

# Nonparametric forecasting of low-dimensional dynamical systems

Tyrus Berry,<sup>1</sup> Dimitrios Giannakis,<sup>2</sup> and John Harlim<sup>1,3</sup>

<sup>1</sup>*Department of Mathematics, the Pennsylvania State University, University Park, PA 16802-6400, USA*

<sup>2</sup>*Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA*

<sup>3</sup>*Department of Meteorology, the Pennsylvania State University, University Park, PA 16802-5013, USA*

(Dated: January 13, 2015)

This letter presents a non-parametric modeling approach for forecasting stochastic dynamical systems on low-dimensional manifolds. The key idea is to represent the discrete shift maps on a smooth basis which can be obtained by the diffusion maps algorithm. In the limit of large data, this approach converges to a Galerkin projection of the semigroup solution to the underlying dynamics on a basis adapted to the invariant measure. This approach allows one to quantify uncertainties (in fact, evolve the probability distribution) for non-trivial dynamical systems with equation-free modeling. We verify our approach on various examples, ranging from an inhomogeneous anisotropic stochastic differential equation on a torus, the chaotic Lorenz three-dimensional model, and the Niño-3.4 data set which is used as a proxy of the El-Niño Southern Oscillation.

A significant challenge in modeling is to account for physical processes which are often not well understood but for which large data sets are available. A standard approach is to perform regression fitting of the data into various parametric models. While this approach is popular and successful in many applied domains, the resulting predictive skill can be sensitive to the choice of models, the parameter fitting algorithm, and the complexity of the underlying physical processes. An alternative approach is to avoid choosing particular models and/or parameter estimation algorithms and instead apply a non-parametric modeling technique. In particular, nonparametric modeling based on local linearization of discrete shift maps on paths has been successful in predicting mean statistics of data sets generated by dynamical systems with low-dimensional attractors [1–3].

In this letter, we generalize this nonparametric approach to quantify the evolving probability distribution of the underlying dynamical system. The key idea is to project the shift map on a set of basis functions that are generated by the diffusion maps algorithm [4] with a variable bandwidth diffusion kernel [5]. This approach has connections with a recently developed family of kernels [6], which utilize small shifts of a deterministic time series to estimate the dynamical vector field. The method also generalizes a recently introduced non-parametric modeling framework for gradient systems [7] to inhomogeneous stochastic systems having non-gradient drift and anisotropic diffusion. Consider a dynamical system,

$$dx = a(x) dt + b(x) dW_t \quad (1)$$

where  $W_t$  is a standard Brownian process,  $a(x)$  a vector field, and  $b(x)$  a diffusion tensor, all defined on a manifold  $\mathcal{M} \subset \mathbb{R}^n$ . Given a time series  $x_i = x(t_i)$ , generated by (1) at discrete times  $\{t_i\}_{i=1}^{N+1}$ , we are interested in constructing a forecast model so that given an initial density  $p_0(x)$  we can estimate the density  $p(x, t) = e^{t\mathcal{L}^*} p_0(x)$  at time  $t > 0$ , without the Fokker-Planck operator,  $\mathcal{L}^*$ , of (1) and without knowing or estimating  $a$  and  $b$ . We will

assume that (1) is ergodic so that  $\{x_i\}$  are sampled from the invariant measure  $p_{\text{eq}}(x)$  of (1). Note that all probability densities are relative to a volume form  $dV$  which  $\mathcal{M}$  inherits from the ambient space, and the generator  $\mathcal{L}$  of (1) is the adjoint of  $\mathcal{L}^*$  with respect to  $L^2(\mathcal{M}, dV)$ .

We note that our approach differs significantly from previous approaches such as [8, 9], which estimate  $a$  and  $b$  explicitly in the ambient space  $\mathbb{R}^n$ , relying on the Kramer-Moyal expansion. In contrast, our approach directly estimates the semi-group solution  $e^{\tau\mathcal{L}}$  on the manifold  $\mathcal{M} \subset \mathbb{R}^n$ , so that  $a$  and  $b$  are represented implicitly. The advantages of our approach are that the data requirements are independent of the ambient space dimension and only depend on the intrinsic dimension of  $\mathcal{M}$ , and we will be able to estimate the semi-group solution,  $e^{\tau\mathcal{L}}$ , directly from the data for any sampling time  $\tau$ .

Our approach is motivated by a rigorous connection between the shift map,  $S$  which we define by  $Sf(x_i) = f(x_{i+1})$  for a function  $f \in L^2(\mathcal{M}, p_{\text{eq}})$ , and the semi-group solution,  $e^{\tau\mathcal{L}}$ , of the underlying dynamical system (1). Applying Itô's Lemma to  $f(x)$ , one can show,

$$Sf(x_i) = e^{\tau\mathcal{L}} f(x_i) + \int_{t_i}^{t_{i+1}} \nabla f^\top b dW_s + \int_{t_i}^{t_{i+1}} Bf ds, \quad (2)$$

where  $\tau = t_{i+1} - t_i$ ,  $Bf = \mathcal{L}f - \mathbb{E}[\mathcal{L}f]$  and the expectation is with respect to paths of (1) conditional to  $x(t_i)$ . The detailed derivation of (2) is in Appendix B in the supplementary material. Since  $\mathbb{E}[Sf(x_i)] = e^{\tau\mathcal{L}} f(x_i)$ , we can use the shift map,  $S$ , to directly estimate the semi-group solution,  $e^{\tau\mathcal{L}}$ . However,  $S$  is a noisy estimate of  $e^{\tau\mathcal{L}}$  and our key contribution is to minimize the error by representing  $S$  on a basis of smooth functions.

Minimizing the error requires choosing a basis which minimizes the functional  $\|\nabla f\|_{p_{\text{eq}}}$ , which is shown in Appendix B. Intuitively, this is because we want to bound the stochastic integral in (2), whose integrand contains  $\nabla f$ . The functional  $\|\nabla f\|_{p_{\text{eq}}}$  is minimized by the eigenfunctions of the generator,  $\hat{\mathcal{L}}$ , of a stochastically forced

gradient flow with potential  $U(x) = -\log(p_{\text{eq}}(x))$ . Let  $\lambda_j$  and  $\varphi_j$  be the eigenvalues and eigenfunctions of  $\hat{\mathcal{L}}$ ;  $\{\varphi_j\}$  are orthonormal on  $L^2(\mathcal{M}, p_{\text{eq}})$ . Since  $\hat{\mathcal{L}}$  is the generator of a gradient flow systems, it is easy to check that  $\psi_j = p_{\text{eq}}\varphi_j$  are the eigenfunctions of the adjoint operator  $\hat{\mathcal{L}}^*$ , which are orthonormal on  $L^2(\mathcal{M}, p_{\text{eq}}^{-1})$ . Numerically, we obtain  $\varphi_j(x_i)$  as eigenvectors of a stochastic matrix, constructed by evaluating a variable bandwidth kernel on all pairs of data points and then applying an appropriate normalization [5]. We summarize this procedure and the related results in Appendix A in the Supplementary material. We emphasize that the operator  $\hat{\mathcal{L}}$  is used only to estimate  $\{\varphi_j\}$  which is the optimal basis for smoothing the shift operator  $S$  approximating semi-group solution  $e^{\tau\mathcal{L}}$  of the full system (1).

We write the solution,  $p(x, \tau) = e^{\tau\mathcal{L}^*} p_0(x)$ , as follows,

$$p(x, \tau) = \sum_{l=1}^{\infty} \langle e^{\tau\mathcal{L}^*} p_0, \psi_l \rangle_{p_{\text{eq}}^{-1}} \psi_l(x) = \sum_{l=1}^{\infty} \langle p_0, e^{\tau\mathcal{L}} \varphi_l \rangle \psi_l(x).$$

Similarly, the initial density is  $p_0(x) = \sum_j c_j(0) \psi_j(x)$ , where  $c_j(0) = \langle p_0, \psi_j \rangle_{p_{\text{eq}}^{-1}}$ . We therefore obtain,

$$p(x, \tau) = \sum_{l=1}^{\infty} \sum_{j=1}^{\infty} c_j(0) A_{lj}(\tau) p_{\text{eq}}(x) \varphi_l(x), \quad (3)$$

where  $A_{lj}(\tau) := \langle \varphi_j, e^{\tau\mathcal{L}} \varphi_l \rangle_{p_{\text{eq}}}$ . Based on the discussion after (2), we will use  $\langle \varphi_j, S\varphi_l \rangle_{p_{\text{eq}}}$  to estimate  $A_{lj}$ . Numerically, we estimate  $A_{lj}$  by a Monte-Carlo integral,

$$\hat{A}_{lj} = \frac{1}{N} \sum_{i=1}^N \varphi_j(x_i) \varphi_l(x_{i+1}), \quad (4)$$

such that, the *diffusion forecast* is defined as follows:

$$p(x, \tau) \approx \sum_{l=1}^{\infty} p_{\text{eq}}(x) \varphi_l(x) \sum_{j=1}^{\infty} \hat{A}_{lj}(\tau) c_j(0). \quad (5)$$

One can show that  $\mathbb{E}[\hat{A}_{lj}] = A_{lj}$  which means that  $\hat{A}_{lj}$  is an unbiased estimate of  $A_{lj}$ . Moreover, the error of this estimate is of order  $\lambda_l \sqrt{\tau/N}$  in probability assuming that  $x_i$  are independent samples of  $p_{\text{eq}}$ . This shows that we can apply a diffusion forecast for any sampling time  $\tau$  given a sufficiently large data set  $N$ . For more details, see Appendix B in the Supplementary Material.

*Non-gradient drift anisotropic diffusion:* We first verify the above approach for a system of SDE's of the form (1) on a torus defined in the intrinsic coordinates  $(\theta, \phi) \in [0, 2\pi)^2$  with drift and diffusion coefficients,

$$a(\theta, \phi) = \left( \frac{1}{2} + \frac{1}{8} \cos(\theta) \cos(2\phi) + \frac{1}{2} \cos(\theta + \pi/2) \right),$$

$$b(\theta, \phi) = \begin{pmatrix} \frac{1}{4} + \frac{1}{4} \sin(\theta) & \frac{1}{4} \cos(\theta + \phi) \\ \frac{1}{4} \cos(\theta + \phi) & \frac{1}{40} + \frac{1}{40} \sin(\phi) \cos(\theta) \end{pmatrix}.$$

This example is chosen to exhibit non-gradient drift, anisotropic diffusion, and multiple time scales. Since it is a system of the form (1) on a smooth manifold, our theory shows that the shift operator  $S$  is an unbiased estimator for the semigroup solution  $e^{\tau\mathcal{L}}$ , and in the limit of large data the diffusion forecast will capture all aspects of the evolution of the density  $p(x, t)$ .

We now verify this theory using a training data set of 20000 points generated by numerically solving the SDE in (1) with a discrete time step  $\Delta t = 0.1$  and then mapping this data into the ambient space,  $\mathbb{R}^3$ , via the standard embedding of the torus given by  $(x, y, z) = ((2 + \sin(\theta)) \cos(\phi), (2 + \sin(\theta)) \sin(\phi), \cos(\theta))$ . We define a Gaussian initial density  $p_0(\theta, \phi)$  with a randomly selected mean and a diagonal covariance matrix with variance 1/10. The initial density is projected into a basis of  $M = 1000$  eigenfunctions giving coefficients  $c_j(0) = \langle p_0/p_{\text{eq}}, \varphi_j \rangle_{p_{\text{eq}}} \approx \sum_{i=1}^{20000} p_0(\theta_i, \phi_i) \varphi_j(\theta_i, \phi_i) / p_{\text{eq}}(\theta_i, \phi_i)$ . The coefficients  $c_j(0)$  are evolved forward in time in discrete steps of length  $\Delta t = 0.1$  by  $\hat{A}$ , constructed by (4) so that the forecast at time  $\tau = n\Delta t$  is effectively  $\hat{A}(\Delta t)^n$ .

In Figure 1, we show the evolution of the first two moments in the ambient space, for the fast and slow variables,  $x$  and  $z$ , respectively, created by the diffusion forecast in (5). To verify the accuracy of the diffusion forecast, we also show the corresponding moments produced by an ensemble forecast of 50000 initial conditions, randomly sampled from the initial distribution  $p_0(\theta, \phi)$ , evolved using the true dynamical system. Notice the long-time pathwise agreement of both moments constructed via the diffusion forecast and those constructed by an ensemble forecast. See also a video of the evolution of  $p$  in the Supplementary Material.




FIG. 1. Validation of first two moments of the diffusion forecast for a stochastic dynamical system on a torus in  $\mathbb{R}^3$ .

*Lorenz-63 model:* Next, we apply the diffusion forecasting algorithm to data generated by the Lorenz-63 model [10], with error in the initial state. Unlike the

previous example, this model does not technically satisfy the requirements of our theory because the attractor is a fractal set rather than a smooth manifold. Our results indicate that the applicability of the method seems to extend beyond the current theory. We compare our approach to the classical nonparametric prediction method which uses a local linear approximation of the shift map [1–3] and a standard ensemble forecasting method with the true model, applied with 50000 initial conditions, sampled from the same initial distribution  $p_0$ .

We generate 10000 data points with discrete time steps  $\Delta t = 0.1$  and  $\Delta t = 0.5$  by solving the Lorenz-63 model. We use the first 5000 data points as training data for the nonparametric models, and the remaining 5000 data points to verify the forecasting skill. For each of the 5000 verification points,  $x_t$ , we define the initial state  $\hat{x}_t = x_t + \xi_t$  by introducing random perturbations  $\xi_t$  sampled from  $\mathcal{N}(0, 0.01)$ . Each forecast method starts with the same initial density,  $p_0 = \mathcal{N}(\hat{x}_t, 0.01)$  centered at the perturbed verification point. We chose this very small perturbation to demonstrate the diffusion forecast for an initial condition which is almost perfect; as the amount of noise increases the advantage of the diffusion forecast over the linear methods is even more significant.

The diffusion forecast is performed with 4500 eigenfunctions  $\varphi_j$ , constructed with the diffusion maps algorithm with a variable bandwidth [5] (we show examples in Figure 2). The local linear forecast uses ordinary least squares to fit an affine model to the  $n$ -step shift map on the 15 nearest neighbors to the initial state. The iterated local linear forecast completes this process for one step and then recomputes the 15 nearest neighbors to the 1-step forecast and then repeats the process. The variance estimate of the local linear models is given by conjugating the covariance matrix of  $p_0$  with the linear part of the appropriate affine forecast model. We compute the root mean squared error (RMSE) between each mean forecast and the true state, averaged over the verification period of 5000 steps. We also show the standard deviation of the forecast density, so that a forecasting method has good uncertainty quantification (UQ) if the standard deviation agrees with the RMSE.

Of course, the ensemble forecast with the true model gives the best forecast, however the diffusion forecast is a considerable improvement over the local linear forecast. For short  $\Delta t = 0.1$ , the iterated local linear forecast is comparable to the diffusion forecast except in the long term where the iterated local linear forecast exhibits significant bias. Moreover, the iterated local linear forecast significantly overestimates the error variance in the intermediate to long term forecast. This overestimation is due to the positive Lyapunov exponent, which is implicitly estimated by the product of the iterated local linearizations. In contrast, the direct local linearization is unbiased in the long term, but converges very quickly to the invariant measure and underestimates the variance.

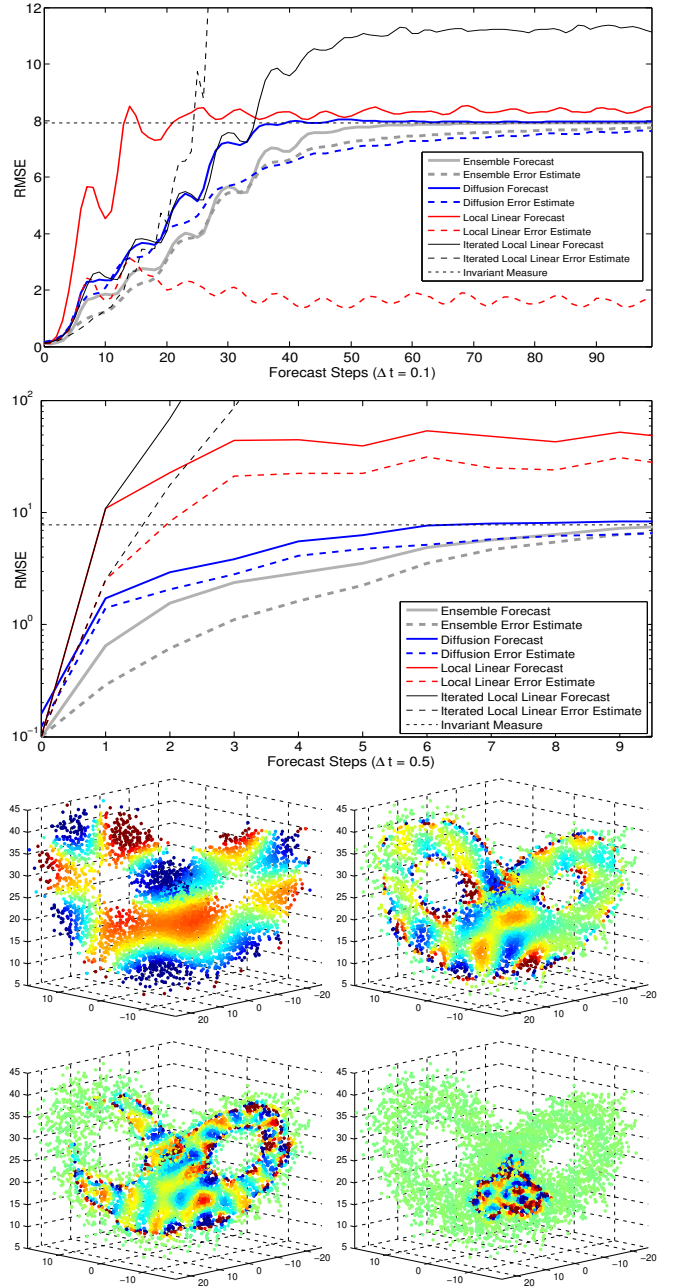


FIG. 2. Comparing forecast methods for Lorenz-63, Top:  $\Delta t = 0.1$ , Middle:  $\Delta t = 0.5$ . Bottom: Eigenfunctions  $\varphi_{40}$ ,  $\varphi_{500}$ ,  $\varphi_{1500}$ , and  $\varphi_{4000}$  of the coarse approximation of the fractal attractor by a manifold.

This underestimation is because no single linearization can capture the information creation introduced by the positive Lyapunov exponent. For long  $\Delta t = 0.5$ , the bias in the local linear models leads them to diverge far beyond the invariant measure for even intermediate term forecasts. The ensemble forecast provides the most consistent UQ since it has access to the true model, however the diffusion forecast produces reasonable estimates without knowing the true model as shown in Figure 2. The local linear forecast error estimates vary widely and




FIG. 3. Forecasting for the El Niño 3.4 index. RMS and correlation (top); 14-month lead forecast (bottom)

do not robustly provide a useful UQ whereas the diffusion forecast is robust across multiple sampling times as suggested by the theory. We include a video showing good long term agreement between the diffusion forecast density and an ensemble in the Supplementary Material.

The diffusion forecast is able to give a reasonable estimate of the evolution of the density by building a consistent finite dimensional Markovian approximation of the dynamics. This Markovian system incorporates global knowledge of the attractor structure via the smoothing with the adapted basis  $\{\varphi_j\}$ . This Markovian approximation of the Lorenz-63 model implicitly uses a small Brownian forcing to replicate the entropy generation of the positive Lyapunov exponent.

*El-Niño data set:* We now apply our method to a real world data set, where the validity of our theory is unverifiable, namely the Niño-3.4 index, which records the monthly anomalies of sea surface temperature (SST) in the central equatorial Pacific region (the raw dataset is available from NOAA <ftp://ftp.cpc.ncep.noaa.gov/wd52dg/data/indices/>). In applying this method, we implicitly assume that there is an underlying dynamics on a low-dimensional manifold that generates these SST anomalies. Since the time series is one-dimensional, we apply the time-delay embedding technique to the data set, following [11–13], which will recover this low dimensional manifold if it exists. We use a 5-lag embedding, empirically chosen to maximize the forecast correlation skill between the true time series and the mean estimate.

We construct the  $A_{ij}$  matrix with 80 eigenfunctions obtained by the diffusion maps algorithm with a variable bandwidth kernel, applied on the lag-embedded data set.

In this experiment, we train our nonparametric model with monthly data between Jan 1950-Dec 1999 (600 data points), following [14] and we verify the forecasting skill on the period of Jan 2000-Sept 2013. The initial distribution  $p_0$  is generated with the same method as in the Lorenz-63 example. Based on the RMSE and correlation measure (see Figure 3), the forecasting skill decays to the climatological error in about 6 months but then the skill improves, peaking at 13-14 month lead time. In fact, our 14-month lead forecast skill, in terms of RMSE 0.60 and correlation 0.64, is significantly better than that of the method proposed in [14] (Fig. 3 in their paper suggests RMSE 1.4 and correlation 0.4) who claimed to beat the current operational forecasting skill. The 14-month lead forecast mean estimate gives a reasonably correct pattern, and the diffusion forecast provides a reasonable error bar (showing one stdev.) which is a useful UQ. We include a movie in the Supplementary Material showing the nontrivial evolution of the forecast distribution starting 14 months before January 2004.

Difficulty in improving the forecasts in this problem may be due to combinations of many factors, including the validity of our assumption of the existence of low-dimensional structures, a transient in the time series, a large stochastic component, and memory effects. One possibility is to combine the local linear models via the diffusion basis to form a global model which respects to invariant measure. Another issue is that both the observational and dynamical noise in the data is currently treated as part of the process, and it would be advantageous to isolate the attractor and build the basis there. Finally, the empirical success of this method suggests that it is possible to approximate a fractal attractor with a smooth manifold, however, there is limited theoretical interpretation for such an approximation.

This research is supported under the ONR MURI grant N00014-12-1-0912. D.G. also acknowledges financial support from ONR DRI grant N00014-14-1-0150. J.H. also acknowledges financial support from ONR grant N00014-13-1-0797 and the NSF grant DMS-1317919.

- 
- [1] J. Farmer and J. Sidorowich, Phys. Rev. Lett. **59**, 845 (1987).
  - [2] T. Sauer, in *Time Series Prediction, Forecasting the Future and Understanding the Past* (1994) pp. 175–193.
  - [3] D. Kugiumtzis, O. C. Lingjrd, and N. Christophersen, Physica D: Nonlinear Phenomena **112**, 344 (1998).
  - [4] R. Coifman and S. Lafon, Appl. Comput. Harmon. Anal. **21**, 5 (2006).
  - [5] T. Berry and J. Harlim, Appl. Comput. Harmon. Anal. [to appear] (2014), <http://arxiv.org/abs/1406.5064>.

- [6] D. Giannakis, SIAM J. Appl. Dyn. Sys. [to appear] (2014), <http://arxiv.org/abs/1403.0361>.
- [7] T. Berry and J. Harlim, SIAM/ASA J. Uncer. Quant. [in review] (2014), <http://arxiv.org/abs/1407.6972>.
- [8] R. Friedrich and J. Peinke, Phys. Rev. Lett. **78**, 863 (1997).
- [9] F. Böttcher, J. Peinke, D. Kleinhans, R. Friedrich, P. G. Lind, and M. Haase, Phys. Rev. Lett. **97**, 090603 (2006).
- [10] E. N. Lorenz, J. Atmos. Sci. **20**, 130 (1963).
- [11] T. Sauer, J. Yorke, and M. Casdagli, J. Stat. Phys. **65**, 579 (1991).
- [12] D. Giannakis and A. J. Majda, Proc. Nat. Acad. Sci. **109**, 2222 (2012).
- [13] T. Berry, J. R. Cressman, Z. G. Ferencsek, and T. Sauer, SIAM J. Appl. Dyn. Syst. **12**, 618 (2013).
- [14] M. D. Chekroun, D. Kondrashov, and M. Ghil, Proc. Nat. Acad. Sci. **108**, 11766 (2011).
- [15] R. R. Coifman, Y. Shkolnisky, F. J. Sigworth, and A. Singer, Image Processing, IEEE Transactions on **17**, 1891 (2008).

## APPENDIX A: ESTIMATING ELLIPTIC OPERATORS WITH DIFFUSION KERNELS

In this appendix we include the details of the numerical algorithm introduced in [5] which is used to estimate the eigenfunctions  $\varphi_j$  of the elliptic operator  $\hat{\mathcal{L}} = \Delta - c_1 \nabla U \cdot \nabla$ . We assume that the data set  $\{x_i\}_{i=1}^N \subset \mathbb{R}^n$  lies on a  $d$ -dimensional submanifold  $\mathcal{M}$  of the ambient Euclidean space  $\mathbb{R}^n$ . In order to describe functions on the manifold, we will represent a function  $f$  by evaluation on the data set, so that  $f$  is represented by the discrete  $N \times 1$  vector  $\vec{f} = (f_1, \dots, f_N)^\top$  where  $f_i = f(x_i)$ . In this framework, operators which map functions to functions are represented as  $N \times N$  matrices, since these take the  $N \times 1$  functions to  $N \times 1$  functions. In this sense, we will approximate the eigenfunctions of  $\hat{\mathcal{L}}$  as eigenvectors of a kernel matrix which provably approximates the operator  $\hat{\mathcal{L}}$  in the limit of large data. In this paper we are interested in data sets generated by dynamical systems, however the general theory of this section is valid for arbitrary data sets. In this case, we will assume that the dynamics are ergodic so that the manifold  $\mathcal{M}$  is an attractor of the system, and the sampling measure  $q$  of the data on  $\mathcal{M}$  is the same as the invariant measure  $p_{\text{eq}}$  of the dynamics.

We first note that  $\hat{\mathcal{L}}$  is the generator of the gradient flow system,

$$dx = -c_1 \nabla U(x) dt + \sqrt{2} dW_t, \quad (6)$$

where  $W_t$  is a Brownian motion on the manifold  $\mathcal{M}$  (so that the Laplacian  $\Delta$  is the infinitesimal generator of  $W_t$ ). The potential function  $U = -\log(q)$  is determined by the sampling measure,  $q$ . In particular, we will be interested in the case  $c_1 = 1$ , in which case the invariant measure of the system (6) is  $e^{-c_1 U} = q = p_{\text{eq}}$ . This means that the invariant measure of the true dynamical system which

governs the evolution of the data set is the same as that of the gradient flow system generated by  $\hat{\mathcal{L}}$ . Since  $\hat{\mathcal{L}}$  is a negative semi-definite elliptic operator, self-adjoint with respect to  $p_{\text{eq}}$ , the eigenfunctions  $\{\varphi_j\}$  form a basis for  $L^2(\mathcal{M}, p_{\text{eq}})$ . Moreover, this is the smoothest basis with respect to  $p_{\text{eq}}$  in the sense that each  $\varphi_j$  minimizes the norm  $\|\nabla f\|_{p_{\text{eq}}} = -\langle f, \hat{\mathcal{L}} f \rangle_{p_{\text{eq}}}$  subject to  $\varphi_j$  being orthogonal to all  $\{\varphi_l\}_{l < j}$ . The minimal value of the norm is given by the corresponding eigenvalue,  $\|\nabla \varphi_j\|_{p_{\text{eq}}} = \lambda_j$ .

Approximating the operator  $\hat{\mathcal{L}}$  was first achieved in [4] for compact manifolds using a fixed bandwidth kernel. However, in order to allow the manifold  $\mathcal{M}$  to be non-compact, it was shown in [5] that we must use a variable bandwidth diffusion kernel of the form,

$$K_\epsilon^S(x, y) = \exp\left(-\frac{\|x - y\|^2}{4\epsilon(q_\epsilon(x)q_\epsilon(y))^\beta}\right), \quad (7)$$

where  $q_\epsilon(x) = q(x) + \mathcal{O}(\epsilon)$  is an order- $\epsilon$  estimate of the sampling density and  $\beta$  should be negative so that the bandwidth function  $q_\epsilon(x)^\beta$  is large in regions of sparse sampling and small in regions of dense sampling. The algorithm presented below is closely related to that presented in [5] and is motivated by Corollary 1 of [5].

**Corollary 1.** *Let  $q \in L^1(\mathcal{M}) \cap \mathcal{C}^3(\mathcal{M})$  be a density that is bounded above on an embedded  $d$ -dimensional manifold  $\mathcal{M} \subset \mathbb{R}^n$  without boundary and let  $\{x_i\}_{i=1}^N$  be sampled independently with distribution  $q$ . Let  $K_\epsilon^S$  be a variable bandwidth kernel of the form (7) with bandwidth function  $q_\epsilon^\beta$  where  $q_\epsilon = q + \mathcal{O}(\epsilon)$  is any order- $\epsilon$  estimate of  $q$ . For a function  $f \in L^2(\mathcal{M}, q) \cap \mathcal{C}^3(\mathcal{M})$  and an arbitrary point  $x_i \in \mathcal{M}$ , define the discrete functionals,*

$$F_i(x_j) = \frac{K_\epsilon^S(x_i, x_j)f(x_j)}{q_\epsilon^S(x_i)^\alpha q_\epsilon^S(x_j)^\alpha}, \quad G_i(x_j) = \frac{K_\epsilon^S(x_i, x_j)}{q_\epsilon^S(x_i)^\alpha q_\epsilon^S(x_j)^\alpha},$$

where  $q_\epsilon^S(x_i) = \sum_l K_\epsilon^S(x_i, x_l)/q_\epsilon(x_i)^{d\beta}$ . Then,

$$\begin{aligned} I_{\epsilon, \alpha, \beta}^S f(x_i) &\equiv \frac{1}{\epsilon m q_\epsilon(x_i)^{2\beta}} \left( \sum_j F_i(x_j) - f(x_i) \right) \\ &= \hat{\mathcal{L}} f(x_i) + \mathcal{O}\left(\epsilon, \frac{q(x_i)^{(1-d\beta)/2}}{\sqrt{N}\epsilon^{2+d/4}}, \frac{\|\nabla f(x_i)\| q(x_i)^{-c_2}}{\sqrt{N}\epsilon^{1/2+d/4}}\right), \end{aligned} \quad (8)$$

with high probability, where  $c_1 = 2 - 2\alpha + d\beta + 2\beta$  and  $c_2 = 1/2 - 2\alpha + 2d\alpha + d\beta/2 + \beta$  and  $m$  is a constant depending on the form of the kernel, and  $m = 2$  for (7).

The key result of [5] is that for the error to be bounded when  $q$  is not bounded below, we require  $c_2 < 0$  (otherwise as  $q \rightarrow 0$  the final error term becomes unbounded). Since we are interested in the case  $c_1 = 1$ , in this paper we will use  $\beta = -1/2$  and  $\alpha = -d/4$ . Notice that this algorithm will require the intrinsic dimension  $d$  of the manifold  $\mathcal{M}$ , however we will determine this empirically as part of the kernel density estimation of the sampling density  $q$ .



To determine the sampling density  $q$  (which also serves as an estimate of the invariant measure  $p_{\text{eq}}$ ) we introduce the ad-hoc bandwidth function  $\rho_0(x) = \left( \frac{1}{k_0-1} \sum_{j=2}^{k_0} \|x_i - x_{I(i,j)}\|^2 \right)^{1/2}$ , where  $I(i,j)$  is the index of the  $j$ -th nearest neighbor of  $x_i$  from the data set. Following [5] we used  $k_0 = 8$  in all of our examples, and empirically the algorithm is not very sensitive to  $k_0$ . With this bandwidth we define the following kernel,

$$K_\epsilon(x, y) = \exp \left( -\frac{\|x - y\|^2}{2\epsilon\rho_0(x)\rho_0(y)} \right), \quad (9)$$

which will be used only for the kernel density estimate of the sampling density  $q$ . The kernel density estimate  $q_\epsilon$  of the sampling density  $q$  is given by the standard formula,

$$\begin{aligned} q_\epsilon(x_i) &\equiv \frac{1}{N(2\pi\epsilon\rho_0(x_i)^2)^{d/2}} \sum_{j=1}^N K_\epsilon(x_i, x_j) \\ &\approx \int_{\mathcal{M}} \frac{K_\epsilon(x_i, y)}{(2\pi\epsilon\rho_0(x_i)^2)^{d/2}} q(y) dV(y) = q(x_i) + \mathcal{O}(\epsilon), \end{aligned} \quad (10)$$

where  $dV$  is the volume form which  $\mathcal{M}$  inherits from the ambient space and  $q$  is the sampling measure relative to this volume form. Note that applying (10) requires choosing the bandwidth  $\epsilon$  and knowing the dimension  $d$  of the manifold.

To determine the bandwidth  $\epsilon$  and the dimension  $d$ , we apply the automatic tuning algorithm, originally developed in [15] and refined in [5]. The idea is that if  $\epsilon$  is not well tuned, the kernel will become trivial; when  $\epsilon$  is too small the kernel (9) is numerically zero when  $x \neq y$ , and when  $\epsilon$  is too large the kernel is numerically one. Forming the double sum  $T(\epsilon) \equiv \frac{1}{N^2} \sum_{i,j=1}^N K_\epsilon(x_i, x_j)$ , when  $\epsilon$  is too small we find  $T$  approaches  $1/N$  and when  $\epsilon$  is too large we find  $T$  approaches 1. As shown in [5, 15] when  $\epsilon$  is well tuned we have  $T(\epsilon) \approx \frac{(4\pi\epsilon)^{d/2}}{\text{vol}(\mathcal{M})}$  so that

$$\log T(\epsilon) \approx \frac{d}{2} \log(4\pi\epsilon) - \log(\text{vol}(\mathcal{M})). \quad (11)$$

Since  $T(\epsilon)$  is monotonically increasing in  $\epsilon$ , we also have  $\log(T(\epsilon))$  monotonically increasing in  $\log(\epsilon)$ , and so the derivative  $\frac{d \log(T(\epsilon))}{d \log(\epsilon)}$  has a unique maximum. Intuitively this maximum corresponds to  $\epsilon$  that gives the maximum ‘resolution’ of the kernel  $K_\epsilon$ . The approximation (11) suggests that the value of the maximum is  $\max_\epsilon \frac{d \log(T(\epsilon))}{d \log(\epsilon)} = \frac{d}{2}$ , and we will use this to determine the intrinsic dimension of the manifold  $\mathcal{M}$ . Since the summation  $T(\epsilon)$  is not very expensive to compute, we simply evaluate  $T(\epsilon)$  for  $\epsilon = 2^l$  where  $l = -30, -29.9, \dots, 9.9, 10$  and then compute the empirical derivative,

$$\frac{d(\log S)}{d(\log \epsilon)} \approx \frac{\log(S(\epsilon + h)) - \log(S(\epsilon))}{\log(\epsilon + h) - \log(\epsilon)}, \quad (12)$$

and choose the value of  $\epsilon$  which *maximizes* this derivative and set  $d = 2 \max_\epsilon \frac{d \log(T(\epsilon))}{d \log(\epsilon)}$ .

Now that we have the empirical estimate  $q_\epsilon(x_i) = q(x_i) + \mathcal{O}(\epsilon)$ , we can form the kernel (7). We reapply the same bandwidth selection to choose the global bandwidth  $\epsilon$  in the kernel (7) and a new estimate of  $d$ , by constructing  $T(\epsilon)$  as a double sum of the kernel (7) over the data set. Notice that these new values of  $\epsilon$  and  $d$  can be different from the previous values used in the kernel (10), although empirically the new dimension  $d$  is typically very similar.

We can now evaluate the kernel (7) on all pairs from the data set and form the matrix  $K_{\epsilon,i,j}^S = K_\epsilon^S(x_i, x_j)$  and we can compute the first normalization factor  $q_\epsilon^S(x_i) = \sum_{j=1}^N K_\epsilon^S(x_i, x_j) / q_\epsilon(x_i)^{d\beta}$  as in Corollary 1. We define a diagonal matrix  $D_{i,i} = q_\epsilon^S(x_i)$ , and the first normalization is to form the matrix  $K_{\epsilon,\alpha}^S = D^{-\alpha} K_\epsilon^S D^{-\alpha}$ . We then compute the second normalization factor  $q_{\epsilon,\alpha}^S(x_i) = \sum_{j=1}^N K_{\epsilon,\alpha,i,j}^S$  and form a diagonal matrix  $D_{\alpha,i,i} = q_{\epsilon,\alpha}^S(x_i)$ . The second normalization is to form the matrix  $\hat{K}_{\epsilon,\alpha}^S = D_\alpha^{-1} K_{\epsilon,\alpha}^S$ . We define the final normalization diagonal matrix  $\hat{D}_{i,i} = 2\epsilon q(x_i)^{2\beta}$ , and by Corollary 1,

$$L_{\epsilon,\alpha,\beta}^S = \hat{D}^{-1} (\hat{K}_{\epsilon,\alpha}^S - \text{Id}) = \hat{D}^{-1} D_\alpha^{-1} D^{-\alpha} K_\epsilon^S D^{-\alpha} - \hat{D}^{-1}$$

approximates the desired operator  $\hat{\mathcal{L}}$  when  $\beta = -1/2$  and  $\alpha = -d/4$ . To find the eigenvectors of  $L_{\epsilon,\alpha,\beta}^S$ , which approximate the eigenfunctions  $\varphi_j$  of  $\hat{\mathcal{L}}$ , we note that setting  $P = \hat{D}^{1/2} D_\alpha^{1/2}$  we can define a symmetric matrix,

$$\hat{L} \equiv P L_{\epsilon,\alpha,\beta}^S P^{-1} = P^{-1} D^{-\alpha} K_\epsilon^S D^{-\alpha} P^{-1} - \hat{D}^{-1}.$$

Since  $\hat{L}$  is a symmetric matrix, which is conjugate to the  $L_{\epsilon,\alpha,\beta}^S$ , we can compute the eigenvectors of  $\hat{L} = \hat{U} \hat{\Lambda} \hat{U}^\top$  efficiently and then the eigenvectors of  $L_{\epsilon,\alpha,\beta}^S$  are given by the column vectors of  $U = P^{-1} \hat{U}$ .

Note that the columns of  $\hat{U}$  will be numerically orthogonal, so the columns of  $U$  are orthogonal with respect to  $P^2$  since  $\text{Id} = \hat{U}^\top \hat{U} = U^\top P^2 U$ . A careful calculation based on the asymptotic expansions in [5] shows that  $P_{ii}^2 = q(x_i)^{c_1-1} + \mathcal{O}(\epsilon)$  and in general the  $q^{c_1}$  is the invariant measure of the gradient flow (6) so that  $P^2$  represents the ratio between the invariant measure  $e^{-c_1 U}$  of (6) and the sampling measure  $q$ . However, in this case since  $c_1 = 1$ , we have  $P = \text{Id} + \mathcal{O}(\epsilon)$ . Thus, for the case  $c_1 = 1$ , we will take the eigenvectors  $\varphi_j$  to be the column vectors of  $\hat{U}$ , since these eigenvectors are numerically orthogonal and are equal to the column vectors of  $U$  up to order- $\epsilon$ . Notice that the orthogonality of these vectors  $\frac{1}{N} \sum_{i=1}^N \varphi_l(x_i) \varphi_j(x_i) \approx \langle \varphi_l, \varphi_j \rangle_q$ , corresponds to the orthogonality of the eigenfunctions  $\varphi_j$  with respect to the sampling measure (and since  $c_1 = 1$  this is also the invariant measure of (6)). Finally, in order to insure that the eigenvectors are orthonormal, we renormalize each column vector so that  $\frac{1}{N} \sum_{i=1}^N \varphi_j(x_i)^2 = 1$ .

## APPENDIX B: ESTIMATING THE SEMI-GROUP SOLUTIONS FOR NON-ELLIPTIC OPERATORS

Consider a dynamical system for a state vector  $x$  on a manifold  $\mathcal{M} \subset \mathbb{R}^n$  given by,

$$dx = a(x) dt + b(x) dW_t \quad (13)$$

where  $W_t$  is a standard Brownian process (generated by the Laplacian on  $\mathcal{M}$ ),  $a$  is a vector field (which is *not necessarily* the gradient of a potential), and  $b$  a diffusion tensor, all defined on  $\mathcal{M}$ . Let  $x_i = x(t_i)$  be a time series realization of (13) at discrete times  $\{t_i\}_{i=1}^{N+1}$  with  $\tau = t_{i+1} - t_i$ . As in the previous section, we represent a smooth function  $f$  on the manifold by a vector  $f_i = f(x_i)$ . Using the time ordering of the data set, we can define the shift map,  $Sf(x_i) = f(x_{i+1})$ . Applying the Itô formula to the process  $y(t) = f(x(t))$  we have,

$$\begin{aligned} dy(s) &= \left( a \cdot \nabla f + \frac{1}{2} \text{Tr}(b^\top H(f)b) \right) ds + \nabla f^\top b dW_s \\ &\equiv \mathcal{L}f ds + \nabla f^\top b dW_s, \end{aligned} \quad (14)$$

where  $H(f)$  denotes the Hessian and the functions and derivatives are evaluated at  $x(s)$ . We first show that the expected value of the discrete time shift map is the semi-group solution  $e^{\tau\mathcal{L}}$  associated to the generator  $\mathcal{L}$  of the system (13).

For all smooth functions  $f$  defined on the data set, the shift map yields a function  $Sf$  which is defined on the first  $N - 1$  points of the data set. Rewriting (14) we have,

$$\begin{aligned} Sf(x_i) &= f(x(t_{i+1})) = y(t_{i+1}) \\ &= f(x_i) + \int_{t_i}^{t_{i+1}} \mathcal{L}f ds + \int_{t_i}^{t_{i+1}} \nabla f^\top b dW_s, \end{aligned} \quad (15)$$

and taking the expectation conditional to the state  $x_i$ ,

$$\mathbb{E}_{x_i}[Sf(x_i)] = f(x_i) + \int_{t_i}^{t_{i+1}} \mathbb{E}_{x_i}[\mathcal{L}f(x_s)] ds. \quad (16)$$

Recall that by the Feynman-Kac connection, the conditional expectation of the functional  $y(t_{i+1}) = f(x_{i+1})$  is given by the semi-group solution  $\mathbb{E}_{x_i}[y(t_{i+1})] = e^{\tau\mathcal{L}}f(x_i)$  and combining this with (16) we find,

$$e^{\tau\mathcal{L}}f(x_i) = f(x_i) + \int_{t_i}^{t_{i+1}} \mathbb{E}_{x_i}[\mathcal{L}f(x_s)] ds. \quad (17)$$

Substituting (17) into (15) we find,

$$\begin{aligned} Sf(x_i) &= e^{\tau\mathcal{L}}f(x_i) + \int_{t_i}^{t_{i+1}} \nabla f^\top b dW_s \\ &\quad + \int_{t_i}^{t_{i+1}} \mathcal{L}f - \mathbb{E}_{x_i}[\mathcal{L}f] ds. \end{aligned} \quad (18)$$

The formula (18) shows that the expectation of the shift map  $S$  is the semi-group solution  $e^{\tau\mathcal{L}}$  as claimed. This suggests that we can use the shift map to estimate the semi-group solution of the generator of (13). We next show that, by representing  $S$  in an appropriate basis, we can minimize the error of this estimate.

From (18), for any smooth function  $g$ , we can define the Monte-Carlo integral,

$$\begin{aligned} \langle g, Sf \rangle_{p_{\text{eq}}} &= \lim_{N \rightarrow \infty} \frac{1}{N-1} \sum_{i=1}^{N-1} g(x(t_i)) Sf(x(t_i)) \\ &= \langle g, e^{\tau\mathcal{L}}f \rangle_{p_{\text{eq}}} + \left\langle g, \int_{t_i}^{t_{i+1}} \nabla f^\top b dW_s \right\rangle_{p_{\text{eq}}} \\ &\quad + \left\langle g, \int_{t_i}^{t_{i+1}} Bf ds \right\rangle_{p_{\text{eq}}}, \end{aligned} \quad (19)$$

where we define  $Bf = \mathcal{L}f - \mathbb{E}_{x_i}[\mathcal{L}f]$ . The Monte-Carlo integral implies that the inner products should be taken with respect to the sampling measure for the training data set, and we assume that the evolution of  $x$  is ergodic so that the sampling measure is the invariant measure  $p_{\text{eq}}$  of the system (13). Note that for smooth functions  $f, g \in L^2(\mathcal{M}, p_{\text{eq}})$ , the final integral in (19) will be order- $\tau$  since it is deterministic and the inner product with  $g$  will be bounded. Therefore, our goal is to choose  $f, g$  from an orthonormal basis for  $L^2(\mathcal{M}, p_{\text{eq}})$  which minimizes the inner product with the stochastic integral, and thereby reduces the variance of our estimates of the coefficients.

We first expand the norm of  $\Omega(f) = \int_{t_i}^{t_{i+1}} \nabla f^\top b dW_s$  by applying the Itô isometry,

$$\begin{aligned} \|\Omega(f)\|_{p_{\text{eq}}}^2 &= \int_{\mathcal{M}} \left( \int_{t_i}^{t_{i+1}} \nabla f^\top b dW_s \right)^2 p_{\text{eq}}(x_i) dV(x_i) \\ &= \int_{\mathcal{M}} \int_{t_i}^{t_{i+1}} (\nabla f^\top b)^2 ds p_{\text{eq}}(x_i) dV(x_i) \\ &= \tau \|\nabla f^\top b\|_{p_{\text{eq}}}^2 + \mathcal{O}(\tau^2). \end{aligned} \quad (20)$$

In order to have a simple generic approach we will avoid estimating the diffusion tensor  $b$  by assuming that the norm  $\|b(x)\| = \sup \frac{\|v^\top b\|}{\|v\|}$  is bounded above on the manifold by a constant  $b_0$ . We can now apply the Cauchy-Schwartz inequality to find that,

$$\begin{aligned} |\langle g, \Omega(f) \rangle_{p_{\text{eq}}}| &\leq \|g\|_{p_{\text{eq}}} \|\Omega(f)\|_{p_{\text{eq}}} \\ &\leq \sqrt{\tau} b_0 \|g\|_{p_{\text{eq}}} \|\nabla f\|_{p_{\text{eq}}} + \mathcal{O}(\tau). \end{aligned} \quad (21)$$

The optimal basis will be the one that minimizes the norm  $\|\nabla f\|_{p_{\text{eq}}} = \int_{\mathcal{M}} |\nabla f|^2 p_{\text{eq}} dV(x)$ . As shown in Appendix A, the norm  $\|\nabla f\|_{p_{\text{eq}}}$  is provably minimized by the eigenfunctions of the generator  $\hat{\mathcal{L}}$  of the gradient flow system (6) and these eigenfunctions can be estimated by the diffusion kernel (7). Letting  $\varphi_i$  be the eigenfunctions of  $\hat{\mathcal{L}}$  with eigenvalues  $\lambda_i$  so that  $\hat{\mathcal{L}}\varphi_i = \lambda_i\varphi_i$ , the set

$\{\varphi_i\}$  is a basis for  $L^2(\mathcal{M}, p_{\text{eq}})$  and the norm we wish to minimize is given by the eigenvalues  $\|\nabla \varphi_i\|_{p_{\text{eq}}} = \lambda_i$ .

Replacing  $f$  and  $g$  in (19) with these eigenfunctions, for a finite data set we define  $\hat{A}_{lj} \equiv \frac{1}{N} \sum_{i=1}^N \varphi_j(x_i) S\varphi_l(x_i)$  and  $A_{lj} \equiv \langle \varphi_j, e^{\tau \mathcal{L}} \varphi_l \rangle_{p_{\text{eq}}}$ . Since  $\mathbb{E}[\varphi_j(x_i) S\varphi_l(x_i)] = A_{lj}$ , we have  $\mathbb{E}[\hat{A}_{lj}] = A_{lj}$  which shows that  $\hat{A}_{lj}$  is an unbiased estimate of  $A_{lj}$ . Using the error bounds derived above, the variance is,

$$\mathbb{E}[(\hat{A}_{lj} - A_{lj})^2] \leq b_0^2 \lambda_l^2 \tau N^{-1} + \mathcal{O}(\tau^2 N^{-1}),$$

assuming that  $x_i$  are independent. Since  $x_i$  form a time

series, they are not independent and the convergence of the Monte-Carlo integral will be slower if the dependence is strong. In that case, one may need to subsample the time series which simultaneously requires a larger data set. Assuming independence, by the Chebyshev bound,

$$P(|\hat{A}_{lj} - A_{lj}| \geq \epsilon) \leq b_0^2 \lambda_l^2 \tau \epsilon^{-2} N^{-1} + \mathcal{O}(\tau^2 \epsilon^{-2} N^{-1})$$

and balancing these error terms requires  $\lambda_l < b_0^{-1} \sqrt{\tau}$  and the errors are of order  $\epsilon = \mathcal{O}(\tau N^{-1/2})$  in probability.