $\sigma^2 = 0.113$ is as before. Under these parameters, y stays near $+1$ or -1 for long periods of time and, because of the term $\frac{a}{\epsilon}y$ in (4.4), the x variable is now highly correlated with y . This system does not exhibit multiscale behaviour.

As with the last example, we computed eigenfunctions and then populated a fiber $\phi^{-1}(\{v\})$ with 10^4 points $\{z_i\}$ distributed in proportion to the computed invariant density of the system, and then computed $\|T(z_i) - z_i\|$ and $\|\hat{T}(z_i) - T(z_i)\|$ for each point. In this case, the average value of $\|T(z_i) - z_i\|$ was 5.36×10^{-3} and the average value of $\|\hat{T}(z_i) - T(z_i)\|$ was 5.34×10^{-3} . These values are nearly equal, showing that the full dynamics does not stay close to the fiber on short time scales, and that $\hat{T} = \pi \circ T$ is not a valid approximation of the dynamics.

Applying Algorithm 1 produces eigenvalues as shown in Table 4.2. One should not read too much into the table, as the fiber dynamics do not meaningfully correspond to any form of “fast” dynamics, but notice that the ratios $\hat{\chi}_i/\chi_i$ are not large in general and that the first ratio $\hat{\chi}_2/\chi_2 \approx 197.37$ is of significantly different order than the others.

TABLE 4.2

Leading eigenvalues λ_i of the Koopman operator for the full two-dimensional dynamics of the “slow-slow” system (4.4)–(4.5) with $\epsilon = 1$, $a = 5$, and $\sigma^2 = 0.113$ with flow time $\tau = 0.04$. Also computed are the flow-independent values $\chi_i = -\frac{1}{\tau} \log |\lambda_i|$ corresponding to the infinitesimal generator of the Koopman operator. Next, are the leading eigenvalues of the Koopman operator of the fiber dynamics also computed with flow time $\hat{\tau} = 0.04$ and the corresponding $\hat{\chi}_i = -\frac{1}{\hat{\tau}} \log |\hat{\lambda}_i|$. The final column shows the ratio $\hat{\chi}_i/\chi_i$.

Full dynamics		Fiber dynamics		Ratio
λ_i	χ_i	$\hat{\lambda}_i$	$\hat{\chi}_i$	$\hat{\chi}_i/\chi_i$
1.0000	0	1.0000	0	—
0.9997	7.50×10^{-3}	0.9425	1.48	197.37
0.9674	8.29×10^{-1}	0.8105	5.25	6.34
0.9479	1.34	0.7019	8.85	6.62
0.9470	1.36	0.5365	15.6	11.43
0.9254	1.94	$0.5090 + 0.0237i$	16.9	8.70
$0.8983 + 0.0411i$	2.66	$0.5090 - 0.0237i$	16.9	6.39
$0.8983 - 0.0411i$	2.66	$0.4903 + 0.0814i$	17.5	6.58
0.8961	2.74	$0.4903 - 0.0814i$	17.5	6.37
0.8647	3.63	$0.4676 + 0.1403i$	17.9	4.93

4.3.4. A distorted example. A major advantage of our approach to identifying and analyzing slow-fast systems is that the techniques are independent of the choice of coordinates and therefore the slow and fast directions do not need to align with coordinate axes in order for the analysis to work. To demonstrate this, we apply a transformation to system (4.4)–(4.5). Recall that in Section 4.3.2, we treated the SDE (4.4)–(4.5) as a black box. Given a point $z \in \mathbb{R}^2$, we have an opaque method for computing the point $T(z)$ which is the evolution after a time-step $\Delta t = 2 \times 10^{-6}$.

In place of $z \mapsto T(z)$ consider instead, $z \mapsto g \circ T \circ g^{-1}(z)$ where g is a diffeomorphism of the domain \mathbb{R}^2 . This changes the coordinates without actually changing the dynamics. We take g to be the composition $h \circ R$, where $h(x, y) = (x + 0.3y^2, y)$




FIG. 4.4. Two finite orbits of the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$ before and after applying the distortion g described in Section 4.3.4. Each orbit represents flowing for one unit of time, and the two orbits were chosen to demonstrate the behaviour of the system near the two slow metastable states.




FIG. 4.5. The invariant density of the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$ after the distortion g described in Section 4.3.4. The density is computed by Ulam’s method applied to the distorted system using a flow time of $\tau = 0.04$.

and R is an irrational rotation of the plane (specifically, a clockwise rotation by one radian). Figure 4.4 shows this distortion applied to two orbits of the system.

It is not immediately clear in the distorted picture where exactly the slow and fast directions lie. Still, we can apply Ulam’s method to calculate the eigenfunctions of the transfer and Koopman operators, as illustrated in Figures 4.5 and 4.6. The leading eigenvalues should, of course, be the same as in the original system, and computation confirms that, as shown in Table 4.3.




FIG. 4.6. The second eigenfunction of the Koopman operator for the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$ after the distortion g described in Section 4.3.4. The eigenfunction is computed by Ulam’s method applied to the distorted system using a flow time of $\tau = 0.04$.

TABLE 4.3

Leading eigenvalues for the Koopman operator of the full system and fiber dynamics as described in Algorithm 1. These eigenvalues are computed using Ulam’s method with axis-aligned boxes for the original system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$ and for the same system conjugated by the diffeomorphism $g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined in Section 4.3.4.

Full Dynamics ($\tau = 0.04$)		Fiber Dynamics ($\hat{\tau} = 4 \times 10^{-5}$)	
original	distorted	original	distorted
1.0000	1.0000	1.0000	1.0000
0.9995	0.9994	0.9979	0.9980
0.9665	0.9663	0.7212	0.7230
0.9429	0.9423	0.5838	0.5838
0.9153	0.9148	0.4694	0.4675
0.8801	0.8787	0.3463	0.3461
0.8397	0.8385	0.2424	0.2423
0.7954	0.7939	0.1619	0.1615
0.7477	0.7454	0.1055	0.1027
0.6969	0.6944	0.0646	0.0634

Using the second eigenfunction of the Koopman operator, we compute fiber dynamics as before. Eigenvalues associated to the fiber dynamics are computed and are nearly identical to the eigenvalues computed for the undistorted system. This shows that this technique of comparing eigenvalues to detect multiscale systems is independent of the choice of coordinates.

As before, we tested the validity of the fiber dynamics by populating a fiber with 10^4 points $\{z_i\}$ distributed in proportion to the computed invariant density of the

system, and then computed $\|T(z_i) - z_i\|$ and $\|\hat{T}(z_i) - T(z_i)\|$ for each point. For the distorted system, the average value of $\|T(z_i) - z_i\|$ was 4.37×10^{-2} and the average value of $\|\hat{T}(z_i) - T(z_i)\|$ was 2.99×10^{-4} . Further, the maximum value of $\|\hat{T}(z_i) - T(z_i)\|$ over all of the points z_i was 1.26×10^{-2} . This shows that, as in the undistorted case, the fiber dynamics \hat{T} closely approximates T .

5. Reduced dynamics. The fiber dynamics approach studied in the previous sections gives a good way to isolate the fast dynamics of the system. We also want to efficiently emulate the slow dynamics of the system using a large time step. To do this, we assume the slow dynamics of the system is well approximated by a drift-diffusion process. The details and justification of this approach will be given later in the section, but we first list the proposed algorithm here:

Algorithm 2

1. Numerically compute the invariant density ρ of the dynamical system T .
2. Numerically compute an approximate eigenfunction ϕ of the Koopman operator associated to an eigenvalue close to one.
3. Reparameterize ϕ to yield a function θ which varies evenly throughout the domain.
4. Compute fibers for θ .
5. For each computed fiber $\theta^{-1}(\{v\})$:
 - (a) Populate $\theta^{-1}(\{v\})$ by an ensemble of points $\{z_i\}_{i=1}^q$ sampled in proportion to the invariant density ρ^1 .
 - (b) Iterate each point z_i forward by T^k and project by θ .
 - (c) Compute $\alpha_k(v)$ and $\beta_k(v)$ as the mean and variance of the ensemble $\{\vartheta_i\}_{i=1}^q = \{\theta(T^k(z_i)) - v\}_{i=1}^q$. That is,

$$\alpha_k(v) = \frac{1}{q} \sum_i \vartheta_i \quad \text{and} \quad \beta_k(v) = \frac{1}{q} \sum_i (\vartheta_i - \alpha_k(v))^2.$$

6. Extend α_k and β_k by linear interpolation to functions on \mathbb{R} .
7. Use $\alpha = \alpha_k/k\Delta t$ and $\beta = \beta_k/k\Delta t$ to define a drift-diffusion process

$$dX_t = \alpha(X_t)dt + \sqrt{\beta(X_t)}dW_t$$

representing the reduced slow dynamics.

A visual depiction of this algorithm is given in Figure 5.1.

Let us first comment on Step 3 of the algorithm and how we define the projection θ . Since we are assuming fast dynamics occurs along the level sets of an eigenfunction ϕ of the Koopman operator, this function ϕ presents a natural choice for projection to a one-dimensional space for the slow dynamics, as discussed in Section 4.1. However, this is not an ideal candidate for numerical calculation in systems with meta-stable slow states: Suppose the dynamical system has two metastable states in the slow direction, as in the equations (4.4)–(4.5). The eigenfunction ϕ , after scaling, will be close to +1 near one of the metastable states, close to -1 near the other, and change sharply in the area where it transitions from one state to the other. Such a sharp change is visible in Figure 4.3. In fact, such regions where an eigenfunction is nearly

¹Alternatively one could use the fast fiber dynamics defined in Section 4.1 to compute the invariant density on each fast fiber F . However, this is computationally much more involved and we therefore use the invariant density of the full dynamics.




FIG. 5.1. An illustration depicting Algorithm 2.

constant can be used to detect the existence of such metastable states and of “almost invariant sets” [17, 29]. For our purpose, however, those regions make it hard numerically to use ϕ as a projection.

To overcome this problem, we replace ϕ by a function which has similar level sets, but which varies evenly throughout the phase space.

Suppose ϕ is represented numerically by its values on a finite grid of points $\{z_i\}_{i=1}^M$ in the phase space \mathcal{Z} . Then, we define the function $\theta : \{z_i\}_{i=1}^M \rightarrow \{1, \dots, M\}$ as a bijection of finite sets and such that $\phi(z_i) < \phi(z_j)$ implies $\theta(z_i) < \theta(z_j)$. That is, from ϕ we have extracted only its ordering of points on the grid. To get a suitable continuous function, extend θ to the rest of the phase space by linear interpolation in each coordinate. Let $B \subset \mathbb{R}^d$ be the box on which the grid of points lies. Then, one way of thinking about this step, is that we are replacing ϕ by some reparameterization $\theta = h \circ \phi$ with $h : \mathbb{R} \rightarrow \mathbb{R}$ a homeomorphism, so that for an interval $[a, b] \subset \mathbb{R}$ the volume of $\theta^{-1}([a, b]) \cap B$ is roughly proportional to the length of $[a, b]$.

For example, the eigenfunction plotted in Figure 4.3 was computed using a 200 by 200 grid G of points $z_i = (x_i, y_i) \in \mathbb{R}^2$. The relative ordering of the values $\phi(x_i, y_i)$ on this grid G defines a bijection $\theta : G \rightarrow \{1, \dots, 40\,000\}$ which extends to the convex hull of G by bilinear interpolation. Such a construction would yield a θ ranging between 1 and 40 000. As a further step, we rescale θ so that it ranges between -4 and $+4$. This makes it easier to allow for an easier comparison between $\theta(T^k(z_i))$ and the original slow x -variable in (4.4)–(4.5).

Algorithm 2 relies on the reduced system being a stochastic differential equation. Diffusion limits of stochastic multiscale systems such as (4.4)–(4.5) have been rigorously established in the framework of homogenization (see, for example, [35, 64])

in the limit $\epsilon \rightarrow 0$. Diffusion limits for deterministic dynamical systems have also been established [55, 37] and have been extended to the discrete time case as well [37]. Assuming, then, that the slow dynamics is well modelled by a one-dimensional drift-diffusion system, our algorithm describes how to estimate the drift term $\alpha(X)$ and the diffusion term $\beta(X)$ in

$$(5.1) \quad dX_t = \alpha(X_t)dt + \sqrt{\beta(X_t)}dW_t$$

so that an orbit of this process closely resembles the projection under θ of an orbit of the full system.

We estimate α and β at $v \in \mathbb{R}$ by averages over an ensemble

$$\{X_{t_k, i}\}_{i=1}^q$$

of size q . The ensemble is generated as follows (cf. Figure 5.1): Populate the fiber F by a number of points z_i , evolve each point forward by k time-steps of the computation to a point to $T^k(z_i)$ at time t_k , and then project onto the slow subspace using the projection θ to yield $\theta(T^k(z_i))$. The initial population of points $\{z_i\}_{i=1}^q$ on F should be chosen in proportion to the invariant density function computed for the full dynamics (or the dynamics on the fast fibers). We remark that this procedure bears resemblance with the equation-free approach proposed in [34, 42, 41, 13, 12].

The drift and diffusion coefficients $\alpha(v)$ and $\beta(v)$ are then approximated for small times t_k by $\alpha_k(v)/t_k$ and $\beta_k(v)/t_k$ where $\alpha_k(t_k)$ and $\beta_k(t_k)$ are the mean and variance of $\{X_{t_k, i} - v\}_{i=1}^q$ [32].

Note that in general the variance of $\{X_{t_k, i} - v\}_{i=1}^q$ would depend on the drift term as well as the diffusion term in (5.1); however for small times t_k the contribution of the drift term is order t_k and can be neglected with respect to the diffusion term.

Let us now comment on the choice of the flow time t_k (or the number of iterations, k). The definition of the drift and diffusion coefficients requires $t_k \rightarrow 0$ [32]. However, to capture time scales which go beyond the fast time scale, i.e., times which are sufficiently large to allow for equilibration on each fast fiber, t_k needs to be sufficiently large. In the context of the particular scaling of the multiscale system (4.4)–(4.5) we require $t_k \sim \mathcal{O}(1/\epsilon)$, which is long enough that the fast variables behave according to the invariant density on each fiber and small enough to allow for (nearly) constant slow variables. For such times the diffusion term in the reduced dynamics is expected to be dominant; hence we chose t_k such that the set of values $\theta(T^k(z_i))$ has a distribution close to normal. To test for normality we use the Lilliefors test for normality [50]. In many cases, the Lilliefors test and the Kolmogorov-Smirnov test on which it is based are used to test the hypothesis that a given finite set of points was sampled from a normal distribution. In our situation, the projection of the dynamics only approximates a normal distribution and it is not helpful to calculate the probability that such a population came from a true normal distribution. Instead, we consider the test statistic used in the Lilliefors test, which is a positive number computed from the sample, and use that this test statistic is closer to zero when the sample is closer to normal. In Figure 5.2 we show how t_k can be chosen as the time where the Lilliefors test statistics has a minimum.

In the framework for slow-fast systems introduced in Section 2, the drift-diffusion process (5.1) now plays the role of the reduced dynamical system S on the reduced space $\mathcal{X} = \mathbb{R}$ and the map $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$ is given by θ . By the assumption $\mathcal{P} \circ T^N \approx S \circ \mathcal{P}$ for large N , if k is sufficiently large, then for a point $z \in \theta^{-1}(\{v\})$, the value $\theta(T^k(z))$




FIG. 5.2. The Lilliefors test statistic for the ensembles $\{\theta(T^k(z_i))\}$ described in Algorithm 2. A lower test statistic indicates the distribution is closer to normal. One iterate corresponds to a time step $\Delta t = 2 \times 10^{-6}$ of the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$.

gives an approximation of the result of evolving (5.1) forward by a time t proportional to k from the starting value $X_0 = v$. Thus, the computation $v \rightsquigarrow z_i \rightsquigarrow \theta(T^k(z_i))$ in Step 5 of Algorithm 2 provides the ensemble $\{X_{t_k,i}\}$ described above.

5.1. Examples.

5.1.1. The skew-product example. We now apply Algorithm 2 to the example system (4.4)–(4.5) from Section 4.3.2. For this system, Figure 5.2 plots how the test statistic varies in the number of iterates k for a sample of points on a fast fiber $F = \theta^{-1}(\{v\})$. For $k = 0$, every point z_i projects to the same value in \mathbb{R} . As k increases, the distribution becomes more normal as the values $\theta(T^k(z_i))$ spread out. However, as k becomes very large, the distribution of the points starts to approach the invariant density of the full system, and this density, when projected by θ , will in general not be normal. From the graph, a value of $k = 10^5$ iterates gives a distribution which is close to normal. That the test statistic is low here (~ 0.02) gives empirical support to the assumption that the slow dynamics can be well approximated by a drift-diffusion process.

For simplicity, we assume that the same value k can be used for all fibers, as is the case for this example. For Step 5 of Algorithm 2, we used 500 fibers with the value v varying between -4 and 4 . Each fiber was populated with $q = 10^4$ points. The resulting functions α_k and β_k are plotted in Figure 5.3.

If k iterations corresponds to a time interval $k\Delta t$, then α and β in (5.1) can be approximated by $\alpha_k/k\Delta t$ and $\beta_k/k\Delta t$. Note that while the diffusion function β would be constant for the homogenization limit of (4.4)–(4.5) as $\epsilon \rightarrow 0$, the computed function β_k in Figure 5.3 is non-constant because it is computed using a non-zero time-step $k\Delta t$ which causes enhanced diffusion near the saddle point which separates




FIG. 5.3. The drift α_k and diffusion β_k computed in Algorithm 2 for the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$.

the two metastable states near $x = \pm 1$.

The drift-diffusion process (5.1) can now be used to simulate the slow dynamics using a time step which would be too large to simulate the full dynamics. Here, we use a time step of $dt = 0.02$ and calculate the orbit by

$$(5.2) \quad v_{n+1} = v_n + \alpha(v_n) dt + \sqrt{\beta(v_n)} dt r_n$$

where the r_n are independent pseudo-random values drawn from the standard normal distribution. Such a simulated orbit is plotted in Figure 5.4. This simulated orbit closely resembles the slow x variable of the original system (shown in Figure 4.1).

5.1.2. A distorted example. Finally, we use Algorithm 2 to simulate a slow orbit of the distorted system of Section 4.3.4. The computed functions for the drift




FIG. 5.4. Simulation of an orbit of the reduced one-dimensional dynamics given by Algorithm 2 for the system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$. The orbit was calculated using (5.2) with $dt = 0.02$.

α_k and diffusion β_k are plotted in Figure 5.5 and a simulated orbit of the reduced system is plotted in Figure 5.6. As in Section 4.3.4, our algorithm is insensitive to non-trivial coordinate transformations and accurately distills the slow variables and their dynamics.

5.2. Statistics. From a visual comparison of Figures 4.1, 5.4, and 5.6, it is not entirely clear if the reduced dynamics of Algorithm 2 is accurately capturing behaviour of the slow variable x in (4.4)–(4.5). To be more certain, we compute statistics related to these time series.

As argued in Section 4, the eigenfunction ϕ of the Koopman operator is approximately equal to $\psi \circ \mathcal{P}$ where ψ is an eigenfunction for the reduced dynamics and $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$ is the mapping down to the reduced space. Algorithm 2 replaces ϕ with a function θ which has similar level sets. Therefore, θ should be of the form $\theta = h \circ \mathcal{P}$ for some unknown h . Because of this, when comparing time series of the full dynamics with time series for the reduced dynamics, one should only use statistics which are unchanged when \mathcal{P} is replaced with $h \circ \mathcal{P}$ where h is a homeomorphism of the real line.

In our case, an obvious choice for such a statistic is the average switching time between the two metastable states. From a time series $\{v_i\}_{i=1}^M$, such as those plotted in Figures 4.1, 5.4, and 5.6, compute percentiles v^- and v^+ such that 40% of the points in $\{v_i\}$ are below v^- and 60% are below v^+ . Then, declare each point v_i to be in either the “high” state or the “low” state by the following heuristic:

- if $v_i \geq v^+$, then v_i is in the high state,
- if $v_i \leq v^-$, then v_i is in the low state,
- otherwise, v_i is in the same state as v_{i-1} .

This gives us a way to measure the switching times which is unaffected by applying a homeomorphism $h : \mathbb{R} \rightarrow \mathbb{R}$ to the time series.

For the full dynamics (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$, we computed a time series $\{x_i\}_{i=1}^M$ of the x -variable with $M = 2 \times 10^6$ where each step x_i to x_{i+1} represents a change of time of $\tau = 0.2$. This involved 2×10^{11} individual steps of the Euler-Maruyama method, using $\Delta t = 2 \times 10^6$. For the reduced dynamics, a similar time series $\{v_i\}_{i=1}^M$ with the same $M = 2 \times 10^6$ and $\tau = 0.2$ was obtained. Each step v_i to v_{i+1} was computed by ten applications of (5.2), each with a time step of $dt = 0.02$.

For the full time series, there were 2222 switches between the slow metastable states, with a mean switching time of $\tau_0 = 180.0$. For the reduced dynamics of

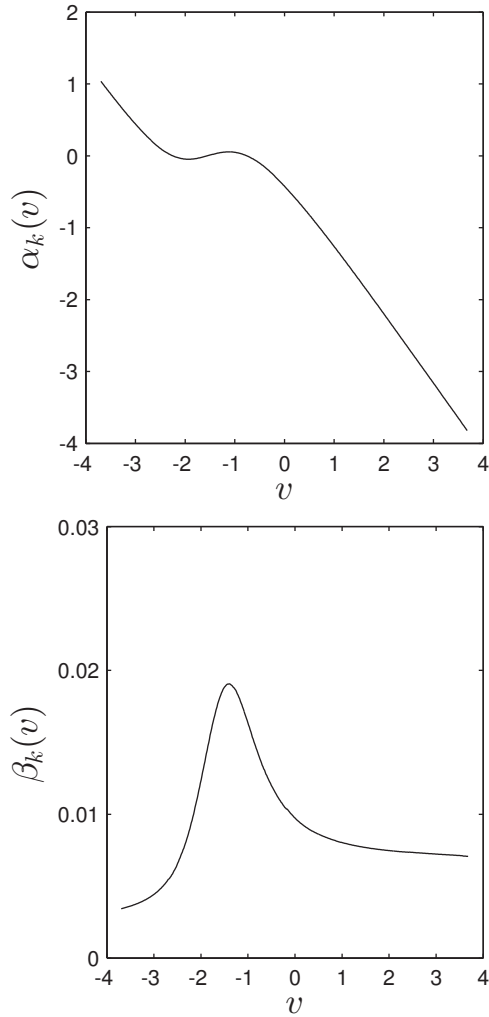


FIG. 5.5. The drift α_k and diffusion β_k computed in Algorithm 2 for system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$ after the distortion g described in Section 4.3.4.




FIG. 5.6. Simulation of an orbit of the reduced one-dimensional dynamics given by Algorithm 2 for system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$ after the distortion g described in Section 4.3.4. The orbit was calculated using (5.2) with $dt = 0.02$.




FIG. 5.7. *Log plot of the normalized histogram of switching times τ_i . This is computed for time series of the x variable of (4.4)–(4.5) as in Section 4.3.2 (red squares), for the reduced dynamics in Section 5.1.1 (green triangles) and for the reduced dynamics of the distorted system in Section 5.1.2 (blue diamonds).*

Section 5.1.1, there were 2280 switches with a mean time of $\tau_0 = 175.5$. In the distorted case in Section 5.1.2, there were 2178 switches with a mean time of $\tau_0 = 183.7$.

In each case, we expect the switching times τ_i to obey a Poisson process, where the cumulative probability function is given by the formula

$$P_c(\tau_i) = 1 - \exp\left(-\frac{\tau_i}{\tau_0}\right),$$

where τ_0 is the mean switching time. To see if this is the case, the computed switching times τ_i are plotted against $-\log(1 - P_c(\tau_i))$, as shown in Figure 5.7. For each of the three time series, the resulting plot is linear, and a least squares fit gives a computed value for $\tau_0 = 177.3$ for the full dynamics $\{x_i\}$ of Section 4.3.2, $\tau_0 = 178.8$ for the reduced dynamics of Section 5.1.1, and $\tau_0 = 182.2$ for the reduced dynamics of the distorted system in Section 5.1.2. This suggests that Algorithm 2 generates dynamics which closely resemble the relevant dynamics of the full system.

6. Computing the eigenfunctions. The techniques in this paper use spectral properties of the transfer operator, and its adjoint, the Koopman operator. Therefore, we must approximate these operators numerically. There are a number of approaches that could be used, and specialized methods other than Ulam, may be significantly faster in certain settings, depending on the smoothness of the system. These include numerical methods based on the infinitesimal generator of Ulam’s method [28], finite-element methods [2], and Galerkin methods using Fourier or Chebyshev bases [8]. See also Chapter 5 of [45]. For simplicity, we employ the standard Ulam’s method [70, 49] for all of the examples in this paper. Eigenpairs and the level sets were computed using the `eigs` and `contour` routines of MATLAB [54].

Recall that for a deterministic dynamical system $T : \mathcal{Z} \rightarrow \mathcal{Z}$, the transfer (or Perron-Frobenius) operator \mathcal{L} on $L^2(\mathcal{Z})$ is defined by

$$(6.1) \quad \int g \cdot \mathcal{L}f \, dm = \int (g \circ T) \cdot f \, dm$$

for all f, g in $L^2(\mathcal{Z})$. To solve numerically, Ulam's method divides the domain \mathcal{Z} into a finite number of boxes B_1, \dots, B_n and reduces the problem to

$$(6.2) \quad \int \mathbb{1}_{B_i} \cdot \mathcal{L}\mathbb{1}_{B_j} \, dm = \int (\mathbb{1}_{B_i} \circ T) \cdot \mathbb{1}_{B_j} \, dm$$

for $i, j = 1, \dots, n$ where $\mathbb{1}_{B_i}$ is the indicator function for the set B_i . The resulting operator, now defined on the finite dimensional space spanned by the functions $\mathbb{1}_{B_i}$, is equivalent to an $n \times n$ matrix defined by

$$P_{ij} = \frac{m(T^{-1}(B_i) \cap B_j)}{m(B_i)}.$$

To estimate P_{ij} numerically, populate each B_j with a large number N of randomly sampled test points $z_{j,k}$ for $k = 1, \dots, N$ and approximate

$$(6.3) \quad P_{ij} \approx \frac{1}{N} \#\{k : T(z_{j,k}) \in B_i\}.$$

In the case of a random dynamical system, the dynamics is defined by a map $T : \mathcal{Z} \times \Omega \rightarrow \mathcal{Z}$ where Ω is a probability space, sometimes called the "noise space" of the system. Each $\omega \in \Omega$ defines a function $T_\omega = T(\cdot, \omega)$ and one can define a transfer operator \mathcal{L}_ω for each T_ω . We then consider the annealed (or integrated) operator $\mathcal{L} : L^2(\mathcal{Z}) \rightarrow L^2(\mathcal{Z})$ defined by

$$\mathcal{L}(f) = \int_{\Omega} \mathcal{L}_\omega(f) \, d\omega.$$

Similarly, the annealed Koopman operator is

$$(\mathcal{K}f)(z) = \int_{\Omega} f(T(z, \omega)) \, d\omega$$

which for a point z , gives the expected value of $f(T(z))$. For a random dynamical system, Ulam's method produces a matrix which can be approximated numerically by a slightly modified version of (6.3) with

$$(6.4) \quad P_{ij} \approx \frac{1}{N} \#\{k : T_{\omega_{j,k}}(z_{j,k}) \in B_i\},$$

where for each $z_{j,k} \in B_j$ an element $\omega_{j,k} \in \Omega$ of the noise space is chosen randomly and independently.

The examples in this paper are defined by stochastic differential equations. In this case, the noise space Ω is the space of continuous paths $C^0(\mathbb{R}, \mathbb{R}^d)$ where the probability measure is given by the Wiener process representing Brownian motion. Once a path $W \in \Omega = C^0(\mathbb{R}, \mathbb{R}^d)$ is chosen, an SDE of the form

$$dx_t = \mu(x)dt + \sigma dW_t$$

at starting value $x(0) = x_0$ is solved by a path satisfying the integral

$$x(t) = x_0 + \int_0^t \mu(x(s)) ds + \sigma W(t)$$

at all times $t > 0$. For a specified flow time $\tau > 0$, this gives a deterministic map $T_W(x_0) = x(\tau)$ for each $W \in \Omega$ and so defines a random dynamical system in the form $T : \mathcal{Z} \times \Omega \rightarrow \mathcal{Z}$. For each point $z_{j,k}$ in (6.4), the image $T_{\omega_{j,k}}(z_{j,k})$ can be computed, for example, by the Euler-Maruyama method [44], and so the annealed operators for these types of systems can be approximated numerically.

For these continuous-time dynamical systems, one must choose some interval of time $\tau > 0$ in order to apply Ulam’s method to the map $T = \varphi_\tau$ given by the flow $\{\varphi_t\}_{t \in \mathbb{R}}$. Picking too small a flow time can lead to a dynamical system appearing to have more diffusion than is actually present, a problem closely related to so-called “numerical diffusion” which appears in finite-difference methods [28]. A common practice is to choose τ large enough, so that for a significant number of sample points $x_{i,k} \in B_i$, the image $T(x_{i,k}) = \varphi_\tau(x_{i,k})$ is no longer inside the box B_i [29]. In other words, so that P_{ii} is not too close to one.

This works well for many systems, but it is not a sufficient criterion for a system with multiscale behaviour. For these systems, choosing a flow time in such manner only guarantees that Ulam’s method captures the fastest dynamics, and there may be significant numerical diffusion present in the direction of the slow dynamics, meaning that the computed invariant density and other eigenfunctions may have significant errors.

The effect is more pronounced when the slow and fast directions do not align with the grid of boxes. Here, Ulam’s method cannot faithfully capture motion happening purely in the fast direction. In some sense, the Galerkin projection “smears” the point $T(x_{i,k})$ over the box B_j in which it resides and this smearing happens in both slow and fast directions. As an example, we use Ulam’s method for the system (4.4)–(4.5) after distortion as defined in Section 4.3.4 with an extremely short time-step of $\tau = 2 \times 10^{-5}$. In the resulting Ulam matrix, the maximum value P_{ii} on the diagonal is less than 0.11. However, the computed density, plotted in Figure 6.1, shows two thick spurious bands, which would be more consistent with a system with a large amount of diffusion in the slow direction (cf. the invariant density in Figure 4.5 computed with significantly larger flow time $\tau = 0.04$).

One way to avoid this is to increase the flow time so that a typical point moves in the slow direction by an amount comparable to the size of the boxes B_i . However, we are assuming that we do not know in advance the directions of the slow and fast dynamics. To handle this, we compute unit eigenvectors using Ulam’s method for a sequence of increasing flow times, and stop when the result stabilizes. For instance, in the axis-aligned example system (4.4)–(4.5) in Section 4.3.2, the invariant density $\rho_\tau \in \text{span}\{\chi_1, \dots, \chi_n\} \subset L^1(\mathbb{R}^2)$ was computed for flow times $\tau = 0.01, 0.02$, and 0.04. Using the L^1 norm $\|f\| = \int_{\mathbb{R}^2} |f(x)| dx$, we compute $\|\rho_{0.01} - \rho_{0.02}\| \cong 0.08$ and $\|\rho_{0.02} - \rho_{0.04}\| \cong 0.03$. More important than the density is the second leading eigenfunction of the Koopman operator, i.e., the eigenfunction shown in Figure 4.3, since this function is used to define the fiber dynamics. The corresponding L^1 differences for this computed eigenfunction are $\|\phi_{0.01} - \phi_{0.02}\| \cong 0.04$ and $\|\phi_{0.02} - \phi_{0.04}\| \cong 0.01$. We view this difference as small enough that we can safely use $\tau = 0.04$ in the computation of the eigenfunctions.

Because we are using the Euler-Maruyama method with a time-step of 2×10^{-6} , this value of τ corresponds to 2×10^4 iterations. However, this value of τ is still much




FIG. 6.1. An incorrect density calculated using Ulam's method with too short a flow time for system (4.4)–(4.5) with $\epsilon = 0.01$, $a = 0.02$, and $\sigma^2 = 0.113$ after the distortion g defined in Section 4.3.4. The flow time used to construct the Ulam matrix is $\tau = 2 \times 10^{-5}$. Compare this to the computed density in Figure 4.5.

smaller than the average time an orbit takes to switch between the metastable states based at $x = -1$ and $x = +1$. That switching happens on the order of hundreds of units of time (as shown in Figure 4.1).

As a further test, we also consider the computed eigenvalues. As argued in Section 3, the leading eigenvalues of the full dynamics correspond to decay rates associated to the slow dynamics. Therefore, if many of the computed eigenvalues are very close to one, it suggests that the flow time may be too short. For $\tau = 0.04$ in the example system, the leading ten computed eigenvalues (as listed in Table 4.1) range from 1.000 to 0.6969 which is far enough from 1 to tell us that advection in the slow direction is being captured in the Ulam matrix.

In addition to choosing the flow time, one must also set the number of sample points N in (6.3) or (6.4). Because of the relatively long flow time, starting from one box P_{ij} the evolved points $T(x_{i,k})$ may lie in a large number of boxes B_j , and N must be chosen large enough that each P_{ij} is computed to a reasonable accuracy. In all our examples, we use $N = 10^4$ and a 200 by 200 grid on the region $[-4, 4] \times [-4, 4] \subset \mathbb{R}^2$. This is a much larger region than where orbits typically lie. However, we compute everything in this larger region so that when fiber dynamics are computed, the orbit will stay in the domain with very high probability.

Once a level set is computed, the same heuristics using L^1 -distance of the computed eigenfunctions and the distance of the eigenvalues from 1 may be used to determine a reasonable number of iterations to apply for the fiber dynamics. These showed that 20 iterations for a total flow time of $20 \times 2 \times 10^{-6} = 4 \times 10^{-5}$ was sufficient for the fiber dynamics of the example system (4.4)–(4.5) in Section 4.3.2.

An eigenfunction ϕ computed using Ulam's method will be given as a linear combination of the indicator functions $\mathbb{1}_{B_i}$. As such functions are piecewise constant,

the gradient $\nabla\phi$ will therefore be zero anywhere it is defined. This gradient is not useful in defining a gradient flow π as used in Step 4 of Algorithm 1. Instead, when computing $\nabla\phi$ for use in the algorithm we replace ϕ by the bilinear interpolation of values it takes on the centers of each box B_i . Further, as solving for the gradient flow is computationally expensive, we employ the following shortcut. To compute the map π at a point z close to the fast fiber, evaluate the gradient at z , then take $\pi(z)$ to be the intersection of the fiber $\phi^{-1}(\{v\})$ with the line given by $z + \mathbb{R} \cdot \nabla\phi(z)$. Since in computing the fiber dynamics the function π is only evaluated at points very close to the level set, the gradient is nearly constant on the short path from z to $\pi(z)$ and so the approximation is valid.

Alternatively, at least in principle, one could have used the eigenfunctions of the transfer operator \mathcal{L} instead of those of the Koopman operator \mathcal{K} . If μ is an invariant density for the dynamical system T , then (6.1) with $m = \mu$ defines a transfer operator \mathcal{L}_μ and the invariance of the measure implies that a constant function $f \equiv c \in \mathbb{R}$ satisfies $\mathcal{L}_\mu(f) = f$. Numerical experiments suggest that, similar to the Koopman operator, leading eigenfunctions of \mathcal{L}_μ are nearly constant in the fast direction of a multiscale system. The problem with this approach is that one must numerically approximate the invariant measure μ , say by a density function, ρ , and use this to compute eigenfunctions of \mathcal{L}_μ . Numerically, this is equivalent to calculating the pointwise quotient ϕ/ρ where ϕ is an eigenvalue of the transfer operator \mathcal{L} defined with respect to Lebesgue measure m . For systems which involve regions of low density with ρ close to zero, this is intractable to do numerically.

7. Discussion and outlook. We have used the transfer operator and the Koopman operator to disentangle the multiple slow and fast scales and provide an effective set of slow variables which can then be efficiently numerically integrated in time. We constructed a projection onto a slow subspace using the eigenfunctions of the Koopman operator. We then devised an algorithm to estimate the drift and diffusion terms of an assumed stochastic differential equation describing the reduced slow dynamics. This constitutes a huge computational advantage as we were able to use time steps of the order of 10^4 larger to propagate forward the slow variables when compared to the full multiscale system.

There are similarities of our approach with isocommittor surfaces in transition path theory (see for example [24] for a review) and with reaction coordinates for constrained molecular dynamics (see for example [38, 39, 47, 48]). Isocommittor surfaces separate metastable states and these methods aim at calculating the associated free energy in order to determine statistics such as most likely transitions and transition rates between the meta-stable states. It was, however, noted in [22] that reaction coordinates and isocommittor functions are not necessarily slow variables. In our method the slow and fast variables are first identified, and then subsequently the reduced dynamics is provided for the slow variables, whereas reaction coordinates and isocommittor surfaces are constructed to obtain statistics of transitions between predefined sets such as meta-stable states.

In the current paper, the algorithms were described and tested in a low-dimensional setting with one-dimensional slow dynamics and two time scales. As an outlook for future research, we now discuss how one might extend these techniques to more complicated situations.

Extension to systems with multi-dimensional slow subspace: In this paper, we considered multiscale examples with one-dimensional slow dynamics, so that the mapping $\mathcal{P} : \mathcal{Z} \rightarrow \mathcal{X}$ could be approximated by a single eigenfunction $\phi : \mathcal{Z} \rightarrow \mathbb{R}$ of the Koopman operator. We outline now a possible extension of our computational method to the case of a higher-dimensional slow subspace. Details are planned for further research. For a system with higher-dimensional slow dynamics, it will be necessary to approximate \mathcal{P} by a product of eigenfunctions

$$\mathcal{P} = \phi_1 \times \phi_2 \times \cdots \times \phi_m : \mathcal{Z} \rightarrow \mathbb{R}^m.$$

Then, a value $v \in \mathbb{R}^m$ defines a fiber $F = \mathcal{P}^{-1}(\{v\}) \subset \mathcal{Z}$. The fiber dynamics can be defined as $\hat{T} : F \rightarrow F$ where $\hat{T} = \pi \circ T$ and $\pi : U \rightarrow F$ is a mapping from a neighbourhood U down to the fiber itself. As explained in Section 6, for a single eigenfunction ϕ the numerical computation of $\pi(z)$ was given by the intersection of F with the line through z tangent to $\nabla\phi(z)$. A possible way to extend this to the general setting is to define $\pi(z)$ as the intersection of F with the m -dimensional affine subspace through z spanned by the vectors $\nabla\phi_i(z)$ for $i = 1, \dots, m$. Then, as in Section 4.3.2, one could test the validity of the fiber dynamics by comparing the distance $\|\hat{T}(z) - T(z)\|$ to the distance $\|T(z) - z\|$ for an ensemble of points on F .

If the dimension of the slow dynamics is not known beforehand, an automated method to find the dimension would be to test products of an increasing number of eigenfunctions until a product is found for which \hat{T} gives a reasonable approximation of fast dynamics. Using this definition of \hat{T} , Algorithm 1 could be implemented much as before.

Algorithm 2 might also be generalized to use a product of eigenvalues. In such a case, the $\alpha_k(v)$ would be computed as the m -dimensional mean of an ensemble of the form $\{\theta(T^k(z_i) - v)\}_{i=1}^q$ using a product $\theta = \theta_1 \times \cdots \times \theta_m$ and $\beta_k(v)$ would be computed as the $m \times m$ covariance matrix.

The main difficulty which would arise in implementing these extensions of the algorithms to high-dimensional systems is accurate and efficient numerical approximation of the transfer operator. Currently, techniques based on Ulam's method are mainly restricted to low-dimensional systems.

Extension to systems with more than two time scales: In this paper, we considered dynamical systems with one fast and one slow time scale. We now suggest an iterative procedure using Algorithms 1 and 2 (or their higher-dimensional variants) which distills the dynamics on several time scales. For ease of exposition, we present this procedure for a system T with three time scales (fast, moderate and slow). Applying Algorithm 1 to such a system produces fiber dynamics $\hat{T} : F \rightarrow F$ which incorporates the dynamics of both the fast and moderate time scales. The slowest time scales can be obtained by applying Algorithm 2 to T . To further isolate the moderate and fast dynamics, apply Algorithms 1 and 2 now to the fiber dynamics \hat{T} in place of the full system T .

Solving for the slow variables without having to integrate forward in time the fast variables: We defined the projection to the slow subspace via the eigenfunctions of the Koopman operator. This requires the integration in time of the full dynamics, including the fast variables, needed for Ulam's method. The transfer operator \mathcal{L} and the infinitesimal generator \mathcal{A} share the same eigenfunctions and this has been exploited to find these eigenfunctions numerically for certain systems without trajectory integration [28, 45]. Thus, it may be possible to reformulate the steps in

Algorithm 1 and 2 based purely on analysing the vector field itself without numerical integration, yielding a significant speed up in computation.

Acknowledgments. The illustrations in Figures 2.1 and 5.1 were produced using the Asymptote vector-graphics language (<http://asymptote.sf.net>). We acknowledge the support by the Australian Research Council under grant DP120104514.

REFERENCES

- [1] V. ARNOLD, V.V. KOZLOV, AND A.I. NEISHTADT, *Mathematical Aspects of Classical and Celestial Mechanics*, Springer-Verlag, New York, 1993.
- [2] I. BABUŠKA AND J. E. OSBORN, *Finite element-Galerkin approximation of the eigenvalues and eigenvectors of selfadjoint problems*, Math. Comp., 52 (1989), pp. 275–297.
- [3] V. BALADI AND L.-S. YOUNG, *On the spectra of randomly perturbed expanding maps*, Comm. Math. Phys., 156 (1993), pp. 355–385.
- [4] NILS BERGLUND AND BARBARA GENTZ, *Noise-Induced Phenomena in Slow-Fast Dynamical Systems*, Springer, 2005.
- [5] MICHAEL BLANK, GERHARD KELLER, AND CARLANGLO LIVERANI, *Ruelle-Perron-Frobenius spectrum for Anosov maps*, Nonlinearity, 15 (2002), pp. 1905–1973.
- [6] CHRIS BOSE AND GARY FROYLAND, *Strange eigenmodes for invertible measure-preserving systems*. in preparation.
- [7] PETRA BOXLER, *A stochastic version of the centre manifold theorem*, Probab. Th. Rel. Fields, 83 (1989), pp. 509–545.
- [8] JOHN P. BOYD, *Chebyshev and Fourier spectral methods*, Dover Publications Inc., Mineola, NY, second ed., 2001.
- [9] J. CARR, *Applications of Centre Manifold Theory*, no. 35 in Applied Mathematical Sciences, Springer, 1981.
- [10] A.J. CHORIN AND O.H. HALD, *Stochastic Tools in Mathematics and Science*, STAMS, Springer-Verlag, New York, 2006.
- [11] A.J. CHORIN, O.H. HALD, AND R. KUPFERMAN, *Optimal prediction and the Mori-Zwanzig representation of irreversible processes*, Proceedings of the National Academy of Science, 97 (2000), p. 29682973.
- [12] R. COIFMAN, I. KEVREKIDIS, S. LAFON, M. MAGGIONI, , AND B. NADLER, *Diffusion maps, reduction coordinates and low dimensional representation of stochastic systems*, SIAM Multiscale Modeling and Simulation, 7 (2008), pp. 842–864.
- [13] R. COIFMAN AND S. LAFON, *Diffusion maps*, Applied and Computational Harmonic Analysis, 21 (2006), pp. 5–30.
- [14] D. CROMMELIN AND E. VANDEN-EIJNDEN, *Diffusion estimation from multiscale data by operator eigenpairs*, Multiscale Modeling & Simulation, 9 (2011), pp. 1588–1623.
- [15] D. T. CROMMELIN AND E. VANDEN-EIJNDEN, *Reconstruction of diffusions using spectral data from timeseries*, Communications in Mathematical Sciences, 4 (2006), pp. 651 – 668.
- [16] MICHAEL DELLNITZ, GARY FROYLAND, AND STEFAN SERTL, *A conjecture on the existence of isolated eigenvalues of the Perron-Frobenius operator*, in International Conference on Differential Equations, Vol. 1, 2 (Berlin, 1999), World Sci. Publ., River Edge, NJ, 2000, pp. 1030–1032.
- [17] MICHAEL DELLNITZ AND OLIVER JUNGE, *On the approximation of complicated dynamical behavior*, SIAM J. Numer. Anal., 36 (1999), pp. 491–515.
- [18] M. DELLNITZ AND O. JUNGE, *On the approximation of complicated dynamical behaviour*, SIAM J. Numer. Anal., 36 (1999), pp. 491–515.
- [19] P. DEUFLHARD, W. HUISINGA, A. FISCHER, , AND C. SCHÜTTE, *Identification of almost invariant aggregates in nearly uncoupled Markov chains*, SIAM J. Numer. Anal., 38 (2000), p. 3959.
- [20] P. DEUFLHARD AND C. SCHÜTTE, *Molecular conformation dynamics and computational drug design*, in Applied Mathematics Entering the 21st Century, J. M. Hill and R. Moore, eds., ICIAM 2003, Sydney, 2004.
- [21] WEINAN E, *Analysis of the heterogeneous multiscale method for ordinary differential equations*, Comm. Math. Sci., 1 (2003), pp. 423–436.
- [22] ———, *Principles of Multiscale Modelling*, Cambridge University Press, 2011.
- [23] WEINAN E, BJÖRN ENQUIST, X. LI, W. REN, AND ERIC VANDEN-EIJNDEN, *Heterogeneous multiscale methods: A review*, Comm. Comp. Phys., 2 (2007), pp. 367–450.
- [24] WEINAN E AND ERIC VANDEN-EIJNDEN, *Metastability, conformation dynamics, and transition*

- pathways in complex systems*, in Multiscale Modelling and Simulation, Sabine Attinger and Petros Koumoutsakos, eds., vol. 39 of Lecture Notes in Computational Science and Engineering, Springer Berlin Heidelberg, 2004, pp. 35–68.
- [25] GARY FROYLAND, *Computer-assisted bounds for the rate of decay of correlations*, Comm. Math. Phys., 189 (1997), pp. 237–257.
- [26] G. FROYLAND, *Statistically optimal almost-invariant sets*, Physica D, 200 (2005), p. 205219.
- [27] GARY FROYLAND, *On Ulam approximation of the isolated spectrum and eigenfunctions of hyperbolic maps*, Discrete Contin. Dyn. Syst., 17 (2007), pp. 671–689 (electronic).
- [28] GARY FROYLAND, OLIVER JUNGE, AND PÉTER KOLTAI, *Estimating long-term behavior of flows without trajectory integration: the infinitesimal generator approach*, SIAM J. Numer. Anal., 51 (2013), pp. 223–247.
- [29] GARY FROYLAND AND KATHRIN PADBERG, *Almost-invariant sets and invariant manifolds—connecting probabilistic and geometric descriptions of coherent structures in flows*, Phys. D, 238 (2009), pp. 1507–1523.
- [30] G. FROYLAND, K. PADBERG, M. ENGLAND, , AND A.M. TREGUIER, *Detection of coherent oceanic structures via transfer operators*, Physical Review Letters, 98 (2007), p. 224503.
- [31] G. FROYLAND, N. SANTITISSADEEKORN, AND A. MONAHAN, *Transport in time-dependent dynamical systems: Finite-time coherent sets*, Chaos, 20 (2010), p. 043116.
- [32] C. W GARDINER, *Handbook of Stochastic Methods for Physics, Chemistry, and the Natural Sciences*, Springer, New York, 3rd ed., 2003.
- [33] B. GAVEAU AND L. S. SCHULMAN, *Multiple phases in stochastic dynamics: Geometry and probabilities.*, Physical Review E, 73 (2006), p. 036124.
- [34] C.W. GEAR AND I. KEVREKIDIS, *Projective methods for differential equations*, SIAM J. Sci. Comp., 24 (2003), pp. 1091–1106.
- [35] DROR GIVON, RAZ KUPFERMAN, AND ANDREW STUART, *Extracting macroscopic dynamics: model problems and algorithms*, Nonlinearity, 17 (2004), pp. R55–R127.
- [36] G. A. GOTTWALD AND J. HARLIM, *The role of additive and multiplicative noise in filtering complex dynamical systems*, Proc. Roy. Soc. A, 469 (2013), p. 20130096.
- [37] GEORG A. GOTTWALD AND IAN MELBOURNE, *Homogenization for deterministic maps and multiplicative noise*, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science, 469 (2013).
- [38] CARSTEN HARTMANN AND CHRISTOF SCHÜTTE, *A geometric approach to constrained molecular dynamics and free energy*, Commun. Math. Sci., 3 (2005), pp. 1–20.
- [39] CARSTEN HARTMANN AND CHRISTOF SCHÜTTE, *Comment on two distinct notions of free energy*, Physica D: Nonlinear Phenomena, 228 (2007), pp. 59–63.
- [40] W. JUST, H. KANTZ, C. RÖDENBECK, AND M. HELM, *Stochastic modelling: replacing fast degrees of freedom by noise*, J. Phys. A, 34 (2001), pp. 3199–3213.
- [41] I. KEVREKIDIS AND G. SAMAËY, *Equation-free multiscale computation: algorithms and applications*, Ann. Rev. Phys. Chem., 60 (2009), pp. 321–344.
- [42] I. G KEVREKIDIS, C. W. GEAR, J. M. HYMAN, G. K. PANAGIOTIS, O. RUNBORG, AND C. THEODOROPOULOS, *Equation-free, coarse-grained multiscale computation: Enabling microscopic simulators to perform system-level analysis*, Comm. Math. Sci., 1 (2003), pp. 715–762.
- [43] R. Z. KHASMINSKY, *On stochastic processes defined by differential equations with a small parameter*, Theory of Probability and its Applications, 11 (1966), pp. 211–228.
- [44] P. E. KLOEDEN AND E. PLATEN, *Numerical methods for stochastic differential equations*, in Nonlinear dynamics and stochastic mechanics, CRC Math. Model. Ser., CRC, Boca Raton, FL, 1995, pp. 437–461.
- [45] P. KOLTAI, *Efficient approximation methods for the global long-term behavior of dynamical systems - Theory, algorithms and examples*, PhD thesis, Technische Universitt Mnchen, 2010.
- [46] T. G. KURTZ, *A limit theorem for perturbed operator semigroups with applications to random evolutions*, Journal of Functional Analysis, 12 (1973), pp. 55–67.
- [47] TONY LELIÈVRE, *Free Energy Computations: A Mathematical Perspective*, Imperial College Press, 2010.
- [48] TONY LELIÈVRE, MATHIAS ROUSSET, AND GABRIEL STOLTZ, *Langevin dynamics with constraints and computation of free energy differences*, Mathematics of Computation, 81 (2012), pp. 2071–2125.
- [49] TIEN YIEN LI, *Finite approximation for the Frobenius-Perron operator. A solution to Ulam’s conjecture*, J. Approximation Theory, 17 (1976), pp. 177–186.
- [50] H. W. LILLIEFORS AND C. C. MCBRIDE, *Kolmogorov-Smirnov and Anderson-Darling statistics with parameters estimated by linear combinations of order statistics and some robustness*

- properties*, in Goodness-of-fit (Debrecen, 1984), vol. 45 of Colloq. Math. Soc. János Bolyai, North-Holland, Amsterdam, 1987, pp. 369–376.
- [51] W. LIU AND G. HALLER, *Strange eigenmodes and decay of variance in the mixing of diffusive tracers*, *Physica D*, 188 (2004), pp. 1–39.
- [52] J. MACLEAN AND G. A. GOTTWALD, *On convergence of the projective integration method for stiff ordinary differential equations*, *Communications in Mathematical Sciences*, 12 (2014), pp. 235–255.
- [53] ANDREW J. MAJDA, ILYA TIMOFEYEV, AND ERIC VANDEN EIJNDEN, *A mathematical framework for stochastic climate models*, *Communications on Pure and Applied Mathematics*, 54 (2001), pp. 891–974.
- [54] MATLAB, *version 7.10.0 (R2010a)*, The MathWorks Inc., Natick, Massachusetts, 2010.
- [55] IAN MELBOURNE AND ANDREW STUART, *A note on diffusion limits of chaotic skew-product flows*, *Nonlinearity*, 24 (2011), pp. 1361–1367.
- [56] IGOR MEZIĆ, *Spectral properties of dynamical systems, model reduction and decompositions*, *Nonlinear Dynam.*, 41 (2005), pp. 309–325.
- [57] ———, *Analysis of fluid flows via spectral properties of the Koopman operator*, in Annual review of fluid mechanics. Volume 45, 2013, vol. 45 of Annu. Rev. Fluid Mech., Annual Reviews, Palo Alto, CA, 2013, pp. 357–378.
- [58] IGOR MEZIĆ AND ANDRZEJ BANASZUK, *Comparison of systems with complex behavior*, *Phys. D*, 197 (2004), pp. 101–133.
- [59] L. MITCHELL AND G. A. GOTTWALD, *Data assimilation in slow-fast systems using homogenized climate models*, *Journal of the Atmospheric Sciences*, 69 (2012), pp. 1359–1377.
- [60] H. MORI, H. FUJISAKA, AND H. SHIGEMATSU, *A new expansion of the master equation*, *Prog. Theor. Phys.*, 51 (1974), pp. 109–122.
- [61] N. NAMACHCHIVAYA AND G. LENG, *Equivalence of stochastic averaging and stochastic normal forms*, *J. Appl. Mech.*, 57 (1990), pp. 1011–1017.
- [62] N. NAMACHCHIVAYA AND Y.K. LIN, *Method of stochastic normal forms*, *Int. J. Non-Linear Mechanics*, 26 (1991), pp. 1011–1017.
- [63] G C PAPANICOLAOU, *Some probabilistic problems and methods in singular perturbations*, *Rocky Mountain Journal of Mathematics*, 6 (1976), pp. 653–674.
- [64] GRIGORIS A PAVLIOTIS AND ANDREW M STUART, *Multiscale Methods: Averaging and Homogenization*, Springer, New York, 2008.
- [65] ANTHONY J. ROBERTS, *Normal form transforms separate slow and fast modes in stochastic dynamical systems*, *Physica A*, 387 (2008), pp. 12–38.
- [66] J.A. SANDERS AND F. VERHULST, *Averaging methods in nonlinear dynamical systems*, vol. 59 of Applied Mathematical Sciences, Springer-Verlag, New York, 1985.
- [67] C. SCHÜTTE, W. HUISINGA, AND P. DEUFLHARD, *Transfer operator approach to conformational dynamics in biomolecular systems*, in Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems, B. Fiedler, ed., Springer, Berlin, 2001, pp. 191–223.
- [68] PETER SJÖGREN, *Operators associated with the Hermite semigroup—a survey*, in Proceedings of the conference dedicated to Professor Miguel de Guzmán (El Escorial, 1996), vol. 3, 1997, pp. 813–823.
- [69] P. STINIS, *Higher order Mori-Zwanzig models for the Euler equations*, *SIAM Multiscale Modeling and Simulation*, 6 (2007), pp. 741–760.
- [70] S. M. ULAM, *A collection of mathematical problems*, Interscience Tracts in Pure and Applied Mathematics, no. 8, Interscience Publishers, New York-London, 1960.
- [71] J. WALTER AND C. SCHÜTTE, *Conditional averaging for diffusive fast-slow systems: A sketch for derivation*, in Analysis, Modeling and Simulation of Multiscale Problems, A. Mielke, ed., Springer, Berlin, 2006, pp. 647–682.
- [72] R. ZWANZIG, *Ensemble method in the theory of irreversibility*, *J. Chem. Phys.*, 33 (1960), pp. 1338–1341.
- [73] ROBERT ZWANZIG, *Nonequilibrium Statistical Mechanics*, Oxford University Press, Oxford, 2001.