

Neighborhoods of periodic orbits and the stationary distribution of a noisy chaotic system

Jeffrey M. Heninger,¹ Domenico Lippolis,^{2,*} and Predrag Cvitanović³

¹*Department of Physics, University of Texas, Austin, Texas 78712, USA*

²*Institute for Advanced Study, Tsinghua University, Beijing 100084, China*

³*Center for Nonlinear Science and School of Physics, Georgia Institute of Technology, Atlanta, Georgia 30332, USA*

(Received 17 July 2015; published 23 December 2015)

The finest state-space resolution that can be achieved in a physical dynamical system is limited by the presence of noise. In the weak-noise approximation, the stochastic neighborhoods of deterministic periodic orbits can be computed from distributions stationary under the action of a local Fokker-Planck operator and its adjoint. We derive explicit formulas for widths of these distributions in the case of chaotic dynamics, when the periodic orbits are hyperbolic. The resulting neighborhoods form a basis for functions on the attractor. The global stationary distribution, needed for calculation of long-time expectation values of observables, can be expressed in this basis.

DOI: [10.1103/PhysRevE.92.062922](https://doi.org/10.1103/PhysRevE.92.062922)

PACS number(s): 05.45.-a, 45.10.Db, 45.50.Pk, 47.11.-j

I. INTRODUCTION

This paper investigates the interplay of deterministic chaotic dynamics and weak stochastic noise, and proposes a definition of the neighborhood of a noisy hyperbolic state-space point. Such neighborhoods are conjectured to partition the state space in an optimal way and provide a basis function set for the evaluation of the stationary distribution.

A. Width of a noisy trajectory

The basic idea of a stochastic “neighborhood” is that the balance between the noise broadening of a trajectory and the deterministic contraction leads to a probability distribution of finite width, as opposed to one that spreads with time (diffusion only). For an orbit that converges to a linearly stable, attractive equilibrium, this neighborhood was computed in 1810 by Laplace [1,2] and is today known as a solution to the Lyapunov equation [3], or the Ornstein-Uhlenbeck process [4]: for a one-dimensional flow, the deterministic equilibrium point is smeared into a Gaussian probability density centered on it, whose covariance $Q = -D/\lambda$ is a balance of the expansion rate D (diffusion constant) against the contraction rate $\lambda < 0$. Fokker-Planck equation [5] generalizations to higher-dimensional stable equilibria and limit cycles (stable periodic orbits) are immediate, provided proper care is taken of the diffusion along the periodic orbit [6,7].

What if a periodic orbit is unstable? Both the diffusion rate and the linearized stability rate $\lambda > 0$ now expand forward in time and cannot balance each other. This problem was solved in Refs. [8,9] for repelling periodic orbits with *no contracting directions*, by balancing the stochastic diffusion against the contraction by the *adjoint* Fokker-Planck operator. The resulting covariance matrix defines the stochastic neighborhood for a repelling orbit, while the Ornstein-Uhlenbeck covariance defines it for a stable orbit. However, neither these stable nor repelling orbits play a role in chaotic dynamics. The long-time attractors of chaotic dynamics are organized by an infinity of *hyperbolic* periodic orbits [10–12], i.e., orbits

which an ergodic trajectory visits by approaching them along their stable eigendirections and leaves along their unstable eigendirections.

The central result of this paper is that techniques developed for solving the Lyapunov equations [13–15] enable us to define the neighborhood of a *hyperbolic* periodic point by splitting the covariance matrix Q into two (mutually nonorthogonal) covariance matrices: Q_{cc} for the contracting directions and Q_{ee} for the expanding directions.

There are two immediate applications of the notion of the neighborhood of a hyperbolic point: (a) “optimal partition” of the attractor, and (b) construction of a basis set for the stationary distribution of a noisy chaotic flow.

B. An optimal partition from periodic orbits

While in the idealized deterministic dynamics the state space can be resolved arbitrarily finely, in physical systems noise always limits the attainable state-space resolution.

This observation had motivated the many limiting resolution estimates for state-space granularity of chaotic systems with background noise. The idea of an optimal partition in this context was first introduced in 1983 by Crutchfield and Packard [16], who formulated a state-space resolution criterion in terms of a globally averaged “attainable information.” The approach was later generalized and applied to time-series analysis, where the underlying dynamics is unknown [17,18]. A different strategy consists of computing a transfer matrix between intervals of a uniform grid and estimating averages of observables from its eigenvalues and eigenfunctions. First introduced by Ulam [19], this technique has been developed over the years [20,21]. All of these approaches (see Ref. [9] for a review) are based on *global* averages and assume that granularity is uniform across the state space. In contrast, the main, computationally precise lesson of our work is that even when the external noise is white, additive, and globally homogenous, the interplay of noise and nonlinear dynamics always results in a *local* stochastic neighborhood, whose covariance depends on both the past and the future noise integrated and nonlinearly convolved with deterministic evolution along the trajectory. The optimal resolution thus varies from neighborhood to neighborhood and has to be computed *locally*. As was shown in Ref. [8] for a strictly

*Present address: Faculty of Science, Jiangsu University, Zhenjiang 212013, China; domenico@ujs.edu.cn

expanding $1d$ chaotic map and a given noise, the maximal set of nonoverlapping neighborhoods of periodic orbits can be used to construct an *optimal partition* of the state space, and compute dynamical averages from the associated approximate matrix Fokker-Planck operator.

C. The stationary distribution

In this paper, we utilize our construction of optimal partitions to approximate the *stationary probability distribution function* by a finite sum over Gaussians, one for each neighborhood. When the dynamics is chaotic, the most one can predict accurately for long times are the statistical properties of the system, given by the state-space averages of observables $a(x)$,

$$\langle a \rangle = \int dx \rho(x) a(x), \quad (1)$$

where the *stationary distribution* (natural measure [22–24]) $\rho(x)$ is the probability of finding the system in the state x on the attractor. For a deterministic system, $\rho(x)$ is a singular, nowhere differentiable distribution with support on a fractal set, and its numerical computation is usually not feasible. However, for any physical system, the noise washes out fine details of the dynamics and the stationary distribution is smooth. Here we propose a smooth function basis for $\rho(x)$, based on the optimal partition of the state space. We develop our formalism for discrete-time dynamical systems and illustrate it by computing the neighborhoods and estimating $\rho(x)$ for the Lozi map [25], a simple two-dimensional discrete-time chaotic system. The idea is to first partition the attractor into an optimal partition set of neighborhoods, and then use the associated local Gaussian distributions as a finite set of basis functions for the global stationary distribution. In the two-dimensional Lozi example, our estimates for the stationary distribution are consistent with those obtained by the direct numerical estimation of the lattice-discretized probability densities computed from long stochastic (Langevin) trajectories.

II. THE NEIGHBORHOOD OF A HYPERBOLIC POINT

An autonomous discrete-time stochastic dynamical system (\mathcal{M}, f, Δ) can be defined by specifying a state space \mathcal{M} , a deterministic map $f: \mathcal{M} \rightarrow \mathcal{M}$, and an additive noise covariance matrix (diffusion tensor) $\Delta = \Delta(x)$. In one time step, an initial Dirac δ density distribution $\rho_a(x)$ located at x is smeared out into a Gaussian ellipsoid $\rho_{a+1}(y)$ centered at $y = f(x)$, with covariance $\Delta(x)$. This defines the kernel of the Fokker-Planck evolution operator in d dimensions [5],

$$\begin{aligned} \mathcal{L}_{FP}(y, x) dx &= e^{-\frac{1}{2}[y-f(x)]^\top \frac{1}{\Delta(x)} [y-f(x)]} [dx], \\ [dx] &= d^d x / \det(2\pi \Delta)^{1/2}. \end{aligned} \quad (2)$$

Consider a trajectory $\{x_a\} = (x_a, x_{a+1}, x_{a+2}, \dots)$ generated by the deterministic evolution rule $x_{a+1} = f(x_a)$, and shift the coordinates in each x_a neighborhood to $x = x_a + z_a$. In the vicinity of x_a , the dynamics can be linearized as $z_{a+1} = M_a z_a$, where $M_a = \partial f(x_a)$ is the one time-step Jacobian matrix.

Prepare the initial density of trajectories $\rho_a(z_a)$ in the x_a neighborhood as a normalized Gaussian distribution $\rho(z_a, Q_a)$,

centered at $z_a = 0$, with a strictly positive-definite covariance matrix Q_a . The support of density $\rho(z_a, Q_a)$ can be visualized as an ellipsoid with axes oriented along the eigenvectors of Q_a . The linearized Fokker-Planck operator,

$$\mathcal{L}(z_{a+1}, z_a) dz_a = e^{-(z_{a+1} - M_a z_a)^\top \frac{1}{2\Delta_a} (z_{a+1} - M_a z_a)} [dz_a],$$

maps this distribution one step forward in time into another Gaussian,

$$\rho(z_{a+1}, Q_{a+1}) = \int dz_a \mathcal{L}(z_{a+1}, z_a) \rho(z_a, Q_a), \quad (3)$$

with the covariance matrix deformed by the dynamics and spread out by the noise, as given by the discrete Lyapunov equation [3,26],

$$Q_{a+1} = M_a Q_a M_a^\top + \Delta_a. \quad (4)$$

In other words, the two covariance matrices, i.e., (i) the deterministically transported $Q_a \rightarrow M_a Q_a M_a^\top$ and (ii) the noise diffusion tensor Δ_a , add together in the usual manner as squares of errors.

Similarly, the density evolution for dynamics with strictly expanding Jacobian matrices M_a can be described by the action of the adjoint Fokker-Planck operator [8,9], with kernel

$$\mathcal{L}^\dagger(y, x) dy = e^{-\frac{1}{2}(y-f(x))^\top \frac{1}{\Delta(x)} (y-f(x))} [dy].$$

The adjoint Fokker-Planck operator expresses the current density ρ_a as the convolution of its image ρ_{a+1} with the noisy dynamics

$$\rho_a(z_a, Q_a) = \int dz_{a+1} \mathcal{L}^\dagger(z_a, z_{a+1}) \rho_{a+1}(z_{a+1}, Q_{a+1}).$$

Like in the forward evolution, we may substitute a Gaussian density into this equation to obtain the discrete adjoint Lyapunov equation for the covariance matrices,

$$M_a Q_a M_a^\top = Q_{a+1} + \Delta_a. \quad (5)$$

We show in what follows that if the Jacobian matrices M_a have all eigenvalues strictly contracting (expanding), any initial Gaussian converges to an invariant density under the action of the (adjoint) Fokker-Planck operator. Consider first the case of a map $f(x)$ with a stable fixed point at x_a (at $z = z_a = 0$). The covariance matrix transforms as

$$\begin{aligned} Q &= \Delta + M \Delta M^\top + M^2 \Delta (M^\top)^2 + \dots \\ &= \sum_{m,n=0}^{\infty} M^n \Delta (M^\top)^m \delta_{mn}. \end{aligned} \quad (6)$$

By inserting the Fourier representation of Kronecker δ_{mn} into (6), we can recast this expression into the resolvent form

$$\begin{aligned} Q &= \int_0^{2\pi} \frac{d\theta}{2\pi} \sum_{m,n=0}^{\infty} (e^{-i\theta} M)^n \Delta (e^{i\theta} M^\top)^m \\ &= \int_0^{2\pi} \frac{d\theta}{2\pi} \frac{1}{1 - e^{-i\theta} M} \Delta \frac{1}{1 - e^{i\theta} M^\top}. \end{aligned} \quad (7)$$

We do the same in the expanding case, by using the adjoint evolution

$$\begin{aligned}
 MQM^\top &= \sum_{m,n=0}^{\infty} \frac{1}{M^n} \Delta \frac{1}{(M^\top)^m} \delta_{nm} \\
 &= \int_0^{2\pi} \frac{d\theta}{2\pi} \sum_{m,n=0}^{\infty} \left(\frac{e^{i\theta}}{M} \right)^n \Delta \left(\frac{e^{-i\theta}}{M^\top} \right)^m \\
 &= \int_0^{2\pi} \frac{d\theta}{2\pi} \frac{1}{1 - e^{i\theta}/M} \Delta \frac{1}{1 - e^{-i\theta}/M^\top}, \quad (8)
 \end{aligned}$$

which is then easily reduced to (7), so that the resolvent form is the same regardless of whether M is expanding or contracting. This result becomes particularly handy when we deal with a hyperbolic fixed point, that is, when M_a has both expanding and contracting eigenvalues. The monodromy matrix is not symmetric and it cannot be diagonalized by an orthogonal transformation, but its expanding and contracting parts can be separated with a similarity transformation S that brings M to a block-diagonal form,

$$\Lambda \equiv S^{-1}MS = \begin{pmatrix} \Lambda_e & 0 \\ 0 & \Lambda_c \end{pmatrix}. \quad (9)$$

Here the blocks Λ_e and Λ_c contain all expanding and contracting eigenvalues of the monodromy matrix, respectively. The covariance matrix $Q = S\hat{Q}S^\top$ is not block diagonalized by the above similarity transformation, but consider the four blocks

$$\hat{Q}_{ij} = \int_0^{2\pi} \frac{d\theta}{2\pi} \frac{1}{1 - e^{-i\theta} \Lambda_i} \hat{\Delta}_{ij} \frac{1}{1 - e^{i\theta} \Lambda_j^\top}, \quad (10)$$

where $\hat{\Delta} \equiv S^{-1}\Delta(S^{-1})^\top$, and $i, j \in \{c, e\}$, where $\{c, e\}$ denotes {contracting, expanding}. This expression may be evaluated as a contour integral around the unit circle in the complex plane [13,15],

$$\hat{Q}_{ij} = \frac{1}{2\pi i} \oint_{\Gamma} dz \frac{1}{z - \Lambda_i} \hat{\Delta}_{ij} \frac{1}{1 - z\Lambda_j^\top}. \quad (11)$$

The diagonal blocks \hat{Q}_{cc} , \hat{Q}_{ee} have either all expanding or all contracting eigenvalues, meaning at least one pole inside and one pole outside the unit circle, and the residue theorem yields a nonvanishing result for the integral. Consider next the off-diagonal block \hat{Q}_{ce} with Λ_i contracting and Λ_j expanding: in this case, the poles all lie outside the unit circle and the integral vanishes. The remaining off-diagonal block having Λ_i expanding and Λ_j contracting must also vanish when integrated, due to the symmetry of \hat{Q} , which is therefore block diagonal,

$$Q = S \begin{pmatrix} \hat{Q}_{ee} & 0 \\ 0 & \hat{Q}_{cc} \end{pmatrix} S^\top. \quad (12)$$

These results are easily extended to a periodic orbit p of period n_p , since any point x_a of the orbit is a fixed point of the n_p th iteration of the map. The forward and adjoint evolution equations (4) and (5) for the covariance matrix, as well as the resolvent (7), all still hold, with some changes in the notation: each periodic point x_a has its own neighborhood, with its own covariance matrix Q_a . The monodromy matrix M_a of x_a now

evolves n_p steps along the orbit,

$$M_a^{n_p} = M_{a+n_p-1} \cdots M_{a+2} M_{a+1} M_a,$$

while the diffusion tensor Δ_a now accounts for the total noise accumulated along the periodic orbit,

$$\Delta_{p,a} \equiv \sum_{i=0}^{n_p-1} M_{a+i+1}^{n_p-i-1} \Delta_{a+i} M_{a+i+1}^{n_p-i-1\top}. \quad (13)$$

III. OPTIMAL PARTITION AND STATIONARY DISTRIBUTION

At this point, our strategy is to build a partition out of *neighborhoods* of the periodic points, each defined by means of the stationarity condition (7): solve for the expanding and contracting blocks of (12) separately and draw a parallelogram on the supports of the resulting Gaussians, with axes oriented along the eigenvectors of the covariance matrices Q_{ee} and Q_{cc} and their widths given by one standard deviation along each direction. We say that two neighborhoods overlap if they do so by at least 50% of their areas (consistent with the 1σ confidence interval chosen as the overlapping threshold in Ref. [8]). For a typical chaotic map, periodic points are dense in the deterministic attractor [27], which we now aim to cover entirely with the minimum number of neighborhoods possible. We do so via the following algorithm: (i) Find periodic points of period $n_p = 1, 2, \dots$, and their corresponding neighborhoods. (ii) If any neighborhood overlaps with the neighborhood of a shorter period, then it is discarded and the neighborhood of lower period occupying the same area is instead kept in the partition. (iii) Among groups of neighborhoods of the same period, discard those that overlap, while keep the rest in the partition. (iv) The algorithm stops when the attractor is fully covered and no further nonoverlapping neighborhoods can be found. An example is shown in Fig. 1 for the two-dimensional Lozi attractor [25].

The main utility of a good partition is that it provides a basis for an accurate and efficient estimate of long-time averages of observables defined on the dynamical system, of the form (1). As explained in Sec. I C, our goal here is to determine the stationary distribution $\rho(x)$. For that purpose, we use as our basis the Gaussian ellipsoids that satisfy the *local* stationarity condition (7) in each neighborhood of the optimal partition. A set of Gaussians centered at every point in the state space \mathcal{M} forms an overcomplete, nonorthogonal basis for functions in $L^2(\mathcal{M})$, as is well known from the study of coherent states of quantum harmonic oscillators [28]. Our (also overcomplete and nonorthogonal) set of Gaussians is centered only on periodic points, which are dense in the deterministic attractor, but not in the entire state space. Therefore, our basis is designed to resolve the structure of any function with support on the hyperbolic “strange set” (an attractor or a repeller). The Gaussians are constructed so that their widths balance the noise spreading and the (time-forward or -backward) contraction of the deterministic dynamics. In the transverse directions, the basis gives the width of the global stationary distribution, locally everywhere determined by the balance between noise and dynamics. Along the attractor, the basis determines the

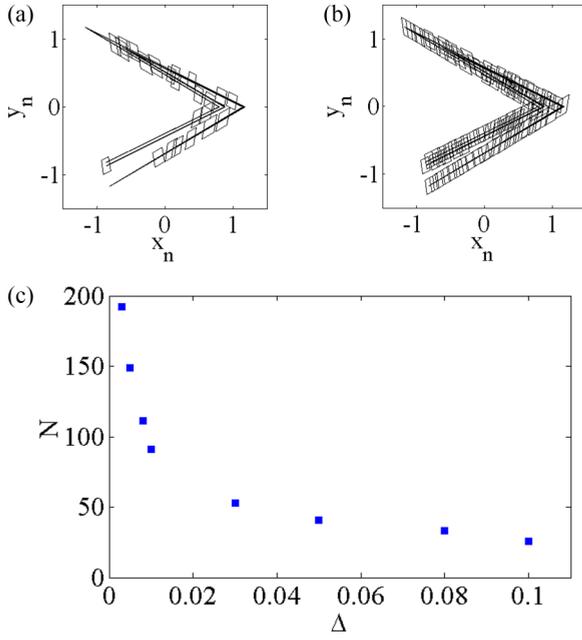


FIG. 1. (Color online) Building the partition for the Lozi attractor, for an isotropic constant diffusion tensor $\delta_{ij}\Delta$. (a),(b) $\Delta = 0.01$. The deterministic attractor is the fractal structure in the background of each picture. Stochastic neighborhoods of a set of periodic points are indicated by their standard deviation parallelograms. (a) An initial partition: with only periodic points of periods ≤ 5 , much of the attractor remains to be covered. (b) The final, optimal partition covers the whole attractor, with no pair of neighborhoods overlapping by more than 50%. (c) N , the number of neighborhoods needed to achieve the optimal partition for a given noise strength Δ .

minimum number of neighborhoods needed to fully resolve the structure of the stationary distribution.

There are numerical methods (such as refinements of Ulam's method [29]) that identify the asymptotic attractor by running long noisy trajectories, dropping the transients, and covering the attractor so revealed by a finite number of boxes. These algorithms have no *a priori* information about how the stationary distribution behaves transversely to the deterministic attractor, and they may easily overestimate the number of basis elements needed to resolve this structure. In contrast, in our approach the transverse structure is automatically accounted for by the local balance between the noise and the deterministic contraction along the stable, transverse directions, given by covariance matrix block \hat{Q}_{cc} in Eq. (12). Furthermore, estimating $\rho(x)$ by binning a long noisy trajectory over a finite number of attractor-covering boxes is feasible only in a low-dimensional state space, while (12) can be computed for a state space of any dimension.

In discrete time dynamics, the stationary distribution is the ground-state eigenfunction of the Fokker-Planck evolution operator (2) with escape rate γ ,

$$\mathcal{L}\rho(x) = e^{-\gamma}\rho(x). \quad (14)$$

In order to estimate the stationary distribution, we write it as a sum over the neighborhoods of the periodic points:

$$\rho(x) = \sum_{a=1}^N h_a \phi_a(x), \quad (15)$$

where $\phi_a = e^{-x_a^\top Q_a x_a}$ are the Gaussian basis functions, with Q_a given by (12), and the coefficients $\{h_a\}$ to be determined. The truncation of the expansion (15) to N basis functions follows from our optimal partition. We estimate the coefficients h_a by minimizing the cost function,

$$\int \left[\sum_{a=1}^N h_a (\mathcal{L} - e^{-\gamma}) \phi_a(x) \right]^2 dx, \quad (16)$$

together with the normalization constraint for $\rho(x)$. We can also estimate the escape rate of the system by minimizing the error with respect to $e^{-\gamma}$.

As an example, we apply the procedure to the Lozi map [25],

$$\begin{aligned} x_{n+1} &= 1 - a|x_n| + b y_n, \\ y_{n+1} &= x_n, \end{aligned} \quad (17)$$

with parameters $a = 1.85, b = 0.3$ and isotropic, constant diffusion tensor $\delta_{ij}\Delta$, with Δ ranging in the interval $[0.003, 0.1]$. Figure 1(c), which shows the number N of neighborhoods required by the optimal partition for a given Δ , illustrates the efficiency of our method: N goes from tens to few hundreds in the noise range considered. In order to test our algorithm, we also estimate $\rho(x)$ and γ by a direct numerical simulation. The (x, y) state space is divided into uniform mesh 6.4×10^5 bins; we follow long stochastic trajectories and count how many times they visit each bin. The stationary distribution $\rho_B(x)$ is then the normalized frequency distribution of the whole grid. The deterministic Lozi map has a fixed point at the edge of the attractor, whose stable manifold is the boundary of the deterministic basin of attraction. The noise makes it possible for a stochastic trajectory to cross this boundary and escape. We compute the escape rate as the ratio of the total number of points in the noisy trajectories to the number of escapes. Figure 2 shows an example of the stationary distribution estimated with both methods, while in Fig. 3 we quantitatively compare the two procedures. In particular, we estimate the relative error between the stationary distribution ρ computed with the optimal partition and ρ_B computed on the uniform grid, by using a normalized L^2 distance, as

$$d(\rho, \rho_B) = \frac{\int [\rho(x) - \rho_B(x)]^2 dx}{\int [\rho_B(x)]^2 dx}. \quad (18)$$

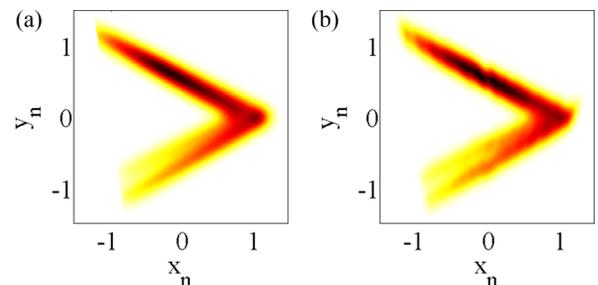


FIG. 2. (Color online) The stationary distribution of the Lozi map with $\Delta = 0.01$. (a) A direct numerical calculation obtained by running noisy trajectories for a long time. (b) The stationary distribution calculated with the optimal partition method.

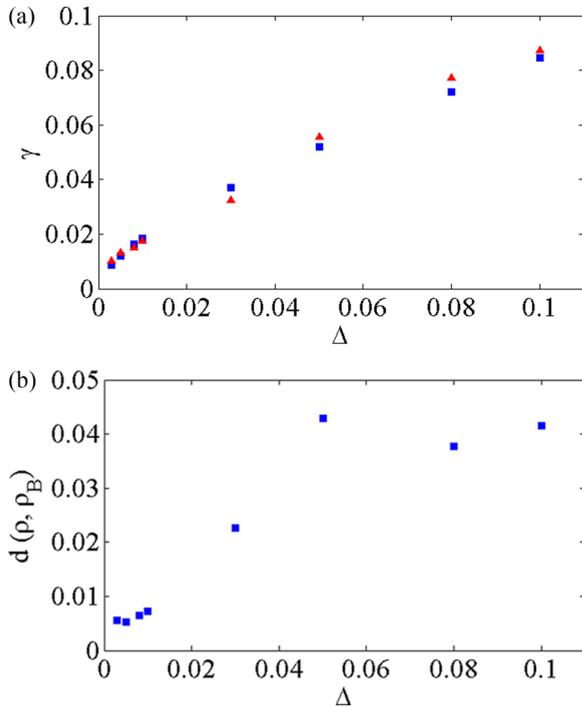


FIG. 3. (Color online) (a) The escape rate from the attractor as the function of the noise strength Δ . Squares: uniform grid discretization method. Triangles: optimal partition. (b) The normalized L^2 distance $d(\rho, \rho_B)$ between the corresponding stationary distributions.

The two distributions are within 5% of each other, whereas the escape rates differ by at most 10%, over a range of Δ that spans two orders of magnitude.

IV. DISCUSSION

In conclusion, we have generalized the optimal partition hypothesis first formulated in [8] to hyperbolic maps in arbitrary dimension and tested the method on a two-dimensional system with weak white isotropic noise. As noise induces a finite resolution of the state space of any physical system, finite numbers of neighborhoods suffice to partition the state space explored by chaotic dynamics and to estimate long-time averages of observables. Here we have used the deterministic unstable periodic orbits as the skeleton on which to build an optimal partition for the noisy state space. First we determine a local stationary distribution in the neighborhood of each periodic point by balancing the noise against the deterministic expansion or contraction. From the separation of expanding and contracting blocks in the covariance matrix

that characterizes the Gaussian approximation to the local stationary distribution, we carve out a precise definition of neighborhood, i.e., the constituent of our partition, which is then used to approximate the global stationary distribution, estimate the escape rate (for open systems that allow escape), and estimate any long-time averaged observable. Numerical tests confirm that the accuracy of our method is comparable to that of a uniform grid discretization, but the number of neighborhoods required for our optimal partition (~ 10 to 100) is three to four orders of magnitude smaller than the number of bins used in the uniform grid discretization method ($\sim 10^5$).

The problems that dynamical chaos (or “turbulence”) theory faces nowadays are not two dimensional, but high and even infinite dimensional. Today it is possible to compute numerically exact periodic orbits (“recurrent flows” [30]) in a variety of physically realistic turbulent fluid flows [31,32], but these calculations are at the limit of what current codes can do, and we hope that the methods presented here can provide sharp criteria for when a sufficient number of such solutions has been computed. Furthermore, unlike the uniform grid discretization, our partitions are smart, since they rely on the periodic orbits of the deterministic system as skeleton of the dynamics, as well as efficient, due to the finite (and surprisingly optimal) numbers of neighborhoods and corresponding basis functions. This, we believe, should make our algorithm less costly to implement than direct numerical simulations in higher dimensions, where discretizations would be impractical. With some modifications and application of Poincaré sections, the formalism can be applied to continuous time flows as well [9,33]. Outstanding challenges include dealing with the lack of hyperbolicity in higher dimensions (marginal directions were treated in Ref. [9] for 1d maps) as well as extending the definition of neighborhood to other time-invariant sets, such as relative periodic orbits and partially hyperbolic invariant manifolds. Further technical issues, such as improving the efficiency of the minimization algorithm by modifying the basis of functions used in the computation of the stationary distribution, are also part of our agenda.

ACKNOWLEDGMENTS

We are indebted to P. M. Svetlichnyy for suggesting that we use error minimization to find the global stationary distribution. J.M.H. is thankful for the partial support provided by a Presidential Undergraduate Research Award, and to D.L. for the hospitality at IASTU, Beijing. P.C. thanks the family of the late G. Robinson, Jr. and NSF Grant No. DMS-1211827 for partial support. D.L. acknowledges support from the National Science Foundation of China (NSFC), International Young Scientists (Grant No. 11450110057-041323001).

- [1] P. S. Laplace, Mémoire sur les intégrales définies et leur application aux probabilités, et spécialement à la recherche du milieu qu’il faut choisir entre les résultats des observations, Mem. Acad. Sci. (I), XI, Sec. V, 375 (1810).
 [2] M. Jacobsen, Laplace and the origin of the Ornstein-Uhlenbeck process, *Bernoulli* **2**, 271 (1996).

- [3] A. M. Lyapunov, The general problem of the stability of motion, *Int. J. Control* **55**, 531 (1992).
 [4] G. E. Uhlenbeck and L. S. Ornstein, On the theory of the Brownian motion, *Phys. Rev.* **36**, 823 (1930).
 [5] H. Risken, *The Fokker-Planck Equation* (Springer, New York, 1996).

- [6] N. G. van Kampen, *Stochastic Processes in Physics and Chemistry* (North-Holland, Amsterdam, 1992).
- [7] H. Nakanishi, T. Sakaue, and J. Wakou, Hamilton-Jacobi method for molecular distribution function in a chemical oscillator, *J. Chem. Phys.* **139**, 214105 (2013).
- [8] D. Lippolis and P. Cvitanović, How Well can One Resolve the State Space of a Chaotic Map? *Phys. Rev. Lett.* **104**, 014101 (2010).
- [9] P. Cvitanović and D. Lippolis, Knowing when to stop: How noise frees us from determinism, in *Let's Face Chaos through Nonlinear Dynamics*, edited by M. Robnik and V. G. Romanovski (American Institute of Physics, Melville, NY, 2012), pp. 82–126.
- [10] M. C. Gutzwiller, Periodic orbits and classical quantization conditions, *J. Math. Phys.* **12**, 343 (1971).
- [11] D. Ruelle, Zeta functions for expanding maps and Anosov flows, *Invent. Math.* **34**, 231 (1976).
- [12] P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, and G. Vattay, *Chaos: Classical and Quantum* (Niels Bohr Institute, Copenhagen, 2015).
- [13] K. Zhou, G. Salomon, and E. Wu, Balanced realization and model reduction for unstable systems, *Intl. J. Robust Nonlin. Control* **9**, 183 (1999).
- [14] B. F. Farrell and P. J. Ioannou, Accurate low-dimensional approximation of the linear dynamics of fluid flow, *J. Atmos. Sci.* **58**, 2771 (2001).
- [15] A. Varga, On balancing and order reduction of unstable periodic systems, in *Periodic Control Systems 2001*, edited by S. Bittanti and P. Colaneri (Elsevier, New York, 2001), pp. 177–182.
- [16] J. P. Crutchfield and N. H. Packard, Symbolic dynamics of noisy chaos, *Physica D* **7**, 201 (1983).
- [17] C. S. Daw, C. E. A. Finney, and E. R. Tracy, A review of symbolic analysis of experimental data, *Rev. Sci. Instrum.* **74**, 915 (2003).
- [18] M. Buhl and M. B. Kennel, Statistically relaxing to generating partitions for observed time-series data, *Phys. Rev. E* **71**, 046213 (2005).
- [19] S. M. Ulam, *A Collection of Mathematical Problems* (Interscience, New York, 1960).
- [20] G. Froyland, Extracting dynamical behaviour via Markov models, in *Nonlinear Dynamics and Statistics: Proceedings of the Newton Institute, Cambridge 1998*, edited by A. Mees (Birkhäuser, Boston, 2001), pp. 281–321.
- [21] D. J. Chappell, G. Tanner, D. Löchel, and N. Søndergaard, Discrete flow mapping: Transport of phase space densities on triangulated surfaces, *Proc. R. Soc. A* **469**, 20130153 (2013).
- [22] Y. G. Sinai, Gibbs measures in ergodic theory, *Russian Math. Surveys* **27**, 21 (1972).
- [23] R. Bowen, *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms* (Springer, Berlin, 1975).
- [24] D. Ruelle, *Statistical Mechanics, Thermodynamic Formalism* (Addison-Wesley, Reading, MA, 1978).
- [25] R. Lozi, Un attracteur étrange (?) du type attracteur de Hénon, *J. Phys. (Paris), Colloq.* **39**, C5 (1978).
- [26] Z. Gajić and M. Qureshi, *Lyapunov Matrix Equation in System Stability and Control* (Academic, New York, 1995).
- [27] R. L. Devaney, *An Introduction to Chaotic Dynamical Systems* (Addison-Wesley, Redwood City, 1989).
- [28] K. Gottfried and T. Yan, *Quantum Mechanics: Fundamentals*, 2nd ed. (Springer, New York, 2003), pp. 181–184.
- [29] L. Ermann and D. Shepelyansky, Ulam method and fractal Weyl law for Perron-Frobenius operators, *Eur. Phys. J. B* **75**, 299 (2010).
- [30] P. Cvitanović, Recurrent flows: The clockwork behind turbulence, *J. Fluid Mech. Focus Fluids* **726**, 1 (2013).
- [31] M. Nagata, Three-dimensional finite-amplitude solutions in plane Couette flow: Bifurcation from infinity, *J. Fluid Mech.* **217**, 519 (1990).
- [32] J. F. Gibson, J. Halcrow, and P. Cvitanović, Visualizing the geometry of state-space in plane Couette flow, *J. Fluid Mech.* **611**, 107 (2008).
- [33] P. Gaspard, Trace formula for noisy flows, *J. Stat. Phys.* **106**, 57 (2002).