

Prediction of dynamical time series using kernel based regression and smooth splines

Raymundo Navarrete and Divakar Viswanath

Department of Mathematics, University of Michigan
e-mail: raymundo@umich.edu; divakar@umich.edu

Abstract: Prediction of dynamical time series with additive noise using support vector machines or kernel based regression is consistent for certain classes of discrete dynamical systems. Consistency implies that these methods are effective at computing the expected value of a point at a future time given the present coordinates. However, the present coordinates themselves are noisy, and therefore, these methods are not necessarily effective at removing noise. In this article, we consider denoising and prediction as separate problems for flows, as opposed to discrete time dynamical systems, and show that the use of smooth splines is more effective at removing noise. Combination of smooth splines and kernel based regression yields predictors that are more accurate on benchmarks typically by a factor of 2 or more. We prove that kernel based regression in combination with smooth splines converges to the exact predictor for time series extracted from any compact invariant set of any sufficiently smooth flow. As a consequence of convergence, one can find examples where the combination of kernel based regression with smooth splines is superior by even a factor of 100. The predictors that we analyze and compute operate on delay coordinate data and not the full state vector, which is typically not observable.

Received December 2017.

1. Introduction

The problem of time series prediction is to use knowledge of a signal $x(t)$ for $0 \leq t \leq T$ and infer its value at a future time $t = T + t_f$, where t_f is positive and fixed. A time series is not predictable if it is entirely white noise. Any prediction scheme has to make some assumption about how the time series is generated. A common assumption is that the observation $x(t)$ is a projection of the state of a dynamical system with noise superposed [8]. Since the state of the dynamical system can be of dimension much higher than 1, delay coordinates are used to reconstruct the state. Thus, the state at time t may be captured as

$$(x(t), x(t - \tau), \dots, x(t - (D - 1)\tau)) \quad (1.1)$$

where τ is the delay parameter and D is the embedding dimension. Delay coordinates are (generically) effective in capturing the state correctly provided $D \geq 2d + 1$, where d is the dimension of the underlying dynamics [21].

Farmer and Sidorowich [8] used a linear framework to compute predictors applicable to delay coordinates. It was soon realized that the nonlinear and more general framework of support vector machines would yield better predictors [15, 18, 19]. Detailed computations demonstrating the advantages of kernel based predictors were given by Müller et al. [19] and are also discussed in the textbook of Schölkopf and Smola [22]. Kernel methods still appear to be the best, or among the best, for the prediction of stationary time series [16, 20].

A central question in the study of noisy dynamical time series is how well that noise can be removed to recover the underlying dynamics. Lalley, and later Nobel, [12, 14, 13] have examined hyperbolic maps of the form $x_{n+1} = F(x_n)$, with $F: \mathbb{R}^d \rightarrow \mathbb{R}^d$. It is assumed that observations are of the form $y_n = x_n + \epsilon_n$, where ϵ_n is iid noise. They proved that it is impossible to recover x_n from y_n , even if the available data y_n is for $n = 0, -1, -2, \dots$ and infinitely long, if the noise is normally distributed. However, if the noise satisfies $|\epsilon_n| < \Delta$ for a suitably small Δ , the underlying signal x_n can be recovered. The recovery algorithm does not assume any knowledge of F . The phenomenon of unrecoverability is related to homoclinic points. If the noise does not have compact support, with some nonzero probability, it is impossible to distinguish between homoclinic points.

Lalley [14] suggested that the case of flows could be different from the case of maps. In discrete dynamical systems, there is no notion of smoothness across iteration. In the case of flows, the underlying signal will depend smoothly on time but the noise, which is assumed to be iid at different points in time, will not. Lalley's algorithm for denoising relies on dynamics and, in particular, on recurrences. In the case of flows, we rely solely on smoothness of the underlying signal for denoising. As predicted by Lalley, the case of flows is different. Denoising based on smoothness of the underlying signal alone can handle normally distributed noise or other noise models. Thus, our algorithms are split into two parts: first the use of smooth splines to denoise, and second the use of kernel based regression to compute the predictor. Only the second part relies on recurrences.

Prediction of discrete dynamics, within the framework of Lalley [12], has been considered by Steinwart and Anghel [24] (also see [3]). Suppose $x_n = F^n(x_0)$ and $\tilde{x}_n = x_n + \epsilon_n$ is the noisy state vector. The risk of a function f is defined as

$$\int \int |F(x) + \epsilon_1 - f(x + \epsilon_2)|^2 d\nu(\epsilon_1) d\nu(\epsilon_2) d\mu(x),$$

where ν is the distribution of the noise and μ is a probability measure invariant under F and with compact support. Thus, the risk is a measure of how well the noisy future state vector can be predicted given the noisy current state vector. It is proved that kernel based regression is consistent with respect to this notion of risk for a class of rapidly mixing dynamical systems. Although the notion of risk does not require denoising, consistency of empirical risk minimization is proved for additive noise ϵ_n of compact support as in [12]. In the case of

empirical risk minimization, compactness of added noise is not a requirement imposed by the underlying dynamics but is assumed to make it easier to apply universality theorems.

Our results differ in the following ways. We consider flows and not discrete time maps. In addition, we work with delay coordinate embedding [21] and do not require the entire state vector to be observable. Finally, we prove convergence to the exact predictor, which goes beyond consistency. The convergence theorem we prove is not uniform over any class of dynamical systems. However, we do not assume any type of decay in correlations or rapid mixing. Non-uniformity in convergence is an inevitable consequence of proving a theorem that is applicable to any compact invariant set of a generic finite dimensional dynamical system [1, 9, 25]. This point is further discussed in Section 2, which presents the main algorithm as well as a statement of the convergence theorem. Section 3 presents a proof of the convergence theorem.

In Section 4, we present numerical evidence of the effectiveness of combining spline smoothing and kernel based regression. The algorithm of Section 2 is compared to computations reported in [19] and the spline smoothing step is found to improve accuracy of the predictor considerably. The numerical examples bring up two points that go beyond either consistency or convergence. First, we explain heuristically why it is not a good idea to iterate 1-step predictor k -times to predict the state k steps ahead. Rather, it is a much better idea to learn the k -step predictor directly. Second, we point out that no currently known predictor splits the distance vector between stable and unstable directions, a step which was argued to be essential for an optimal predictor by Viswanath et al. [28]. The heuristic explanation for why iterating a 1-step predictor k times is not a good idea relies on the same principle.

The concluding discussion in Section 5 points out connections to related lines of current research in parameter inference [16, 17] and optimal consistency estimates for stationary data [11].

2. Prediction algorithm and statement of convergence theorem

Let $\frac{dU}{dt} = \mathcal{F}(U)$, where $\mathcal{F} \in C^r(\mathbb{R}^d, \mathbb{R}^d)$, $r \geq 2$, define a flow that may be limited to an open subset of \mathbb{R}^d with compact closure. Let $\mathcal{F}_t(U_0)$ be the time- t map with initial data U_0 . It is assumed that $U(t; U_0)$, $t \in \mathbb{R}$, is a trajectory of the flow whose initial point $U(0; U_0)$ is $U_0 \in \mathbb{R}^d$. Let $\tilde{\mu}$ be a compactly supported invariant probability measure of the flow-map \mathcal{F}_t for $t > 0$ and let \tilde{X} be its support. It is assumed that the initial point $\tilde{\omega}$ is drawn from the measure $\tilde{\mu}$. For $\tilde{\omega} \in \tilde{X}$, the trajectory $U(t; \tilde{\omega})$ exists for all $t \in \mathbb{R}$ and is unique. In addition, the flow is assumed to be ergodic with respect to the measure $\tilde{\mu}$.

Let $\pi : \mathbb{R}^d \rightarrow \mathbb{R}$ be a generic nonlinear projection. Let $u(t; \tilde{\omega}) = \pi U(t; \tilde{\omega})$ be the projection of the random trajectory $U(t; \tilde{\omega})$. By the embedding theorem of Sauer et al. [21], we assume that the delay coordinates give a C^r diffeomorphism into the state space implying that $U(t; \tilde{\omega})$ can be recovered from the delay vector, with delay $\tau > 0$,

$$(u(t; \tilde{\omega}), u(t - \tau; \tilde{\omega}), \dots, u(t - (D - 1)\tau; \tilde{\omega}))$$

for $D \geq 2d + 1$. This delay vector is denoted by $u(t; \tau; \tilde{\omega})$.

As a consequence of the C^r embedding, there is a measure μ compactly supported in \mathbb{R}^D that corresponds to $\tilde{\mu}$. The measure μ is ergodic and invariant under the flow lifted via the embedding. Denote the compact support of μ by X . For every point $\tilde{\omega}$ in \tilde{X} , there corresponds a unique point ω in X and vice versa. Because the prediction algorithm is based on delay coordinates and not the state vector, it is more convenient to work in the embedding space \mathbb{R}^D and in terms of ω and μ . Therefore, we will rely on the bijective correspondence between X and \tilde{X} and use the notation $u(t; \tau; \omega)$ instead of $u(t; \tau; \tilde{\omega})$ and $u(t; \omega)$ instead of $u(t; \tilde{\omega})$. With these conventions, $u(t; \tau; \omega)$ can be thought of as the path in \mathbb{R}^D with $u(0; \tau; \omega) = \omega$. Similarly, $u(t; \omega)$ can be thought of as a real-valued signal with $u(0; \omega) = \omega_1$, where ω_1 is the first component of $\omega \in \mathbb{R}^D$. In later arguments, the assumption that ω is μ -distributed will be significant, and so will be the ergodicity of the flow with respect to μ .

Given the signal $u(t; \omega)$, it is assumed that the recorded observations are $u_\eta(jh; \omega) = u(jh; \omega) + \epsilon_j$, where ϵ_j is iid noise. Following Eggermont and LaRiccia [5, 6], we assume that $\mathbb{E}\epsilon_j = 0$ and $\mathbb{E}|\epsilon_j|^\kappa < \infty$ for some $\kappa > 3$. To avoid inessential technicalities it is assumed that $\tau/h \in \mathbb{Z}^+$ so that the delay is an integral multiple of the time step h . In particular, we set $\tau = nh$. Similarly, we assume $t_f = n_f\tau$, $n_f \in \mathbb{Z}^+$, where t_f is the look-ahead into the future. The noisy delay coordinates $u_\eta(jh; \tau; \omega)$ are assumed to be available for $j = 0, \dots, (N + n_f)n$, which implies that the observation interval of $u_\eta(t; \omega)$ is $t \in [-(D - 1)\tau, N\tau + t_f]$.

The exact predictor $F : \mathbb{R}^D \rightarrow \mathbb{R}$ is a C^r function such that $F(u(t; \tau; \omega)) = u(t + t_f; \omega)$ for $\omega \in X$. Lemma 3 proves uniqueness and existence of the exact predictor F . The exact predictor F corresponds to a fixed $t_f > 0$, but that dependence is not shown in the notation. The problem as considered by Müller et al. [19] is to recover the exact predictor F from the noisy observations $u_\eta(jh; \omega)$. Let $|\cdot|_\epsilon$ denote Vapnik's ϵ -loss function. The algorithm of Müller et al. computes f_m such that the functional

$$\frac{1}{Nn + 1} \sum_{j=0}^{Nn} |f(u_\eta(jh; \tau; \omega)) - u_\eta(jh + \tau; \omega)|_\epsilon + \Lambda \|f\|_{K_\gamma}^2 \quad (2.1)$$

is minimized for $f = f_m$ in the reproducing kernel Hilbert space \mathcal{H}_{K_γ} corresponding to the kernel K_γ . The kernel K_γ is assumed to be given by $K_\gamma(x, y) = \exp\left(-\frac{\sum_{i=1}^D (x_i - y_i)^2}{\gamma^2}\right)$. The kernel bandwidth parameter γ and the Lagrange multiplier Λ are both determined using cross-validation. This method approximates the exact predictor F for $t_f = \tau$. If $t_f = n_f\tau$, $n_f \in \mathbb{Z}^+$, the approximation is iterated n_f times. We will compare our predictor against that of Müller et al. using some of the same examples and the same framework as they do in Section 4.

In our algorithm, the first step is to apply spline smoothing. In particular, we apply cubic spline smoothing [4] to compute a function $u_s(t; \omega)$, $t \in [-(D - 1)\tau,$

$N\tau + t_f]$ such that the functional

$$\frac{1}{(N + n_f + D - 1)n + 1} \sum_{j=-(D-1)n}^{(N+n_f)n} (u_\eta(jh; \omega) - \tilde{u}(jh))^2 + \lambda \int_{-(D-1)\tau}^{N\tau+t_f} \tilde{u}''(t)^2 dt \tag{2.2}$$

is minimum for $\tilde{u} = u_s(\cdot; \omega)$ over $\tilde{u} \in W^{2,2}[-(D-1)\tau, N\tau + t_f]$, where $W^{2,2}[a, b]$ denotes the Sobolev space of twice-differentiable functions $g : [a, b] \rightarrow \mathbb{R}$ with the norm $\|g\|^2 = \|g\|_2^2 + \|g'\|_2^2 + \|g''\|_2^2$. The parameter λ is determined using five-fold cross-validation. The minimizer $u_s(t; \omega)$ depends upon the noise-free signal $u(t; \omega)$ as well as the instantiation of the iid noise in $u_\eta(jh; \omega)$ for $-(D-1)n \leq j \leq (N + n_f)n$. However, the dependence on the iid noise is not shown in the notation.

The second step of our algorithm is similar to the method of Müller et al. The predictor f_1 is computed as

$$f_1 = \operatorname{argmin}_{f \in \mathcal{H}_k} \frac{1}{Nn + 1} \sum_{j=0}^{Nn} (f(u_s(jh; \tau; \omega)) - u_s(jh + t_f; \omega))^2 + \Lambda \|f\|_{K_\tau}^2. \tag{2.3}$$

Both the parameters γ and Λ are determined using five-fold cross-validation. Here n_f and therefore t_f are fixed because we seek to approximate the exact predictor with lookahead fixed at t_f . As explained in Section 4, it is significant that the predictor directly optimizes with a lookahead of t_f . Iterating a τ -step predictor n_f times gives worse predictions.

The second step (2.3) differs from the algorithm of Müller et al. in using the spline smoothed signal $u_s(t; \omega)$ in place of the noisy signal $u_\eta(t; \omega)$. Our algorithm relies mainly on spline smoothing to eliminate noise. Yet another difference is that we use the least squares loss function in place of the ϵ -loss function. This difference is a consequence of relying on spline smoothing to eliminate noise. As explained by Christmann and Steinwart [3], the ϵ -loss function, Huber’s loss, and the L^1 loss function are used to handle outliers. However, spline smoothing eliminates outliers, and we choose the L^2 loss function because of its algorithmic advantages.

We now turn to a discussion of the convergence of the predictor f_1 to the exact predictor F . The first step is to assess the accuracy of spline smoothing. We quote the following lemma, which is a convenient restatement of a result of Eggermont and LaRiccia [5, 6] (see pages 132 and 133 of [6]). In the lemma, $W^{m,2}[a, b]$ denotes the Sobolev space of m -times differentiable functions $g : [a, b] \rightarrow \mathbb{R}$ with norm $\|g\|^2 = \sum_{j=0}^m \|g^{(j)}\|_2^2$.

Lemma 1. *Assume $2 \leq m \leq r$. Suppose that $u(t; \omega)$ is a signal defined for $t \in \mathbb{R}$ with $\omega \in X$. For $j = -(D-1)n, \dots, Nn + n_f$, let $y_j = u(jh; \omega) + \epsilon_j$, where $h = \tau/n$ and where ϵ_j are iid random variables. It is further assumed that $\mathbb{E}\epsilon_j = 0$, $\mathbb{E}\epsilon_j^2 = \sigma^2$, and $\mathbb{E}|\epsilon_j|^\kappa < \infty$ for some $\kappa > 3$. Let $u_s(t) \in W^{m,2}[-(D-1)\tau, N\tau + t_f]$*

be the spline that minimizes the functional

$$\frac{1}{n(N + D - 1) + n_f + 1} \sum_{j=-(D-1)n}^{(N+n_f)n} (\tilde{u}(jh) - y_j)^2 + \lambda \int_{-(D-1)\tau}^{N\tau+t_f} |\tilde{u}^{(m)}(t)|^2 dt$$

over $\tilde{u} \in W^{m,2}[-(D-1)\tau, N\tau]$. Assume

$$\lambda = \left(\frac{\log(n(N + n_f + D - 1))}{n(N + n_f + D - 1)} \right)^{\frac{2m}{2m+1}}.$$

Let $p = \mathbb{P}(n, N, \Delta, \omega)$ be the probability that

$$\|u_s(\cdot; \omega) - u(\cdot; \omega)\|_{\infty} > \Delta > 0,$$

where the ∞ -norm is over the interval $[-(D-1)\tau, N\tau + t_f]$. Then $\lim_{n \rightarrow \infty} \mathbb{P}(n, N, \Delta, \omega) = 0$.

Some remarks about the connection of this lemma to the algorithm given by (2.2) and (2.3) follow. First, the lemma assumes a fixed choice of λ (the relevant theorem in [5, 6] in fact allows λ to lie in an interval). In our algorithm, λ is determined using cross-validation because of its practical effectiveness [29].

Second, the probability $\mathbb{P}(n, N, \Delta, \omega)$ (which may be interpreted as the probability that spline smoothing fails to denoise effectively) depends on ω and therefore on the particular trajectory. If $\mathbb{P}(n, N, \Delta, \omega)$ depends on ω only through a bound on the m -th derivative of $u(t; \omega)$, $t \in [-(D-1)n, nN]$, the bound would be uniform for all trajectories on the compact invariant set X . The achievability part of Stone's optimality result [26] gives such a bound but the algorithm in that proof does not appear practical. Proving a similar result for smooth splines based on the existing literature does not appear entirely straightforward. In the L^2 norm, some uniform bounds have been proved for smooth splines by Györfi et al. [10]. A bound on the L^2 norm can be combined with a bound on the m -th derivative using a Sobolev inequality to obtain an ∞ -norm bound. Although the rate of convergence would be slightly sub-optimal, it would suffice for our purposes. However, the result of Györfi et al. is for expectations and not for convergence in probability, and an argument using Chebyshev's inequality does not give strong bounds.

The convergence analysis of the second half of the algorithm also alters the algorithm slightly. In particular, the use of cross-validation to choose parameters is not a part of the analysis. To state the convergence theorem, we first fix $\epsilon > 0$. By the universality theorem of Steinwart [23], we may choose $F_{\epsilon} \in \mathcal{H}_{K_{\gamma}}$ such that $\|F_{\epsilon} - F\|_{\infty} < \epsilon$ in a compact domain that has a non-empty interior and contains the invariant set X . The convergence theorem also makes the technical assumption $\epsilon^2 / \|F_{\epsilon}\|_{K_{\gamma}}^2 < 1$, which may always be satisfied by taking ϵ small enough.

The choice of the kernel-width parameter γ is important in practice. In the convergence proof, the choice of γ is not explicitly considered. However, γ still plays a role because $\|F_{\epsilon}\|_{K_{\gamma}}$ depends upon γ .

The parameter Λ in (2.3) is fixed as $\Lambda = \epsilon^2 / \|F_\epsilon\|_{K,\gamma}^2$ for the proof. Next we pick $\delta = \epsilon^{1/2}$ and $\ell \in \mathbb{Z}^+$ such that the covering of the invariant set X using boxes of dimension $2^{-\ell}$ ensures that the variation of F_ϵ (as well as that of the exact predictor F and f_3 , which is defined later) within each box is bounded by $\delta/4$.

Suppose A_1, \dots, A_L are boxes of dimension $2^{-\ell}$ that cover X in the manner hinted above. We next choose T^* such that the measure of the trajectories (with respect to the ergodic measure μ) that sample each one of the boxes A_j adequately (in a sense that will be explained) is greater than $1 - \epsilon$ if the time interval of the trajectory exceeds T^* .

The parameter Δ is a bound on the infinite norm accuracy of the smooth spline as in Lemma 1. Choose $\Delta > 0$ small enough that

$$\frac{B_1 \Delta^{1/2}}{\Lambda} = \frac{B_1 \Delta^{1/2} \|F_\epsilon\|_K^2}{\epsilon^2} < \epsilon^{1/2},$$

where B_1 is a constant specified later. The main purpose of increasing n is to make spline smoothing accurate. However, the following condition requiring n to be large enough is assumed in the proof:

$$\frac{B_1 h^{1/2}}{\Lambda} = \frac{B_1 \tau^{1/2} \|F_\epsilon\|_K^2}{\epsilon^2 n^{1/2}} < \epsilon^{1/2}.$$

Within this set-up, we have the following convergence theorem.

Theorem 2. For $\epsilon > 0$, $T > T^*$, $N = T/\tau$, and Λ, Δ chosen as above, we have

$$\mu \left\{ x \in X \mid |f_1(x) - F(x)| > 3\sqrt{\epsilon} \right\} < \frac{8\epsilon}{1 - \epsilon},$$

when f_1 is constructed (or learnt) from the signal $u_\eta(t; \omega)$, $t \in [-(D-1)\tau, N\tau]$, for $\{\omega \in X\}$ of μ -measure greater than $1 - \epsilon$ and with probability $1 - \mathbb{P}(n, N, \Delta, \omega)$ (probability of successful denoising in the spline-smoothing step) tending to 1 in the limit $n \rightarrow \infty$.

Nonuniform bounds implying a form of weak consistency are considered by Steinwart, Hush, and Scovel [25]. However, the algorithm of (2.2) and (2.3) does not fit into the framework of [25]. The application of spline smoothing to produce $u_s(t; \omega)$ means that $u_s(t; \omega)$ may not be stationary, and our method of analysis does not rely on verifying a weak law of large numbers as in [25]. The analysis summarized above and given in detail in the following section relies on ∞ -norm bounds.

3. Proof of convergence

We begin the proof with a more complete account of how the embedding theorem is applied. Let $\frac{dU}{dt} = \mathcal{F}(U)$, where $\mathcal{F} \in C^r(\mathbb{R}^d, \mathbb{R}^d)$, $r \geq 2$, be a flow. Let $\mathcal{F}_t(U_0)$ be the time- t map with initial data U_0 . Let $\tilde{V} \subset \mathbb{R}^d$ be an open set

with compact closure. If $U_0 \in \tilde{V}$, it is assumed that $\mathcal{F}_t(U_0)$ is well-defined for $-\tau D \leq t \leq n_f \tau = t_f$, where D is the embedding dimension.

Assumption: For embedding dimension $D \geq 2d + 1$ and a suitably chosen delay $\tau > 0$, the map

$$x \rightarrow (\pi x, \pi \mathcal{F}_{-\tau} x, \pi \mathcal{F}_{-2\tau} x, \dots, \pi \mathcal{F}_{-(D-1)\tau} x)$$

is a C^r diffeomorphism between \tilde{V} and its image in \mathbb{R}^D . This assumption is generically true [21]. This map is called the delay embedding. Denote the image of \tilde{V} under the delay embedding by V .

The invariant measures $\tilde{\mu}$ and μ as well as \tilde{X} , X , $\tilde{\omega}$, ω , $u(t; \omega)$, and $u(t; \tau; \omega)$ are as defined earlier. It is assumed that $\tilde{X} \subset \tilde{V}$, which implies $X \subset V$.

Lemma 3. Suppose $\frac{dU(t)}{dt} = \mathcal{F}(U(t))$ for $-\tau D \leq t \leq t_f$, $U(0) = U_0 \in \tilde{V}$. Denote the delay vector

$$(\pi U_0, \pi \mathcal{F}_{-\tau} U_0, \dots, \pi \mathcal{F}_{-(D-1)\tau} U_0)$$

by $U_{0,\tau}$ so that $U_{0,\tau} \in V$. There exists a unique and well-defined C^r function $F : V \rightarrow \mathbb{R}$, called the exact predictor, such that

$$F(U_{0,\tau}) = \pi \mathcal{F}_{t_f}(U_0)$$

for all $U_{0,\tau} \in V$. In particular, $F(u(t; \tau; \omega)) = u(t + t_f; \omega)$ for all $t \in \mathbb{R}$ and all $\omega \in X$.

Proof. To map $U_{0,\tau} \in V$ to $\pi \mathcal{F}_{t_f}(U_0)$, first invert the delay map to obtain the point U_0 in \tilde{V} , advance that point by t_f by applying \mathcal{F}_{t_f} , and finally project using π . Each of the three maps in this composition is C^r or better. The predictor must be unique because \mathcal{F}_{t_f} is uniquely determined by the flow. \square

Remark. The embedding theory of Sauer et al. [21] may be applied to the compact invariant set \tilde{X} without enclosing it in the open set \tilde{V} . Indeed, if the box counting dimension of \tilde{X} is d' , the embedding dimension need only satisfy $D \in \mathbb{Z}^+$ and $D > 2d'$. That can be advantageous because we may have d' much smaller than d . However, there are two difficulties if \tilde{X} is a fractal set. First, tangent spaces cannot be defined and we cannot assert the delay map to be a diffeomorphism although it will be one-one generically. Second, we will need to extend F to the closure of an open neighborhood of X in \mathbb{R}^D to apply the universality theorem, and such an extension cannot be made from X if X is a fractal set. Both these difficulties go away if we take \tilde{V} to be a submanifold that contains \tilde{X} . If d' is the dimension of \tilde{V} , we would only require $D > 2d'$. For simplicity, we have assumed \tilde{V} to be an open set.

The following convexity lemma is an elementary result of convex analysis [7]. It is stated and proved for completeness.

Lemma 4. Let $\mathcal{L}_1(f)$ and $\mathcal{L}_2(f)$ be convex and continuous in f , where $f \in \mathcal{H}$ and \mathcal{H} is a Hilbert space. If $w \in \nabla \mathcal{L}_i(f)$, the subgradient at f , assume that

$$\mathcal{L}_i(f + g) - \mathcal{L}_i(f) - \langle w, g \rangle \geq \lambda \langle g, g \rangle / 2$$

for $\lambda > 0$, all $g \in \mathcal{H}$, and $i = 1, 2$. Let $f_1 = \operatorname{argmin} \mathcal{L}_1(f)$ and $f_2 = \operatorname{argmin} \mathcal{L}_2(f)$. Suppose that

$$|\mathcal{L}_1(f) - \mathcal{L}_2(f)| \leq \delta$$

for $\|f\| \leq r$, and assume that $\|f_1\| < r$ and $\|f_2\| < r$. Then,

$$\|f_1 - f_2\|^2 \leq \frac{2\delta}{\lambda}.$$

Proof. Because f_1 minimizes $\mathcal{L}_1(f)$, we have $0 \in \nabla \mathcal{L}_1(f_1)$. Thus,

$$\mathcal{L}_1(f_2) - \mathcal{L}_1(f_1) \geq \lambda \|f_2 - f_1\|^2 / 2.$$

Similarly, $\mathcal{L}_2(f_1) - \mathcal{L}_2(f_2) \geq \lambda \|f_2 - f_1\|^2 / 2$. By adding the two inequalities, we have

$$\|f_2 - f_1\|^2 \leq \frac{|\mathcal{L}_1(f_2) - \mathcal{L}_1(f_1) + \mathcal{L}_2(f_1) - \mathcal{L}_2(f_2)|}{\lambda} \leq \frac{2\delta}{\lambda},$$

proving the lemma. This last step relies on $\|\mathcal{L}_i(f_1) - \mathcal{L}_i(f_2)\| \leq \delta$ and the assumption $\|f_1\|, \|f_2\| < r$. \square

If $u(t; \omega)$, $t \in [-(D-1)\tau, N\tau + t_f]$, is the noise-free signal, our arguments are phrased under the assumption that $|u(t; \omega) - u_s(t; \omega)| \leq \Delta$. This assumption is realized with probability $1 - \mathbb{P}(n, N, \Delta, \omega)$, which tends to 1 as n increases (by Lemma 1). For convenience, we denote $\mathbb{P}(n, N, \Delta, \omega)$ by p . The probability that $u_\eta(t; \omega)$ is successfully denoised by smooth splines so that $|u(t; \omega) - u_s(t; \omega)| \leq \Delta$ is then $1 - p$.

In general, a C^r function defined on an embedded submanifold can be extended to an open neighborhood of the submanifold using a partition of unity. Because $V \subset \mathbb{R}^D$ is an embedded submanifold, $X \subset V$, and the exact predictor F is defined on V , it follows that there exists $M > 0$ such that F can be extended to Y , where

$$Y = \{y \mid \|y - \omega\|_\infty \leq M \text{ for some } \omega \in X\}.$$

We will always assume $\Delta < M$ so that the spline-smoothed signal maps to Y under delay embedding with probability greater than $1 - p$. Without loss of generality, we assume $M \leq 1$. The convergence proof will assess the approximation to F with respect to the measure μ . Therefore, the manner in which the extension is carried out is not highly relevant. The sole purpose of the extension is to facilitate an application of the universality theorem for Gaussian kernels.

Let

$$B = \sup_{\omega \in X} \|\omega\|_\infty + M \tag{3.1}$$

Thus, B is a bound on the size of the embedded invariant set with ample allowance for error in spline smoothing.

Let $u_s(t; \omega)$ denote the spline-smoothed signal and $u(t; \omega)$ the noise-free signal with $\omega \in X$. Define

$$\mathcal{W}_1(f) = \frac{1}{Nn+1} \sum_{j=0}^{Nn} (f(u_s(jh; \tau; \omega)) - u_s(jh + t_f; \omega))^2 + \Lambda \|f\|_K^2,$$

where $t_f = n_f\tau$, $n_f \in \mathbb{Z}^+$, and K is any smooth and positive kernel defined over $Y \times Y$. The kernel K will be specialized to the Gaussian kernel K_γ when applying the universality theorem. Define

$$\mathcal{W}_2(f) = \frac{1}{Nn + 1} \sum_{j=0}^{Nn} (f(u(jh; \tau; \omega)) - u(jh + t_f; \omega))^2 + \Lambda \|f\|_K^2$$

using the noise-free signal $u(t; \omega)$. Let $T = N\tau$ and define

$$\mathcal{W}_3(f) = \frac{1}{T} \int_0^T (f(u(t; \tau; \omega)) - u(t + t_f; \omega))^2 dt + \Lambda \|f\|_K^2.$$

For $\Lambda > 0$, all three functionals are strictly convex and have a unique minimizer. The unique minimizers of \mathcal{W}_1 , \mathcal{W}_2 , and \mathcal{W}_3 are denoted by f_1 , f_2 , and f_3 , respectively. The functional \mathcal{W}_1 is the same as in (2.3), the second step of the algorithm. Thus, f_1 is the computed approximation to the exact predictor F .

The following lemma bounds the minimizers of $\mathcal{W}_1(f)$, $\mathcal{W}_2(f)$, $\mathcal{W}_3(f)$ in norm by $B/\Lambda^{1/2}$.

Lemma 5. *The minimizer f_1 satisfies $\|f_1\|_K \leq \frac{B}{\Lambda^{1/2}}$ with probability greater than $1 - p$. The minimizers f_2 and f_3 satisfy $\|f_2\|_K \leq \frac{B}{\Lambda^{1/2}}$ and $\|f_3\|_K \leq \frac{B}{\Lambda^{1/2}}$.*

Proof. Because f_1 minimizes $\mathcal{W}_1(f)$, we must have $\mathcal{W}_1(f_1) \leq \mathcal{W}_1(0)$. We have $\mathcal{W}_1(0) \leq B^2$ with probability greater than $1 - p$. Thus, $\Lambda \|f_1\|_K^2 \leq \mathcal{W}_1(f_1) \leq \mathcal{W}_1(0) \leq B^2$ and the stated bound for $\|f_1\|_K$ follows. The bounds for f_2 and f_3 are proved similarly. \square

Lemma 6. *Assume $0 < \Lambda \leq 1$ and $|u(t; \omega) - u_s(t; \omega)| \leq \Delta$ for $t \in [-(D - 1)\tau, T]$. For $f \in \mathcal{H}_K$ with $\|f\|_K \leq \frac{B}{\Lambda^{1/2}}$, we have $|\mathcal{W}_1(f) - \mathcal{W}_2(f)| \leq \frac{B_1^2 \Delta}{\Lambda}$. Here B_1 depends only on B and the kernel K . The kernel K is assumed to be C^2 .*

Proof. First, we note that $\|f\|_\infty \leq c_0 \|f\|_K$ and $\|\partial f\|_\infty \leq c_1 \|f\|_K$, where ∂ is the directional derivative of f in any direction. By a result of Zhou (part (c) of Theorem 1 of [30]), we may take $c_0 = \|K(x, y)\|_\infty$ and $c_1 D^{-1/2} = \|K(x, y)\|_\infty + \sum \|\partial_{x_i} K(x, y)\|_\infty + \sum \|\partial_{x_i} \partial_{x_j} K(x, y)\|_\infty$, where D is the embedding dimension and the ∞ -norm is over $x, y \in Y$. If we define B'_1 using

$$B'_1 = \max(B, c_0 B, c_1 B), \tag{3.2}$$

it follows that both $\|f\|_\infty$ and $\|\partial f\|_\infty$ (where ∂ is a directional derivative in any direction) are bounded above by $B'_1/\Lambda^{1/2}$.

We may write

$$|\mathcal{W}_1(f) - \mathcal{W}_2(f)| \leq \frac{1}{Nn + 1} \sum_{j=0}^{Nn} \frac{4B'_1}{\Lambda^{1/2}} \left(|f(u_s(jh; \tau; \omega)) - f(u(jh; \tau; \omega))| + |u_s(jh; \tau; \omega) - u(jh; \tau; \omega)| \right). \tag{3.3}$$

Here $\frac{4B'_1}{\Lambda^{1/2}}$ is used an upper bound on $|f(u_s(jh; \tau))| + |f(u(jh; \tau))| + |u_s(jh; \tau)| + |u(jh; \tau)|$. The bound of $B'_1/\Lambda^{1/2}$ on $|f|$ is justified by the previous paragraph. The same bound on $|u_s|$ and $|u|$ follows from $B'_1 < B$ and $\Lambda \leq 1$.

Now, $|u_s(jh; \tau; \omega) - u(jh; \tau; \omega)| \leq \Delta$ implies that

$$|f(u_s(jh; \tau; \omega)) - f(u(jh; \tau; \omega))| \leq B'_1 \Delta / \Lambda^{1/2}$$

by the bound on $\|\partial f\|_\infty$. By replacing B'_1 with $\max(B'_1, 1)$ if necessary, we have

$$|u_s(jh; \tau; \omega) - u(jh; \tau; \omega)| \leq \Delta \leq B'_1 \Delta / \Lambda^{1/2}.$$

The proof is completed by utilizing these bounds in (3.3) and defining B_1 as $B_1 = \sqrt{8}B'_1$. \square

Lemma 7. Assume $0 < \Lambda \leq 1$. With probability greater than $1 - p$, $\|f_1 - f_2\|_K \leq \frac{B_1 \Delta^{1/2}}{\Lambda}$.

Proof. Follows from Lemmas 5, 6, and 4. Lemma 4 is applied with $r = \frac{B}{\Lambda^{1/2}}$, $\delta = \frac{B_1^2 \Delta}{\Lambda}$, and $\lambda = 2\Lambda$. The choice of r is justified by Lemma 5 and the choice of δ is justified by Lemma 6. To justify the choice of λ , note that $\mathcal{W}_1(f)$ and $\mathcal{W}_2(f)$ can both be written as $\mathcal{W}_i(f) = \mathcal{L}(f) + \Lambda \|f\|_K^2$ with \mathcal{L} a convex functional. The identity $\langle f + g, f + g \rangle_K = \langle f, f \rangle_K + \langle 2f, g \rangle_K + \langle g, g \rangle_K$ shows that $2f$ is the unique subgradient at f for $\Lambda \|f\|_K^2$. Thus, if $w \in \nabla \mathcal{W}_i(f)$ (the subgradient of \mathcal{W}_i is unique and may be obtained explicitly), we must have $\mathcal{W}_i(f + g) - \mathcal{W}_i(f) - \langle w, g \rangle_K \geq \Lambda \langle g, g \rangle_K$, justifying the choice of λ . \square

Lemma 8. Assume $0 \leq \Lambda \leq 1$. For $f \in \mathcal{H}_K$ and $\|f\|_K \leq \frac{B}{\Lambda^{1/2}}$, we have $|\mathcal{W}_2(f) - \mathcal{W}_3(f)| \leq \frac{B_1^2 h}{\Lambda}$.

Proof. We will argue as in Lemma 6 and assume that $\|f\|_\infty$, and $\|\partial f\|_\infty$ are bounded by $B'_1/\Lambda^{1/2}$.

Suppose $\alpha \in [0, 1]$. In the difference

$$\begin{aligned} & \frac{1}{h} \int_{kh}^{(k+1)h} (f(u(t; \tau; \omega)) - u(t + t_f; \omega))^2 dt \\ & - (1 - \alpha) (f(u(kh; \tau; \omega)) - u(kh + t_f; \omega))^2 \\ & - \alpha (f(u((k + 1)h; \tau; \omega)) - u((k + 1)h + t_f; \omega))^2, \end{aligned}$$

we may apply the mean value theorem to the integral and argue as in Lemma 6 to upper bound the difference by $(B'_1)^2 h/\Lambda$. The proof is completed by summing the differences from $k = 0$ to $k = Nn - 1$ and dividing by Nn . \square

Lemma 9. Assume $0 \leq \Lambda < 1$. Then $\|f_2 - f_3\| \leq \frac{B_1 h^{1/2}}{\Lambda}$.

Proof. Follows from Lemmas 5, 8, and 4. Lemma 4 is applied with $r = \frac{B}{\Lambda^{1/2}}$, $\delta = \frac{B_1^2 h}{\Lambda}$, and $\lambda = 2\Lambda$. The choices of r , δ , and Λ are justified using Lemmas 5 and 8 and an additional argument as in the proof of Lemma 7. \square

Choose $\epsilon > 0$. At this point, we specialize K to a kernel for which the universality theorem of Steinwart applies. For example, $K = K_\gamma$. We may then find $F_\epsilon \in \mathcal{H}_K$ such that $\|F_\epsilon - F\|_\infty \leq \epsilon$, where the ∞ -norm is over Y . In fact, we will need the difference $|F_\epsilon(x) - F(x)|$ to be bounded by ϵ only for $x \in X$. The larger compact space Y is needed to apply the universality theorem and for other RKHS arguments.

Lemma 10. *Let $\Lambda = \epsilon^2 / \|F_\epsilon\|_K^2 \leq 1$. If f_3 minimizes $\mathcal{W}_3(f)$, we have*

$$\frac{1}{T} \int_0^T (f_3(u(t; \tau; \omega)) - u(t + t_f; \omega))^2 dt \leq \Lambda \|F_\epsilon\|_K^2 + \epsilon^2 = 2\epsilon^2.$$

In addition, $\|f_3\|_K^2 \leq 2\|F_\epsilon\|_K^2$.

Proof. We have

$$\frac{1}{T} \int_0^T (f_3(u(t; \tau; \omega)) - u(t + t_f; \omega))^2 dt \leq \mathcal{W}_3(f_3),$$

$\mathcal{W}_3(f_3) \leq \mathcal{W}_3(F_\epsilon)$ because f_3 is the minimizer, and

$$\mathcal{W}_3(F_\epsilon) \leq \epsilon^2 + \Lambda \|F_\epsilon\|_K^2.$$

This last inequality uses $\int (F_\epsilon(u(t; \tau; \omega)) - u(t + t_f; \omega))^2 dt = \int (F_\epsilon(u(t; \tau; \omega)) - F(u(t; \tau; \omega)))^2 dt$. The proof of the first part of the lemma is completed by combining the inequalities. To prove the second part, we argue similarly after noting $\|f_3\|_K^2 \leq \mathcal{W}_3(F_\epsilon) / \Lambda$. □

Consider half-open boxes in \mathbb{R}^D of the form

$$A_{j_1, j_2, \dots, j_D} = \left[\frac{j_1}{2^\ell}, \frac{j_1 + 1}{2^\ell} \right) \times \dots \times \left[\frac{j_D}{2^\ell}, \frac{j_D + 1}{2^\ell} \right),$$

with $\ell \in \mathbb{Z}^+$ and $j_i \in \mathbb{Z}$. The whole of \mathbb{R}^D is a disjoint union of such boxes. Because X is compact, we can assume that $X \subset \cup_{j=1}^L A_j$, where the union is disjoint, each A_j is a half-open box of the form above, and $A_j \cap X \neq \emptyset$ for $1 \leq j \leq L$.

We will pick ℓ to be so large, that each box has a diameter that is bounded as follows:

$$\frac{\sqrt{D}}{2^\ell} < \frac{\delta}{4\sqrt{2}D^{1/2} \|\partial^2 K\|_{2,\infty}^{1/2} \|F_\epsilon\|_K}.$$

Here $\delta > 0$ is determined later, and $\|\partial^2 K\|_{2,\infty}$ is the ∞ -norm in the function space $C^2(Y \times Y)$. Lemma 10 tells us that $\|f_3\|_K \leq \sqrt{2}\|F_\epsilon\|_K$, and therefore (by part (c) of Theorem 1 of [30]) $\|\partial f_3\|_\infty \leq \sqrt{2}D^{1/2} \|\partial^2 K\|_{2,\infty}^{1/2} \|F_\epsilon\|_K$. As a consequence of our choice of ℓ , $x, y \in A_j$ implies that

$$|f_3(x) - f_3(y)| < \delta/4, \tag{3.4}$$

bounding the variation of f_3 within a single cell A_j . Because the exact predictor F is C^r , $r \geq 2$, and X is compact, we may also assert that

$$|F(x) - F(y)| < \delta/4 \tag{3.5}$$

for $x, y \in A_j$ by taking ℓ larger if necessary.

The next lemma is about taking a trajectory that is long enough that each of the sets A_j is sampled accurately. By assumption X is the support of μ . However, we may still have $\mu(A_j) = 0$ for some j . In the following lemma and later, it is assumed that all A_j with $\mu(A_j) = 0$ are eliminated from the list of boxes covering X .

Lemma 11. *Let χ_{A_j} denote the characteristic function of the set A_j . There exist $T^* > 0$ and a Borel measurable set*

$$S_{\epsilon, T^*} \subset X$$

such that $\omega \in S_{\epsilon, T^*}$ implies that for all $T \geq T^*$ and $j = 1, \dots, L$

$$\left| \frac{1}{T} \int_0^T \chi_{A_j}(u(t; \tau; \omega)) dt - \mu(A_j) \right| \leq \epsilon \mu(A_j).$$

and with $\mu(S_{\epsilon, T^*}) > 1 - \epsilon$.

Proof. To begin with, consider the set A_1 . By the ergodic theorem,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \chi_{A_1}(u(t; \tau; \omega)) dt = \mu(A_1)$$

for $\omega \in S \subset X$ with $\mu(S) = 1$. Let

$$A_{s, \epsilon} = \left\{ \omega \in X \left| \left| \frac{1}{T} \int_0^T \chi_{A_1}(u(t; \tau; \omega)) dt - \mu(A_1) \right| > \epsilon \mu(A_1) \text{ for some } T \geq s \right. \right\}.$$

The sets $A_{s, \epsilon}$ shrink with increasing s . Then the measure of $\bigcap_{s=1}^\infty A_{s, \epsilon}$ under μ is zero. Therefore, there exists $s_1 \in \mathbb{Z}^+$ such that $\mu(A_{s_1, \epsilon}) < \epsilon/L$.

We can find s_2, \dots, s_L similarly by considering the sets A_2, \dots, A_L . The lemma then holds with $T^* = \max(s_1, \dots, s_L)$. □

Lemma 12. *Suppose that $\omega \in S_{\epsilon, T^*}$, $T \geq T^*$, and $\Lambda = \epsilon^2 / \|F_\epsilon\|_K^2 \leq 1$. Suppose that f_3 minimizes $\mathcal{W}_3(f)$, which is defined using $u(t; \omega)$, T , and Λ . Then*

$$\mu \{ x \in X \mid |f_3(x) - F(x)| \geq \delta \} < \frac{8\epsilon^2}{\delta^2(1 - \epsilon)}.$$

Proof. Denote the set $\{x \in X \mid |f_3(x) - F(x)| \geq \delta\}$ by S_δ . Let J be the set of all $j = 1, \dots, L$ such that $|f_3(x) - F(x)| \geq \delta$ for some $x \in A_j$. Evidently, $S_\delta \subset \cup_{j \in J} A_j$, and it is sufficient to bound the measure of $\cup_{j \in J} A_j$.

By (3.4) and (3.5), if $|f_3(x) - F(x)| \geq \delta$ for some $x \in A_j$ then for any $y \in A_j$, we have

$$\begin{aligned} |f_3(y) - F(y)| &\geq |f_3(x) - F(x)| - |f_3(x) - f_3(y)| - |F(x) - F(y)| \\ &> \frac{\delta}{2}. \end{aligned} \quad (3.6)$$

For $\omega \in S_{\epsilon, T^*}$, we have

$$\begin{aligned} &\frac{1}{T} \int_0^T (f_3(u(t; \tau; \omega)) - u(t + t_f; \omega))^2 dt \\ &= \frac{1}{T} \int_0^T (f_3(u(t; \tau; \omega)) - F(u(t; \tau; \omega)))^2 dt \\ &\geq \frac{1}{T} \int_0^T (f_3(u(t; \tau; \omega)) - F(u(t; \tau; \omega)))^2 \sum_{j \in J} \chi_{A_j}(u(t; \tau; \omega)) dt \\ &= \frac{1}{T} \sum_{j \in J} \int_0^T (f_3(u(t; \tau; \omega)) - F(u(t; \tau; \omega)))^2 \chi_{A_j}(u(t; \tau; \omega)) dt \\ &\geq \frac{\delta^2}{4T} \sum_{j \in J} \int_0^T \chi_{A_j}(u(t; \tau; \omega)) dt \\ &\geq \frac{\delta^2}{4} \mu(\cup_{j \in J} A_j) (1 - \epsilon), \end{aligned}$$

where the first inequality holds because A_j are disjoint, the second inequality holds because $|f_3(y) - F(y)| > \delta/2$ follows from (3.6) for $y = u(t; \tau; \omega) \in A_j$ with $j \in J$, and the final inequality is a consequence of Lemma 11 and $\omega \in S_{\epsilon, T^*}$.

Applying Lemma 10, we get

$$\frac{\delta^2}{4} \mu \left(\bigcup_{j \in J} A_j \right) (1 - \epsilon) \leq 2\epsilon^2,$$

completing the proof of the lemma. \square

Lemma 13. Suppose $\omega \in S_{\epsilon, T^*}$ and that the signals $u(t; \omega)$ and $u_\eta(t; \omega)$ are used to define $\mathcal{W}_i(f)$, $i = 1, 2, 3$. Suppose that f_1 , f_2 , and f_3 minimize $\mathcal{W}_1(f)$, $\mathcal{W}_2(f)$, and $\mathcal{W}_3(f)$, respectively, with $T \geq T^*$ and $\Lambda = \epsilon^2 / \|F_\epsilon\|_K^2 \leq 1$. Then

$$\mu \left\{ x \in X \left| |f_1(x) - F(x)| > \delta + \frac{B_1 h^{1/2} + B_1 \Delta^{1/2}}{\Lambda} \right. \right\} < \frac{8\epsilon^2}{\delta^2(1 - \epsilon)}$$

with probability greater than $1 - p$.

Proof. Follows from Lemmas 7, 9, and 12. \square

The above lemma implies Theorem 2 with the choice of δ , n , and Δ specified above it.

4. Numerical illustrations

We compare three methods to compute an approximate predictor f . The first method is that of Müller et al. [19] given in (2.1). The second method is exactly the same but with the least squares regression function. The third method is the convergent algorithm given by (2.2) and (2.3).

When comparing the methods, we always used the same noisy data for all three methods. There can be some fluctuation due to the instance of noise that is added to the exact signal $\tilde{x}(t)$ as well as the segment of signal that is used. The effect of this fluctuation on comparison is eliminated by using the same noisy data in each case. In addition, reported results are averages over multiple datasets. For all three methods, the error in the approximate predictor is estimated by applying it to a noise-free stretch of the signal as in [19], which is standard because the object of each method is to approximate the exact predictor.

The first signal we use is the same as in [19], except for inevitable differences in instantiation. The Mackey-Glass equation

$$\frac{d\tilde{x}(t)}{dt} = -0.1\tilde{x}(t) + \frac{0.2\tilde{x}(t-D)}{1 + \tilde{x}(t-D)^{10}},$$

with $D = 17$, is solved with time step $\Delta t = 0.1$ and transients are eliminated to produce the exact signal $\tilde{x}(t)$. This signal will of course have rounding errors and discretization errors, but those are negligible compared to prediction errors. The standard deviation of the Mackey-Glass signal is about 0.23. An independent normally distributed quantity of mean zero is added at each point so that the ratio of the variance of the noise to that of the signal (0.23^2) is equal to the desired signal-to-noise ratio (SNR).

To confirm with [19], the Mackey-Glass signal was down-sampled so that $nh = 1$ and $n = 1$. The spline smoothing method would fare even better if we chose $h = .1$. The delay and the embedding dimension used for delay coordinates were $\tau = 6$ and $D = 6$, as in [19]. The size of the training set was $N = 1000$. For cross-validation, the $\gamma/2D$ parameter was varied over $\{0.1, 1.5, 10.0, 50.0, 100.0\}$, and the Λ parameter was varied over $\{10^{-8.5}, 10^{-8}, \dots, 10^{-0.5}\}$ for least squares with or without spline smoothing but over $\{10^{-10}, 10^{-6}, 10^{-2}, 10^2\}$ for the more expensive support vector regression. For support vector regression, the ϵ was varied over $\{0.01, 0.05, 0.25\}$. The phenomenon we will demonstrate is far more pronounced than the slight gains obtained using more extensive cross-validation. For support vector regression, we were able to reproduce the relevant results reported in [19].¹

Figure 4.1 demonstrates that (2.1) produces predictors that are corrupted by errors in the inputs or delay coordinates. The method with spline smoothing is more accurate and deteriorates less with increasing SNR. For the Mackey-Glass

¹The RMS error of 0.017 reported for $t_f = 1$ with SNR of 22.15% in [19] appears to be a consequence of an unusually favorable noise or signal. The typical RMS error is around 0.03. We eliminate the effect of unusual datasets by taking averages over multiple datasets.




FIG 4.1. Root mean square errors in the prediction of the Mackey-Glass signal with $t_f = 1$ as a function of the signal to noise ratio. The superiority of the method using smooth splines is evident.




FIG 4.2. Comparison of the 1-step least squares predictor (without spline smoothing) iterated t_f times with the t_f -step predictor (without spline smoothing). The latter is seen to be superior.

plots in Figures 4.1, 4.2, and 4.3, each point is an average over 480 independent datasets in the case of least squares with or without spline smoothing and over 48 data sets in the case of support vector regression. In all cases, using half as many datasets does not change the picture.

A $t_f = n_f \tau$ predictor can be obtained by iterating a τ -step predictor n_f times, and this strategy is sometimes used to save cost [19]. This is not a good idea as explained in [28] and as shown in Figure 4.2. An optimal predictor would need to roughly split the distance to the nearest training sample such that the component of the distance along unstable directions is small and with the component along stable directions allowed to be much larger. The balance between the two components depends upon t_f , and therefore, iterating a one-step predictor is not a good strategy.




FIG 4.3. The plot on the left uses SNR of 0.2 and the plot on the right uses 0.4. The method using smooth splines does better in all instances.

In Figure 4.1, we see that spline smoothing becomes more and more advantageous as noise increases. The situation in Figure 4.3 is a little different. When t_f is small, spline smoothing does help more for the noisier SNR of 0.4 compared to 0.2. However, for larger t_f , even though spline smoothing helps, it does not help more when the noise is higher. This could be because as t_f increases capturing the correct geometry of the predictor becomes more and more difficult, and this difficulty may be constraining the accuracy of the predictor.

The MacKey-Glass example is a delay-differential equation and does not come under the purview of our convergence theorem. The Lorenz example, $\dot{x} = 10(y - x)$, $\dot{y} = 28x - y - xz$, $\dot{z} = -8z/3 + xy$, is a dynamical system with a compact invariant set and comes under the purview of the convergence theorem. The Lorenz signal has a standard deviation of 7.9. For the Lorenz plots of Figure 4.4, each point is an average over 160 datasets each with $N = 1000$. The picture did not change even with many fewer datasets.

Figure 4.4 compares $h = .01$ and $h = .1$ for Lorenz. In both cases, the




FIG 4.4. The advantage of spline smoothing for Lorenz is much less on the right with $h = 0.1$ than on the left with $h = 0.01$.

embedding dimension is $d = 10$, the delay parameter is $\tau = 1$, and the lookahead is $t_f = h$. It may be seen that spline smoothing is less effective when $h = 0.1$ as compared to $h = 0.01$. A typical Lorenz oscillation has a period of about 0.75, and when $h = 0.1$ the resolution is too low causing too much discretization error. Smooth splines are less effective in reconstructing the noise-free signal if the grid on the time axis does not have sufficient resolution. The left half of Figure 4.4 shows an example where prediction using spline smoothing improves accuracy by a factor of 100 with $h = 0.01$.

5. Discussion

For the prediction of dynamical time series, we have shown that flows are quite different from maps. In the case of flows, the time series can be denoised by relying solely on the smoothness of the underlying flow. The predictor can be derived by applying kernel-based regression to the denoised signal. The result-

ing predictor converges to the exact predictor under conditions described by Theorem 2.

As far as dynamical time series are concerned, the parameter estimation problem [16, 17] is complementary to prediction. Much of the existing theory is for maps and with the assumption of rapid mixing. For flows, smooth splines or a similar technique may prove an effective method to denoise in the context of parameter estimation as well.

The convergence theorem given here does not give rates and is not uniform. Obtaining rates with uniformity over a class of flows will probably require rapid mixing assumptions as in the case of maps [11, 24]. Rapid mixing results for flows may be found in [2] for example.

With respect to rates and uniformity, there are two more issues that would need to be considered. First, convergence of smooth splines in the ∞ -norm must be proved with explicit bounds that depend only on the norm of the m -th derivative. A more significant point is that rates of convergence for a given lookahead t_f may not be the best direction. As pointed out in [28], the question of how large t_f can be given a signal of length T appears to have implications for the prediction algorithm and not just to its analysis. There is no evidence that existing algorithms including the one in this paper are capable of predicting as far into the future as an optimal algorithm should.

The smooth spline idea is primarily local and so are the optimality results of Stone [26]. Stone's algorithm for achievability is to find a local scale and to fit a polynomial using linear least squares within that local region. It is perhaps worth noting that the same idea has a dynamical analog. In its dynamical version [27], the noisy dynamical time series is embedded within Euclidean space using delay coordinates. The embedding will be necessarily noisy. However, the embedded manifold can be smoothed locally using linear techniques.

Acknowledgements

The authors thank the reviewer for a valuable report and the editors for their comments.

References

- [1] T.M. Adams and A.B. Nobel. On density estimation from ergodic processes. *The Annals of Probability*, 25:794–804, 1998. [MR1626511](#)
- [2] V. Araújo and I. Melbourne. Exponential decay of correlations for nonuniformly hyperbolic flows with a $c^{1+\alpha}$ stable foliation, including the classical Lorenz attractor. *Ann. Henri Poincaré*, 17:2975–3004, 2015. [MR3556513](#)
- [3] A. Christmann and I. Steinwart. Consistency and robustness of kernel-based regression in convex risk minimization. *Bernoulli*, 3:799–819, 2007. [MR2348751](#)
- [4] C. de Boor. *A Practical Guide to Splines*. Springer, New York, revised edition, 2001. [MR1900298](#)

- [5] P.P.B. Eggermont and V.N. LaRiccia. Uniform error bounds for smoothing splines. In *High Dimensional Probability*, volume 51 of *IMS Lecture Notes-Monograph Series*, pages 220–237. Institute of Mathematical Statistics, 2006. [MR2387772](#)
- [6] P.P.B. Eggermont and V.N. LaRiccia. *Maximum Penalized Likelihood Estimation*, volume II. Springer, New York, 2009. Springer Series in Statistics. [MR2817245](#)
- [7] I. Ekeland and R. Temam. *Convex Analysis and Variational Problems*. SIAM, Philadelphia, 1987. [MR1727362](#)
- [8] J. Doyne Farmer and J.J. Sidorowich. Predicting chaotic time series. *Physical Review Letters*, 59:845–848, 1987. [MR0902527](#)
- [9] L. Györfi, W. Härdle, P. Sarda, and P. Vieu. *Nonparametric Curve Estimation from Time Series*, volume 60 of *Lecture Notes in Statistics*. Springer, New York, 1989. [MR1027837](#)
- [10] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer, New York, 2002. [MR1920390](#)
- [11] H. Hang, Y. Feng, I. Steinwart, and J.A.K. Suykens. Learning theory estimates with observations from general stationary stochastic processes. *Neural Computation*, 28:2853–2889, 2016.
- [12] S. P. Lalley. Beneath the noise, chaos. *The Annals of Statistics*, 27:461–479, 1999. [MR1714721](#)
- [13] S. P. Lalley and A. B. Nobel. Denoising deterministic time series. *Dynamics of PDE*, 3:259–279, 2006. [MR2271730](#)
- [14] S.P. Lalley. Removing the noise from chaos plus noise. In A.I. Mees, editor, *Nonlinear Dynamics and Statistics*, pages 233–244. Birkhäuser, Boston, 2001. [MR1937487](#)
- [15] D. Mattera and S. Haykin. Support vector machines for dynamic reconstruction of a chaotic system. In B. Schölkopf, C.J.C. Burges, and A.J. Smola, editors, *Advances in Kernel Methods*, pages 211–242. MIT Press, MA, 1999.
- [16] K. McGoff, S. Mukherjee, and N. Pillai. Statistical inference for dynamical systems: A review. *Statistics Surveys*, 9:209–252, 2015. [MR3422438](#)
- [17] K. McGoff and A.B. Nobel. Empirical risk minimization and complexity of dynamical models. www.arxiv.org, 2016.
- [18] S. Mukherjee, E. Osuna, and F. Girosi. Nonlinear prediction of chaotic time series using a support vector machine. In *Neural Networks for Signal Processing VII—Proceedings of the 1997 IEEE Workshop*, pages 511–520. IEEE, 1997.
- [19] K.-S. Müller, A.J. Smola, G. Rätsch, B. Schölkopf, J. Kohlmorgen, and V.N. Vapnik. Using support vector machines for time series prediction. In B. Scholkopf, C.J.C. Burges, and S. Mika, editors, *Advances in Kernel Methods*, pages 243–253. MIT Press, 1998.
- [20] N.I. Sapankevych and R. Sankar. Time series prediction using support vector machines: a survey. *IEEE Computational Intelligence Magazine*, pages 24–38, May 2009.
- [21] T. Sauer, J.A. Yorke, and M. Casdagli. Embedology. *Journal of Statistical*

- Physics*, 65:579–616, 1991. [MR1137425](#)
- [22] B. Schölkopf and A.J. Smola. *Learning with Kernels*. MIT Press, MA, 2002.
- [23] I. Steinwart. On the influence of the kernel on the consistency of support vector machines. *Journal of Machine Learning Research*, 2:67–93, 2001. [MR1883281](#)
- [24] I. Steinwart and M. Anghel. Consistency of support vector machines for forecasting the evolution of an unknown ergodic dynamical system from observations with unknown noise. *The Annals of Statistics*, 37:841–875, 2009. [MR2502653](#)
- [25] I. Steinwart, D. Hush, and C. Scovel. Learning from dependent observations. *Journal of Multivariate Analysis*, 100:175–194, 2009. [MR2460486](#)
- [26] C. Stone. Optimal rates of convergence for nonparametric regression. *Annals of Statistics*, 10:1040–1063, 1982. [MR0673642](#)
- [27] K. Urbanowicz and J.A. Holyst. Noise-level estimation of time series using coarse-grained entropy. *Physical Review E*, 67:046218, 2003.
- [28] D. Viswanath, X. Liang, and K. Serkh. Metric entropy and the optimal prediction of chaotic signals. *SIAM Journal on Applied Dynamical Systems*, 12:1085–1113, 2013. [MR3071149](#)
- [29] G. Wahba. *Spline Models for Observational Data*. SIAM, Philadelphia, 1990. [MR1045442](#)
- [30] D.-X. Zhou. Derivative reproducing properties for kernel methods in learning theory. *Journal of Computational and Applied Mathematics*, 220:456–463, 2008. [MR2444183](#)