

Embeddings and Prediction of Dynamical Time Series

by
Raymundo Navarrete

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied and Interdisciplinary Mathematics)
in The University of Michigan
2018

Doctoral Committee:

Professor Divakar Viswanath, Co-Chair
Associate Professor Clayton Scott, Co-Chair
Associate Professor Lydia Bieri
Professor Ralph Spatzier
Associate Professor Shuheng Zhou

Raymundo Navarrete

raymundo@umich.edu

ORCID iD: 0000-0002-9317-0943

TABLE OF CONTENTS

LIST OF FIGURES	iii
ABSTRACT	iv
CHAPTER	
I. Introduction	1
1.1 Introduction	1
1.2 Contributions of this thesis	8
1.3 Overview of Chapter 2	10
1.4 Overview of Chapter 3	14
1.5 Overview of Chapter 4	19
II. Delay embeddings of periodic orbits using a fixed observation function	24
2.1 Introduction	24
2.2 Embedding periodic signals in \mathbb{R}^3	28
2.3 Embedding periodic orbits in \mathbb{R}^3	40
2.4 Discussion	50
III. Prevalence of delay embeddigs with a fixed observation function	52
3.1 Introduction	52
3.2 Transfer of volume	55
3.3 Rank lemmas	59
3.4 Review of Sauer et al [47]	63
3.5 Perturbing the dynamical system	75
3.6 The setting for injectivity and immersivity theorems	81
3.7 Proof of injectivity	83
3.8 Proof of immersivity	90
3.9 Discussion	91
IV. Prediction of dynamical time series using kernel based regression and smooth splines	93
4.1 Introduction	93
4.2 Prediction algorithm and statement of convergence theorem	96
4.3 Proof of convergence	103
4.4 Numerical illustrations	113
4.5 Discussion	116
BIBLIOGRAPHY	120

LIST OF FIGURES

Figure

2.1	A periodic signal (only a single period is shown) and its delay embedding in \mathbb{R}^3 with delay τ . The points a, b, c map to A, B, C with delay coordinates.	28
2.2	The points a_1 and a_2 , and likewise b_1, b_2 and c_1, c_2 , map to the same point in \mathbb{R}^3 under delay embedding with the delay shown. The fundamental period of this signal is half of what is shown. However, by modifying the signal in the box shown, its fundamental period becomes equal to the interval shown and the delay map still fails to be injective because c_1 and c_2 map to the same point in \mathbb{R}^3	29
2.3	An infinitely differentiable (bump) function (dashed line), which is zero outside (α, β) and 1 near the middle of that interval, subtracted from a constant value of $\frac{d\phi}{dt}$. If the amount subtracted is adjusted, the integral of $\frac{d\phi}{dt}$ over one full period becomes zero as shown.	31
2.4	The function $f(x) = x^3 \sin\left(\frac{1}{x}\right)$ and its derivative.	33
2.5	A periodic orbit of the classical Lorenz system and its x -coordinate as a function of time (over a single period). The periodic orbit shown is $A^{24}B$ in the nomenclature of [58].	40
2.6	A periodic orbit with a tube around it.	41
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.	114
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.	115
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.	115
4.4	The advantage of spline smoothing for Lorenz is much less on the right with $h = 0.1$ than on the left with $h = 0.01$	117

ABSTRACT

A dynamical time series is a sequence of real-valued observations of a dynamical system. Commonly in applications, the dynamical system of interest is unknown and only a dynamical time series is observed. Dynamical time series arise in ergodic theory in mathematics, nonlinear dynamics in physics, state space modeling in statistics, and control theory in engineering. We consider two common goals in the analysis of dynamical time series.

First, it is desirable to construct a faithful representation of the state of the dynamical system using only the time series. Delay-time coordinates are widely used for this purpose. Under certain conditions, the delay map taking the state of the dynamical system to its corresponding delay-time coordinates is generically an embedding of the state space. More precisely, current work shows that for a fixed dynamical system, delay-time coordinates result in embeddings of the state space generically with respect to the observation function. Motivated by common usage of delay-time coordinates, we consider the more difficult situation where the observation function is fixed and genericity is studied with respect to the dynamical system. We prove that delay-time coordinates result in embeddings of the state space for polynomial perturbations of the dynamical system with probability one over the perturbing coefficients.

Second, it is desirable to predict time series accurately. Prediction of dynamical time series with additive noise using kernel-based regression is consistent for certain

classes of discrete dynamical systems. These methods are effective at computing the expected value of the time series at a future time given the present delay-time coordinates. However, the present coordinates themselves are noisy, so these methods are only optimal when it is not possible to remove noise. We consider the problem of prediction for flows, and show that the use of smoothing splines to reduce noise before using kernel-based regression results in increased prediction accuracy. We prove that our method is consistent, converging to the exact predictor based on the noiseless time series, in the limit as sampling frequency and sampling time go to infinity.

CHAPTER I

Introduction

1.1 Introduction

A dynamical time series is a sequence of real numbers generated from observations of a dynamical system. Dynamical time series arise in ergodic theory in mathematics, nonlinear dynamics in physics, state space modeling in statistics, and control theory in engineering. Commonly in applications, only the dynamical time series is known, and the goal of analysis is to reconstruct the state space of the unknown dynamical system or to make predictions about future observations. In this thesis, we focus on these two aspects of the analysis of dynamical time series.

We consider both discrete-time and continuous-time dynamical systems. A discrete-time dynamical system (X, ϕ) consists of a compact manifold X containing the states of the system and a diffeomorphism $\phi : X \rightarrow X$ dictating the evolution of the system with time. If a system (X, ϕ) is in the state $\mathbf{x} \in X$ at some time $t \in \mathbb{Z}$, then this system is in the state $\phi(\mathbf{x})$ at time $t + 1$. We call X the state space and we call ϕ the dynamical map. The trajectory of a state $\mathbf{x}_0 \in X$ is the sequence $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots$, where $\mathbf{x}_{t+1} = \phi(\mathbf{x}_t)$ for $t = 0, 1, 2, \dots$. Hence, a trajectory $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots$ is completely determined by the map ϕ and the initial state $\mathbf{x}_0 \in X$. In a continuous-time dynamical system $(X, \{\phi^t\}_{t \geq 0})$, the map is replaced by a flow $\{\phi^t\}_{t \geq 0}$, which consist

of a family $\phi^t : X \rightarrow X$ of diffeomorphisms dictating the evolution of the system at all times $t \in \mathbb{R}$, satisfying $\phi^0 = \text{id}$ and $\phi^{t_1} \circ \phi^{t_2} = \phi^{t_1+t_2}$. The trajectory of a state $\mathbf{x}_0 \in X$ is the curve $\{\mathbf{x}(t)\}_{t \geq 0}$, where $\mathbf{x}(t) = \phi^t(\mathbf{x}_0)$. Of particular interest are dynamical systems described by differential equations: if $f : X \rightarrow TX$ is a smooth vector field on X , then we can define a flow $\{\phi^t\}_{t \geq 0}$ by $\phi^t(\mathbf{x}_0) = \mathbf{x}(t)$, where $\mathbf{x}(t)$ is the unique solution to the differential equation $\dot{\mathbf{x}} = f(\mathbf{x})$ with initial condition $\mathbf{x}(0) = \mathbf{x}_0$.

We consider situations in which the full state of the dynamical system cannot be observed, and only real-valued observations of the dynamical system are available. For a discrete-time dynamical system (X, ϕ) , given a trajectory $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots \in X$ of ϕ , we observe a dynamical time series $s_0, s_1, s_2, \dots \in \mathbb{R}$, where each s_t is an observation of \mathbf{x}_t , for each $t = 0, 1, 2, \dots$. More precisely, in the noiseless setting, we assume that the time series is given by an observation function $h : X \rightarrow \mathbb{R}$, so that $s_t = h(\mathbf{x}_t)$ for $t = 0, 1, 2, \dots$. In the observational noise setting, we assume that the observations are contaminated by additive noise, in which case $s_t = h(\mathbf{x}_t) + \epsilon_t$ for $t = 0, 1, 2, \dots$, where the ϵ_t are i.i.d. realizations of a mean-zero random variable. For continuous-time dynamical systems, we assume that observations of the system occur at equally spaced time intervals with frequency n . Given a trajectory $\{\mathbf{x}(t)\}_{t \geq 0}$ of a flow $\{\phi^t\}_{t \geq 0}$, the observed time series s_0, s_1, s_2, \dots is given by $s_t = h(\mathbf{x}(t/n))$, in the noiseless setting, or by $s_t = h(\mathbf{x}(t/n)) + \epsilon_t$, in the observational noise setting, for $t = 0, 1, 2, \dots$.

The first two chapters of this thesis deal with the problem of state space reconstruction from a time series using delay-time coordinates. Let us now overview the idea and aim of state space reconstruction and delay-time coordinates. Let (X, ϕ) be a dynamical system and let h be an observation function, and suppose that we

observe the noiseless time series $h(\mathbf{x}_0), h(\mathbf{x}_1), h(\mathbf{x}_2), \dots$, where $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ is a trajectory of ϕ . We assume that X is compact and that ϕ and h are C^1 . Suppose that $\xi \in C^1(X, \mathbb{R}^D)$ is an embedding, that is, ξ is diffeomorphic onto its image $\xi(X)$ (which will also be a compact manifold of the same dimension as X). Here D is a parameter called the embedding dimension. Then, observations of the trajectory resulting from the dynamical system $(\xi(X), \xi \circ \phi \circ \xi^{-1})$, the observation function $h \circ \xi^{-1}$, and the initial state $\xi(\mathbf{x}_0)$ result in the same time series as before. Hence, the basic goal of state space reconstruction is to produce a representation $\xi(\mathbf{x}_t) \in \mathbb{R}^D$ of each state \mathbf{x}_t based only on the observed time series, so that the map $\xi : X \rightarrow \mathbb{R}^D$ is an embedding of X . Delay-time coordinates are a popular and easy-to-apply method for state space reconstruction. Here, a representation of the state \mathbf{x}_t is constructed from the time series $h(\mathbf{x}_0), h(\mathbf{x}_1), h(\mathbf{x}_2), \dots, h(\mathbf{x}_t)$ in the form of the delay vector $(h(\mathbf{x}_t), h(\mathbf{x}_{t-1}), \dots, h(\mathbf{x}_{t-D+1}))$. The resulting map $\xi : X \rightarrow \mathbb{R}^D$, which we call the delay map, is

$$(1.1) \quad \xi(\mathbf{x}) := (h(\mathbf{x}), h(\phi^{-1}\mathbf{x}), \dots, h(\phi^{-D+1}\mathbf{x})).$$

Our focus is on the question of when do delay-time coordinates result in faithful representations of the state space, that is, when is the delay map an embedding.

Before overviewing current work on delay-time coordinates, we first connect this work to a classical theorem of Whitney. A main idea behind state space reconstruction is that in order to construct a good representation of a state $\mathbf{x} \in X$, all that is needed are sufficiently many independent observations of \mathbf{x} . This is a consequence of Whitney's embedding theorem: let X be a compact C^r manifold of dimension d , where $r \geq 1$; for some $D > 2d$, consider the space $C^r(X, \mathbb{R}^D)$ of C^r functions from X to \mathbb{R}^D (the space is given the standard strong topology); then it follows that the set

of C^r embeddings from X to \mathbb{R}^D is open and dense in $C^r(X, \mathbb{R}^D)$ [25]. As a consequence, a generic set of $D > 2d$ observation functions $h_1, h_2, \dots, h_D \in C^r(X, \mathbb{R})$ can be combined into a map $\mathbf{x} \rightarrow (h_1(\mathbf{x}), h_2(\mathbf{x}), \dots, h_D(\mathbf{x}))$ that is an embedding of X into \mathbb{R}^D . A subset of $C^r(X, \mathbb{R})$ is generic if it is equal to a countable intersection of open and dense subsets of $C^r(X, \mathbb{R})$. This result does not apply to delay-time coordinates, since the components of the delay map are not independent, but it suggests that a similar result can hold.

We now mention the main results on the question of when delay-time coordinates result in embeddings of the state space. The first results are due to Aeyels [2] and Takens [54]. Takens paper is particularly well known, and the result that delay maps are embeddings for $D > 2d$ and generic dynamical systems and observation functions is sometimes called Takens' theorem. Both authors rely on standard tools from differential topology, mainly parametric transversality. Their arguments do not consider some subtleties (see chapter 3). Sauer, Yorke and Casdagli [47] strengthened these results in a few ways. The method of proof is more careful and genericity is replaced by the stronger notion of prevalence.

Let us explain what is meant by prevalence. First, we consider perturbations $h + \sum_{\alpha} c_{\alpha} p_{\alpha}$ of the observation function h , where p_{α} are polynomials, c_{α} are coefficient, and the index α varies over a finite set. A property is said to be prevalent with respect to the observation function, if it holds with probability one with respect to the perturbing coefficients c_{α} . Likewise, a notion of prevalence can be defined by perturbing the vector field f or the map ϕ by polynomials. Sauer et al considered prevalence for a fixed map ϕ with perturbations applied to the observation function h . We consider the more difficult but also more natural situation where h is fixed (for example, as a projection π_1 to the first coordinate) but perturbations are applied to

f or ϕ . This situation is the most common in applications of delay-time coordinates, where the dynamical system and initial conditions are unknown, but the state space and observation function are known and fixed.

In chapter 2, we study delay maps with a fixed observation function, for flows and for the special case where the delay maps are restricted to a periodic trajectory. Let $r \geq 2$. Fix a C^r observation function h on \mathbb{R}^d . Consider the space $C^{r-1}(\mathbb{R}^d, \mathbb{R}^d)$ of vector fields in \mathbb{R}^d . Given a vector field $f \in C^{r-1}(\mathbb{R}^d, \mathbb{R}^d)$, there is a corresponding flow $\{\phi^t\}_{t \geq 0}$ defined by the differential equation $\dot{\mathbf{x}} = f(\mathbf{x})$. Suppose that the flow $\{\phi^t\}_{t \geq 0}$ has a hyperbolic periodic orbit $\mathbf{p} : [0, T) \rightarrow \mathbb{R}^d$ with period $T > 0$. Let $o : [0, T) \rightarrow \mathbb{R}$ be the periodic signal $o : [0, T) \rightarrow \mathbb{R}$ given by $o(t) := h(\mathbf{p}(t))$. For a delay parameter $\tau > 0$, define the delay map $o(\cdot; \tau) : [0, T) \rightarrow \mathbb{R}^3$ by $o(t; \tau) := (o(t), o(t - \tau), o(t - 2\tau))$. Here the embedding dimension is $D = 3$, which is the lowest possible. We ask if $o(\cdot; \tau)$ is an embedding of the circle $[0, T)$ for generic vector fields f . We must restrict the analysis to nearby perturbations $f' \in C^{r-1}(\mathbb{R}^d, \mathbb{R}^d)$ of the vector field f , for which we are guaranteed a corresponding perturbed hyperbolic periodic orbit \mathbf{p}' [41]. We show that, for suitable $\tau > 0$, there exist arbitrarily close perturbations f' of f for which the corresponding delay map o' is an embedding, and secondly, if f is such that o is an embedding, there exists an open neighborhood of f in $C^{r-1}(\mathbb{R}^d, \mathbb{R}^d)$ such that the corresponding delay maps are embeddings.

In chapter 3, we again study delay maps with a fixed observation function, but for maps and where the delay maps are restricted to $K \subset \mathbb{R}^d$, a compact ball centered at the origin. Let the projection onto the first component $\pi_1 : \mathbb{R}^d \rightarrow \mathbb{R}$ act as the fixed observation function. Let ϕ be a diffeomorphism on \mathbb{R}^d having finitely many periodic points of period less than $2D$ in K , where the embedding dimension D is to be specified. Let F_0 be the delay map corresponding to ϕ and π_1 . Consider perturbations

of ϕ given by $\phi(\mathbf{x}) + \mathbf{e}_1 \sum_{\alpha=1}^N c_\alpha p_\alpha(\mathbf{x})$, where $\mathbf{e}_1 = (1, 0, \dots, 0) \in \mathbb{R}^d$, p_1, p_2, \dots, p_N is a basis of the polynomials of degree at most $2D-1$, and c_1, c_2, \dots, c_N are perturbation coefficients. Let F_c be the delay map corresponding to the perturbation vector $c = (c_1, c_2, \dots, c_N)$. For a suitable $a > 0$ and $\|c\| \leq a$, we ask whether the delay map F_c is an embedding of K . We show that for $D \geq 4d+2$, the subset of $\|c\| \leq a$ for which the delay map is an embedding has full measure.

The final chapter of this thesis deals with the problem of prediction. Time series prediction consists of using the time series $\{s_0, s_1, \dots, s_T\}$ and predicting future values s_{T+1}, s_{T+2}, \dots . A common approach is to estimate a one-step predictor, that is, a function $f : \mathbb{R}^D \rightarrow \mathbb{R}$ that takes the delay vector $(s_t, s_{t-1}, \dots, s_{t-D+1})$ and returns a prediction to s_{t+1} . This method is motivated by an application of delay-time coordinates: if the time series is given by noiseless observations $s_t = h(\mathbf{x}_t)$ of a trajectory $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots$ of a dynamical system (X, ϕ) , where $h : X \rightarrow \mathbb{R}$ is an observation function, then the map $\xi(\mathbf{x}_t) = (s_t, s_{t-1}, \dots, s_{t-D+1})$ is generically an embedding of X and hence the function $F = h \circ \phi \circ \xi^{-1}$ satisfies

$$F(s_t, s_{t-1}, \dots, s_{t-D+1}) = s_{t+1}$$

for every t . In this noiseless setting, this one-step predictor $F : \xi(X) \rightarrow \mathbb{R}$ can be estimated by interpolating the data $\{(x_t, y_t)\}_t$, where $x_t = (s_t, s_{t-1}, \dots, s_{t-D+1})$ and $y_t = s_{t+1}$. We call F the exact one-step predictor. Methods for estimating F from a noiseless time series exist. For example, Nobel [38] gives a histogram-based algorithm for computing F .

We consider the observational noise model, where the time series is of the form $s_t = h(\mathbf{x}_t) + \epsilon_t$ for $t \geq 0, 1, 2, \dots$, where the ϵ_t are i.i.d samples of a mean-zero random variable. Steinwart and Anghel [50] analyzed a kernel-based algorithm for computing a one-step predictor from the noisy data $\{(x_t, y_t)\}_t$. They consider the minimizer f_T

of

$$(1.2) \quad \frac{1}{T-D-1} \sum_{i=D-1}^{T-1} (y_i - f(x_i))^2 + \Lambda_T \|f\|_{\mathcal{H}_\gamma}^2$$

over $f : \mathbb{R}^D \rightarrow \mathbb{R}$ in the reproducing kernel Hilbert space \mathcal{H}_γ corresponding to the Gaussian kernel $k(x, y) = \exp(-\gamma\|x - y\|^2)$, where $\gamma > 0$ and $\Lambda_T > 0$ are determined using cross validation [48]. They give conditions on the regularization parameter Λ_T to ensure that f_T converges with probability one to the regression function $G(x) := \mathbb{E}[\tilde{y} \mid \tilde{x} = x]$, where (\tilde{x}, \tilde{y}) is the distribution of (x, y) .

At first sight, it seems as if the regression function G is the best possible predictor for the observational noise model. Nonetheless, given that the data is presented in the form of a time series, it is possible to reduce noise before constructing the pairs $\{(x_t, y_t)\}_t$, and hence obtain a better estimate of the states at each point. In particular, if it is possible to eliminate all noise in the time series, then it is possible to estimate the exact one-step predictor F instead of G . We pursue this strategy for dynamical time series generated from flows.

Many time series of interests are given by measurements of continuous-time dynamical systems, such as systems modeled by ordinary differential equations. Lalley [31] suggested that noise removal for time series generated from flows is different than for maps. For general discrete dynamical systems, there is no notion of smoothness with respect to time, so all noise-removal algorithms must be based on recurrences. For continuous-time dynamical systems, a signal $h(\phi^t \mathbf{x}_0)$ will depend smoothly on time, whereas the observational noise will not. Let $\{\phi^t\}_t$ be a flow on X and consider a time series given by sampling with frequency n

$$(1.3) \quad s_t = h(\phi^{t/n}(\mathbf{x}_0)) + \epsilon_t$$

The time series is a noisy discretization of the signal $h(\phi^t(\mathbf{x}_0))$ with noisy samples at times $t = 0, 1/n, 2/n, \dots$, so we can estimate $h(\phi^t(\mathbf{x}_0))$ using any of the usual methods for one-dimensional smoothing. Due to their practical benefits, we consider smooth splines for this purpose. In particular, we use cubic spline smoothing [8] to compute the signal $u_n : [0, T]$ minimizing the functional

$$(1.4) \quad \frac{1}{nT+1} \sum_{i=0}^{nT} (s_t - u(i/n))^2 + \lambda_{nT} \int_0^T u''(t)^2 dt$$

over the Sobolov space $W^{2,2}[0, T]$. The regularization parameter $\lambda_{nT} > 0$ determines the balance between fit to the signal and smoothness of the spline function, and is determined using cross-validation.

In chapter 4, we propose using smooth splines (1.4) before learning a one-step predictor using kernel-based regression (1.2). We compare our approach to other kernel-based algorithms that try to directly estimate the regression function G and show that our approach can improve prediction accuracy significantly. We show that, for an ergodic flow $\{\phi^t\}_t$ on X preserving a compactly supported probability measure $\tilde{\mu}$ and initial state \mathbf{x} drawn from $\tilde{\mu}$, given the time series (1.3), the regularization parameters λ_{nT} and Λ_{nT} can be chosen so that the estimated one-step predictor converges to the exact one-step predictor F with probability one in the limit as $n \rightarrow \infty$ and $T \rightarrow \infty$.

1.2 Contributions of this thesis

1. Suppose that $\mathbf{p}(t)$ is a hyperbolic periodic solution of the differential equation $\frac{dx}{dt} = f(x)$, $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$. We prove that there exists an open set O of vector fields containing f (in C^r with $r \geq 2$) such that the delay map into \mathbb{R}^3 is an embedding of the circle for a set of vector fields dense in O (chapter 2).
2. The theory of delay embeddings originated with the work of Takens [54] and

Aeyels [2]. We argue that the proofs in both papers have significant gaps (chapter 3).

3. A major paper on delay embeddings is due to Sauer et al [47], who considered the situation where the dynamical system is fixed and only the observation function is perturbed. We point out a minor but significant error in that paper and correct it (chapter 3).
4. We analyze the considerably more difficult situation where the dynamical system is perturbed and the observation is fixed (as π_1 the projection to the first coordinate). For diffeomorphisms ϕ defined over K , where K is a compact ball in \mathbb{R}^d , we prove that the delay map with embedding dimension $D \geq 4d + 2$ is an embedding with probability one with respect to polynomial perturbations of ϕ (chapter 3).
5. As a part of that analysis, we introduce a new technique based on the concept of Lebesgue points for proving prevalence. Prevalence, which asserts a property to be true with probability one with respect to polynomial perturbations, is a more powerful concept than genericity.
6. The best methods for time series prediction are based on machine learning. We use smooth splines to derive a method that significantly improves prediction accuracy (chapter 4).
7. For predicting a dynamical time series with a lookahead equal to t_f , we show that a predictor derived assuming a lookahead of t_f from the beginning is superior to iterating a one-step predictor (chapter 4).
8. Prediction of dynamical time series with observational noise using support vector machines with Gaussian kernel has been shown to be consistent by Steinwart

and Anghel [50]. We prove that a similar predictor that uses a preliminary spline smoothing step converges to the exact predictor of the unknown noiseless time series in the limit as sampling time and frequency go to infinity (chapter 4).

1.3 Overview of Chapter 2

Consider a dynamical system defined by the differential equation $\frac{d\mathbf{x}}{dt} = f(\mathbf{x})$, where f is a C^{r-1} vector field on \mathbb{R}^d , where $r \geq 2$. Suppose that this system has a hyperbolic periodic solution $\mathbf{p} : [0, T) \rightarrow \mathbb{R}^d$ of period $T > 0$, which is guaranteed to be C^r . We always think of $[0, T)$ as a topological circle, that is, the interval $[0, T]$ with 0 and T identified. Consider the periodic signal $o : [0, T) \rightarrow \mathbb{R}$ given by $o(t) := \pi(\mathbf{p}(t))$, where $\pi \in C^r(\mathbb{R}^d, \mathbb{R})$ is an observation function. To simplify exposition, we restrict π to linear observation functions $\pi(\mathbf{x}) = \mathbf{a} \cdot \mathbf{x}$, where $\mathbf{a} \in \mathbb{R}^d$. For a delay $\tau > 0$, we define the delay map $o(\cdot; \tau) : [0, T) \rightarrow \mathbb{R}^3$ given by $o(t; \tau) := (o(t), o(t - \tau), o(t - 2\tau))$.

Let \mathcal{O}^r be the set of C^r periodic signals. More precisely, \mathcal{O}^r is the set of pairs (o, T) , where $o : [0, T) \rightarrow \mathbb{R}$ is a periodic C^r signal of period $T > 0$. When there is no room for confusion, we will simply write $o \in \mathcal{O}^r$ instead of $(o, T) \in \mathcal{O}^r$. We define the following metric d on \mathcal{O}^r

$$d(o, o') := \sup_{k=0, \dots, r} \sup_{0 \leq s < 1} |o^{(k)}(sT) - o'^{(k)}(sT')| + |T - T'|.$$

\mathcal{O}^r is given the topology generated by this metric. Similarly, let \mathcal{P}^r be the set of C^r periodic functions on \mathbb{R}^d . More precisely, \mathcal{P}^r is the set of pairs (\mathbf{p}, T) , where $\mathbf{p} : [0, T) \rightarrow \mathbb{R}^d$ is a periodic C^r function on \mathbb{R}^d of period $T > 0$. Again, we will simply write $\mathbf{p} \in \mathcal{P}^r$ instead of $(\mathbf{p}, T) \in \mathcal{P}^r$. We can also define a metric (which we also call d) on \mathcal{P}^r by

$$d(\mathbf{p}, \mathbf{p}') := \sup_{k=0, \dots, r} \sup_{0 \leq s < 1} \|\mathbf{p}^{(k)}(sT) - \mathbf{p}'^{(k)}(sT')\|_2 + |T - T'|$$

and give \mathcal{P}^r the topology generated by this metric. Finally, let \mathcal{F}^{r-1} be the set of C^{r-1} vector fields on $U \subset \mathbb{R}^d$. The set is endowed with the usual C^{r-1} Whitney topology. In this topology, two vector fields are close to each other if their values and first $r-1$ derivatives are close to each other at every point (in local coordinates).

The theorem we prove is the following. Consider a dynamical system $\frac{d\mathbf{x}}{dt} = f(\mathbf{x})$ with a hyperbolic periodic orbit \mathbf{p} of period T and corresponding periodic signal o . Then, there exists $\tau > 0$ an open neighborhood of f in \mathcal{F}^{r-1} such that for an open and dense subset of this neighborhood every vector field g admits a nearby hyperbolic periodic orbit \mathbf{p}' of period T' and a corresponding periodic signal o' such that the map $t \rightarrow o'(t; \tau)$ is an embedding of the circle $[0, T')$ into \mathbb{R}^3 .

Sketch of some of the proofs

The arguments of chapter 2 are divided into two sections, the first dealing with general periodic signals and their embeddings into \mathbb{R}^3 , and the second dealing with periodic orbits and their embeddings into \mathbb{R}^3 . In the next paragraphs, we give an overview of the proof of genericity of general periodic signals and make some comments on the proof of genericity of periodic orbits of flows.

We first sketch the proof of genericity of embeddings of general periodic signals. We begin with the argument that a generic delay map $t \rightarrow o(t; \tau)$ is locally an embedding. To begin, given a general periodic signal $o \in \mathcal{O}^r$, we show that o has at most finitely many critical points, or that o can be approximated arbitrarily well by a signal $o' \in \mathcal{O}^r$ of the same period and with finitely many critical points. More precisely, we show that if 0 is a regular value of \dot{o} , then o has finitely many critical points; otherwise, 0 is a critical value of \dot{o} , and we show how to construct an arbitrarily close perturbation o' of o for which \dot{o}' has 0 as a regular value. The last step uses Sard's theorem, which states that the set of critical values of a smooth

function between manifolds is a null set [17, 25]. Then, given a periodic signal $o \in \mathcal{O}^r$ with finitely many critical points, we show that the map ξ is locally an embedding. Since o has finitely many critical points, the set $[0, T)$ is broken into finitely many sub-intervals of strict monotonicity with minimum length $\mu > 0$. For $\tau < \mu/3$, we show that $\xi(t) \neq \xi(t')$ for $|t - t'| < \mu/3$.

Suppose now that $t \rightarrow o(t; \tau)$ is a local embedding, as described above. In order to find an arbitrarily small perturbation o' of o for which $t \rightarrow o'(t; \tau)$ is a (global) embedding, we make use of a finite-dimensional family of perturbations of o that is rich enough to perturb away points where $t \rightarrow o(t; \tau)$ fails to be injective. Explicitly, let $h = \tau/2$ and $0 \leq j \leq n := \text{floor}(T/h)$, and consider C^∞ bump functions $\lambda_j : [0, T) \rightarrow [0, 1]$ centered at jh and of width $2h = \tau$, such that $\lambda(t) = 1$ for $|t - jh| < h/2$ and $\lambda(t) = 0$ for $|t - jh| > h$. Hence, the family of perturbations o_ϵ of o defined by

$$o_\epsilon(t) = o(t) + \sum_{i=0}^n \epsilon_i \lambda_i(t),$$

has the property that for any $t \in [0, T)$, there is one j for which $\lambda_j(t) = 1$, and if $|t - t'| \geq \tau$, then $\lambda_j(t') = 0$. First, we show that the perturbations o_ϵ are also local embeddings for sufficiently small ϵ . Explicitly, we show that there exists $\epsilon' > 0$ such that for $\|\epsilon\| < \epsilon'$ and for $\tau < 12$, the corresponding delay maps $o_\epsilon(t; \tau)$ satisfy $o_\epsilon(t; \tau) \neq o_\epsilon(t'; \tau)$ for $|t - t'| < \mu/3$.

Suppose now that $o \in \mathcal{O}^r$ is such that $o(t; \tau)$ is a local embedding for $\tau < 12$, as just described. We find an arbitrarily small perturbation o' of o for which $t \rightarrow o'(t; \tau)$ is a (global) embedding. We find o' among the perturbations o_ϵ of o using results from transversality theory. For this, we consider the set

$$\mathcal{T} = \{(t_1, t_2) \mid |t_1 - t_2| > \mu/3, t_1, t_2 \in [0, T)\},$$

which is a submanifold of dimension 2 of the torus $[0, T) \times [0, T)$. The function

$f : \mathbb{R}^{n+1} \times \mathcal{T} \rightarrow \mathbb{R}^3 \times \mathbb{R}^3$ taking $(\epsilon, (t_1, t_2))$ to $(o_\epsilon(t_1; \tau), o_\epsilon(t_2; \tau))$ is shown to be transverse to the diagonal of $\mathbb{R}^3 \times \mathbb{R}^3$. By the parametric transversality theorem, there are ϵ arbitrarily small so that $(t_1, t_2) \rightarrow (o_\epsilon(t_1; \tau), o_\epsilon(t_2; \tau))$ is transversal to the diagonal of $\mathbb{R}^3 \times \mathbb{R}^3$, which implies that $o_\epsilon(t_1; \tau) \neq o_\epsilon(t_2; \tau)$ for $(t_1, t_2) \in \mathcal{T}$.

Finally, we show that if o is such that $o(t; \tau)$ is an embedding, then there is an open neighborhood of o in \mathcal{O}^r such that every corresponding delay map is an embedding. This is done with an application of the inverse function theorem. It follows that the set of periodic signals o for which $o(t; \tau)$ is an embedding for sufficiently small τ is open and dense in \mathcal{O}^r .

Let us now mention some of the arguments used in analyzing periodic orbits of dynamical systems and the genericity of their embeddings via delay coordinates. The openness part of the argument follows easily from standard results and the results above. Given a system $\frac{d\mathbf{x}}{dt} = f(\mathbf{x})$ with a hyperbolic periodic orbit \mathbf{p} , it is known that an open neighborhood $A \subset C^{r-1}(\mathbb{R}^d, \mathbb{R}^d)$ of f exists such that every $f' \in A$ contains a nearby hyperbolic periodic orbit \mathbf{p}' . If in addition the delay map $o(\cdot; \tau)$ of \mathbf{p} is an embedding for some τ , then the previous result about the openness of \mathcal{O}^r imply that the delay maps $o'(\cdot; \tau)$ will be embeddings for every $f' \in A$. It is the denseness part of the argument that requires more work. As already explained, given a periodic orbit $\mathbf{p} \in \mathcal{P}$ of a system $\frac{d\mathbf{x}}{dt} = f(\mathbf{x})$ with observed signal $o \in \mathcal{O}^r$, we can find an arbitrarily close signal $o' \in \mathcal{O}^r$ to o such that $o'(t; \tau)$ is an embedding. Hence, the difficulty is in showing that this can be used to construct an arbitrarily close vector field f' to f with periodic orbit $\mathbf{p}' \in \mathcal{P}$ arbitrarily close to \mathbf{p} such that o' is the corresponding observed signal.

1.4 Overview of Chapter 3

As in chapter 2, we are interested in the case where the observation function is fixed and the genericity of delay map embeddings is studied with respect to perturbations of the dynamical system. Whereas chapter 2 deals with flows, we now focus on maps. We expand the results of chapter 2 in two ways. First, we no longer restrict our analysis to periodic orbits of the dynamical system. Second, following Sauer et al [47], we replace genericity by prevalence. In exchange for these improvements, we must weaken the lower bound on the embedding dimension from chapter 2.

Consider a dynamical system defined by the diffeomorphism $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and an observation function given by $h : \mathbb{R}^d \rightarrow \mathbb{R}$. Both ϕ and h are assumed to be C^r for some $r \geq 2$. We simply write the iterates of a state $\mathbf{x} \in \mathbb{R}^d$ as $x_1 := \mathbf{x}$, $x_2 := \phi(x_1)$, $x_3 := \phi(x_2)$, and so on. The delay map $F_0 : \mathbb{R}^d \rightarrow \mathbb{R}^D$, where D is the embedding dimension, is given by

$$F_0(\mathbf{x}) = (h(x_1), h(x_2), \dots, h(x_D))^T.$$

We consider the question of whether perturbations of F_0 resulting from perturbations of ϕ are embeddings on $K \subset \mathbb{R}^d$, a closed ball centered at the origin. Since K is compact, F_0 is an embedding on K if and only if it is injective and the tangent map $T_{\mathbf{x}}F_0$ is injective at every $\mathbf{x} \in K$.

Sketch of some of the arguments

In what follows, $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a diffeomorphism and $h : \mathbb{R}^d \rightarrow \mathbb{R}$ is an observation function, both assumed to be C^r for some $r \geq 2$. The corresponding delay map is

given by

$$F_0(\mathbf{x}) = \begin{pmatrix} h(x_1) \\ \vdots \\ h(x_D) \end{pmatrix}.$$

The domain of F_0 is restricted to $K \subset \mathbb{R}^d$, a closed ball centered at the origin.

Given a multi-index $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d) \in \{0, 1, 2, \dots\}^d$, let $p_\alpha : \mathbb{R}^d \rightarrow \mathbb{R}$ be the monomial given by

$$p_\alpha(\mathbf{x}) := (\pi_1 \mathbf{x})^{\alpha_1} (\pi_2 \mathbf{x})^{\alpha_2} \cdots (\pi_d \mathbf{x})^{\alpha_d},$$

where $\pi_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is the projection onto the i -th coordinate, for $i = 1, 2, \dots, d$. The degree of p_α is given by $|\alpha| := \alpha_1 + \alpha_2 + \cdots + \alpha_d$. The set of polynomials \mathbb{P}_d^N in \mathbb{R}^d and of degree at most N consists of all functions of the form $\sum_{|\alpha| \leq N} c_\alpha p_\alpha$, where $c_\alpha \in \mathbb{R}$ for each index α . This is a finite-dimensional subspace of $C^r(\mathbb{R}^d, \mathbb{R})$, with dimension $D_d^N := \binom{N+d}{N}$. By enumerating the multi-indices $|\alpha| \leq N$, we can define $\mathbf{p} : \mathbb{R}^d \rightarrow \mathbb{R}^{D_d^N}$ by $\mathbf{p}(\mathbf{x}) = (p_\alpha(\mathbf{x}))_{|\alpha| \leq N}$, which we treat as a row vector. The elements of \mathbb{P}_d^N can then be expressed as $\mathbf{p}(\mathbf{x})c$, where $c := (c_\alpha)_{|\alpha| \leq N} \in \mathbb{R}^{D_d^N}$ is the perturbation (column) vector.

We first sketch the arguments used in proving the prevalence of delay map embeddings with regards to perturbations of the observation function h , as done in [47]. Here it is enough to consider perturbations by polynomials of degree at most $N = 2D - 1$. For a given perturbation vector $c \in \mathbb{R}^{D_d^N}$, the perturbed observation h_c is given by

$$(1.5) \quad h_c(\mathbf{x}) := h(\mathbf{x}) + \sum_{|\alpha| \leq 2D-1} c_\alpha p_\alpha(\mathbf{x}) = h(\mathbf{x}) + \mathbf{p}(\mathbf{x})c$$

and the corresponding delay map is

$$(1.6) \quad \begin{aligned} F_c(\mathbf{x}) &:= \begin{pmatrix} h_c(x_1) \\ \vdots \\ h_c(x_D) \end{pmatrix} = \begin{pmatrix} h(x_1) \\ \vdots \\ h(x_D) \end{pmatrix} + \sum_{|\alpha| \leq 2D-1} c_\alpha \begin{pmatrix} p_\alpha(x_1) \\ \vdots \\ p_\alpha(x_D) \end{pmatrix} \\ &= F_0(\mathbf{x}) + \begin{pmatrix} \mathbf{p}(x_1) \\ \vdots \\ \mathbf{p}(x_D) \end{pmatrix} c. \end{aligned}$$

Note that, for a fixed $\mathbf{x} \in K$, $F_c(\mathbf{x})$ is an affine function of c . In the proof, we assume that the perturbation coefficients are bounded by $\|c\| \leq a$ for some $a > 0$, and then take $a \rightarrow \infty$.

Let us first deal with the argument for injectivity of 1.6. Here we restrict ϕ to contain only finitely many periodic points of period less than $2D$. Define the set $\mathcal{K} := \{(\mathbf{x}, \mathbf{y}) \in K \times K \mid \mathbf{x} \neq \mathbf{y}\}$ and function $G_c : \mathcal{K} \rightarrow \mathbb{R}^D$ given by $G_c(\mathbf{x}, \mathbf{y}) = F_c(\mathbf{x}) - F_c(\mathbf{y})$. The goal is then to show that $0 \notin \text{range}(G_c)$ for a subset of perturbations c of full measure. Due to compactness of K and the ball $\|c\| \leq a$, the maps G_c are Lipschitz with a global Lipschitz constant L . We can write $G_c(\mathbf{x}, \mathbf{y}) = M_{\mathbf{x}, \mathbf{y}}c + b_{\mathbf{x}, \mathbf{y}}$ where

$$M_{\mathbf{x}, \mathbf{y}} = \begin{pmatrix} p(x_1) - p(y_1) \\ \vdots \\ p(x_D) - p(y_2) \end{pmatrix} \quad \text{and} \quad b_{\mathbf{x}, \mathbf{y}} := F_0(\mathbf{x}) - F_0(\mathbf{y}).$$

The matrix $M_{\mathbf{x}, \mathbf{y}}$ can be rewritten as

$$M_{\mathbf{x}, \mathbf{y}} = \begin{pmatrix} I_{D \times D} & -I_{D \times D} \end{pmatrix} \begin{pmatrix} \mathbf{p}(x_1) \\ \vdots \\ \mathbf{p}(x_D) \\ \mathbf{p}(y_1) \\ \vdots \\ \mathbf{p}(y_D) \end{pmatrix},$$

which has full rank D whenever all the iterates $x_1, \dots, x_D, y_1, \dots, y_D$ are distinct, as a consequence of the interpolation property of polynomials and that $N = 2D - 1$. Even if some of the iterates are the same, as long as neither \mathbf{x} or \mathbf{y} are periodic with period less than $2D$, we show that $M_{\mathbf{x}, \mathbf{y}}$ has rank D . Most pairs $(\mathbf{x}, \mathbf{y}) \in \mathcal{K}$ are of this form, since only a finite number of periodic points with period less than $2D$ exist. We denote the set of pairs in \mathcal{K} where containing one periodic point of period less than $2D$ by \mathcal{K}_1 . This set is dealt with separately. For the other points $\mathcal{K}_2 := \mathcal{K} \setminus \mathcal{K}_1$, set

$$\mathcal{K}_2(\delta) := \{(\mathbf{x}, \mathbf{y}) \in \mathcal{K}_2 \mid \|\mathbf{x} - \mathbf{y}\| \geq \delta, \text{dist}((\mathbf{x}, \mathbf{y}), \mathcal{K}_1) \geq \delta\}.$$

Due to compactness, for every $(\mathbf{x}, \mathbf{y}) \in \mathcal{K}_2(\delta)$, the matrix $M_{\mathbf{x}, \mathbf{y}}$ has full rank D with D -th singular values bounded below by some $\sigma_\delta > 0$. To show that G_ϵ is injective over $\mathcal{K}_2(\delta)$ with probability one, we use that $\mathcal{K}_2(\delta)$ can be covered by at most $C_K \epsilon^{-2d}$ ϵ -balls centered at $\mathcal{K}_2(\delta)$. For a given pair $(\mathbf{x}, \mathbf{y}) \in \mathcal{K}_2(\delta)$, the probability that $G_\epsilon(\mathbf{x}, \mathbf{y}) = 0$ is no more than the probability that $\|G_\epsilon(\mathbf{x}', \mathbf{y}')\| \leq L\epsilon$, where $(\mathbf{x}', \mathbf{y}')$ is the center of an ϵ -ball in the cover containing (\mathbf{x}, \mathbf{y}) , which can be shown to be bounded by $\frac{D_\alpha! L^D \epsilon^D}{\sigma_\delta^D a^D}$ using elementary arguments. Then the probability that $G_\epsilon(\mathbf{x}, \mathbf{y}) = 0$ for any $(\mathbf{x}, \mathbf{y}) \in \mathcal{K}_2(\delta)$ is bounded by $\frac{C_K}{\epsilon^{2d}} \frac{D_\alpha! L^D \epsilon^D}{\sigma_\delta^D a^D}$, which for $D > 2d$ goes to 0 as $\epsilon \rightarrow 0$. Since $\mathcal{K}_2 = \bigcup_{n=1}^{\infty} \mathcal{K}_2(1/n)$, the result can be extended to \mathcal{K}_2 and

finally to $\mathcal{K} = \mathcal{K}_1 \cup \mathcal{K}_2$.

The proof of immersivity of 1.6 uses similar arguments as those sketched above. To talk about immersions, we now require K to be a submanifold of \mathbb{R}^d , with T_1K as its unit tangent bundle. We consider the function $H_c : T_1K \rightarrow \mathbb{R}^d$ given by $H_c(\mathbf{x}, v) := (dF_c)_{\mathbf{x}}v$ and show that $H_c(\mathbf{x}, v) \neq 0$ for any $(\mathbf{x}, v) \in T_1K$ with probability one over perturbations c . Combining this and the results for injectivity, we obtain that for a fixed dynamical system satisfying the restrictions above, the resulting delay maps are embeddings with probability one over observation functions in the space of polynomials of degree at most $2D - 1$ (and any finite dimensional function space containing it).

We now make some comments regarding the proof of prevalence of delay map embeddings with respect to perturbations of the dynamical system. For simplicity of exposition, we set $h = \pi_1$, so we only need to consider perturbations of ϕ along its first coordinate. Here we make use of polynomials of degree at most $N = 4D + 1$. For a given $c \in \mathbb{R}^{D^N}$, the corresponding perturbed map ϕ_c is

$$\phi_c(\mathbf{x}) := \phi(\mathbf{x}) + \mathbf{e}_1 \sum_{|\alpha| < N} c_\alpha p_\alpha(\mathbf{x}) = \phi(\mathbf{x}) + \mathbf{e}_1 \mathbf{p}(\mathbf{x})c.$$

The trajectory starting at $\mathbf{x} \in \mathbb{R}^d$ under this new map is denoted by $\tilde{x}_1(c) = \mathbf{x}$, $\tilde{x}_2(c) = \phi_c(\tilde{x}_1)$, and so on (sometimes we drop the dependency on c). The delay map is

$$F_c(\mathbf{x}) = \begin{pmatrix} \pi_1(\tilde{x}_1(c)) \\ \vdots \\ \pi_1(\tilde{x}_D(c)) \end{pmatrix}.$$

Here $F_c(\mathbf{x})$ is no longer an affine function of c , for fixed \mathbf{x} . Using Taylor's theorem,

we obtain

$$F_c(\mathbf{x}) = F_0(\mathbf{x}) + \begin{pmatrix} \mathbf{0} \\ \mathbf{p}(x_1) \\ \mathbf{p}(x_2) + \pi_1 \rho_2(x_2, \mathbf{p}(x_1)) \\ \vdots \\ \mathbf{p}(x_{D-1}) + \pi_1 \rho_{D-1}(x_2, \dots, x_{D-1}, \mathbf{p}(x_1), \dots, \mathbf{p}(x_{D-1})) \end{pmatrix} c + \mathcal{O}(\|c\|^2)$$

where the ρ_i depend linearly on the $\mathbf{p}(x_j)$. This is similar to 1.6, except for the nonlinear terms in c .⁴

1.5 Overview of Chapter 4

For the purpose of predicting time series generated from noisy observations of continuous-time dynamical systems, we describe a new algorithm that can significantly improve prediction accuracy compared to other prediction methods. This algorithm makes use of the fact that flows exhibit smoothness with respect to time, meaning that it is possible to reduce noise in the time series before attempting to reconstruct the dynamics. We prove that this method can consistently learn the exact predictor function in the limit as the sampling frequency and length of measurements go to infinity.

Consider the flow $\phi : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ defined by the differential equation $\frac{d\mathbf{x}}{dt} = \mathcal{F}(\mathbf{x})$, where $\mathcal{F} \in C^r(\mathbb{R}^d, \mathbb{R}^d)$ for some $r \geq 2$. As customary, we write $\phi^t \mathbf{x}_0 := \phi(\mathbf{x}_0, t) = \mathbf{x}(t)$, where $\mathbf{x}(t)$ is the unique solution to $\frac{d\mathbf{x}}{dt} = \mathcal{F}(\mathbf{x})$ with initial condition $\mathbf{x}_0 \in \mathbb{R}^d$. We further assume that the flow ϕ preserves a probability measure μ on \mathbb{R}^d with compact support $\mathbf{X} \subset \mathbb{R}^d$ and that ϕ is ergodic with respect to μ . We will consider a trajectory $\{\phi^t \mathbf{x}_0\}_{t \in [0, \infty)}$ in \mathbf{X} , where the initial state $\mathbf{x}_0 \in \mathbf{X}$ is drawn from the measure μ .

Consider the observation function $\pi \in C^r(\mathbb{R}^d, \mathbb{R})$. For some embedding dimension D and delay-time $\tau > 0$, suppose that the delay-coordinate map $\xi : \mathbb{R}^d \rightarrow \mathbb{R}^D$ taking a state in \mathbf{X} to its corresponding delay-coordinate

$$\xi(\mathbf{x}) := (\pi(\mathbf{x}), \pi(\phi^{-\tau}\mathbf{x}), \dots, \pi(\phi^{-\tau(D-1)}\mathbf{x}))$$

is a C^r diffeomorphism. This embedding defines naturally a flow $\tilde{\phi} = \xi \circ \phi \circ \xi^{-1}$ on the image $\xi(\mathbb{R}^d) \subset \mathbb{R}^D$. The flow $\tilde{\phi}$ preserves the push-forward measure $\mu = \xi_*\boldsymbol{\mu}$ of $\boldsymbol{\mu}$, which has compact support $X = \xi(\mathbf{X})$, and is ergodic. For a fixed $t_f > 0$, we define the t_f -step predictor function $F = \pi \circ \phi^{t_f} \circ \xi^{-1}$, which for a given state $\mathbf{x} \in X$, takes the delay-coordinate $\xi(\mathbf{x})$ to the future observation $\pi(\phi^{t_f}\mathbf{x})$ after the system has evolved by time t_f . When there is not risk of confusion, we will simply write $x = \xi(\mathbf{x})$.

Suppose that a time series $\{s_t\}_{t=0}^\infty$ is generated from noisy observations of the trajectory $\{\phi^t\mathbf{x}_0\}_{t \in [0, \infty)}$ by sampling at regular time intervals with frequency n

$$s_t = \pi(\phi^{t/n}\mathbf{x}_0) + \varepsilon_t,$$

where the ε_t are i.i.d. samples from a noise distribution ε . Hence, if the system is observed up to a time $T \in \mathbb{Z}^+$, the times series obtained is $\{s_t\}_{t=0}^{nT}$. We assume that $\varepsilon = \mathcal{N}(0, \sigma^2)$, or more generally, that ε has variance σ^2 and a finite κ -moment for some $\kappa > 3$.

The predictor function F is normally estimated from the time series $\{s_t\}_{t=0}^{nT}$ as follows. For easiness of notation, we assume that the delay-time τ and the prediction time t_f are integers. Then, the delay-coordinate and future observation at time t become $x_t := (s_t, s_{t-n\tau}, \dots, s_{t-n(D-1)\tau})$ and $y_t := s_{t+nt_f}$. We can produce $N = n(T - t_f - \tau(D - 1))$ such pairs of points, and to simplify matters we can re-index the pairs to obtain the data pairs $\{(x_t, y_t)\}_{t=1}^N$. This data is then used to estimate

the regression function $G(x) = \mathbb{E}[\bar{Y}|\bar{X} = x]$, where (\bar{X}, \bar{Y}) is the distribution of the noisy delay-coordinate/future-observation pairs.

Algorithms following this protocol make use of the time series only to construct the delay-coordinate/future pairs. This is not optimal, as the regression function G , when restricted to X , will generally be different from the exact predictor F . In situations where it is possible to reduce noise in observations by other methods, such as in our set-up for flows, it is possible to obtain a regression function \tilde{F} that improves prediction accuracy and actually tries to approximate F .

The algorithm we study consists of two parts. First, the time series $\{s_t\}_{t=0}^{nT}$ is smoothed out to reduce noise. We choose the method of cubic splines for this purpose, since it uses the same methodology as the kernel-based methods that we use in the second part. In this step, we find the signal $\tilde{s} : [0, T] \rightarrow \mathbb{R}$ that minimizes the functional

$$\frac{1}{nT+1} \sum_{t=0}^{nT} (s_t - f(t/n))^2 + \lambda_{nT} \int_0^T f''(t)^2 dt$$

among functions $f \in W^{2,2}[0, T]$, the Sobolev space of twice-differentiable functions on $[0, T]$. Here, $\lambda_{nT} > 0$ is a regularization parameter balancing fit to the time series data and smoothness of the spline function \tilde{s} . In practice, the parameter is determined using cross-validation.

In the second part of our algorithm, the smoothed out signal \tilde{s} is used instead of s to learn the predictor function. As before, delay-coordinate/future pairs $\{(\tilde{x}_t, \tilde{y}_t)\}_{t=1}^N$ are constructed from \tilde{s} . The predictor \tilde{F} is obtained by minimizing the functional

$$\frac{1}{N} \sum_{i=1}^N (\tilde{y}_i - f(\tilde{x}_1))^2 + \Lambda_{nT} \|f\|_{\mathcal{H}_\gamma}^2$$

over functions $f \in \mathcal{H}_\gamma$, where \mathcal{H}_γ is the reproducing kernel Hilbert space corresponding to the Gaussian kernel $k(x, y) = \exp(-\gamma\|x - y\|_2^2)$, for some fixed $\gamma > 0$. In

practice, the bandwidth parameter $\gamma > 0$ and the regularization parameter $\Lambda_{nT} > 0$ are determined using cross-validation.

In the analysis of consistency, we set $\lambda_{nT} = \left(\frac{\log(nT)}{nT}\right)^{4/5}$, for which a theorem of Eggermont and LaRiccia [10, 11] assures that for a fixed $\Delta > 0$, the probability $p = p(n, T, \Delta, \mathbf{x})$ that $\sup_{t \in [0, T]} |u(t) - \pi(\phi^t \mathbf{x})| > \Delta$ goes to 0 and $n \rightarrow \infty$. In the analysis of consistency, we set $\Lambda_{nT} = \epsilon^2 / \|F_\epsilon\|_{\mathcal{H}_\gamma}^2$, where $\epsilon > 0$ and $F_\epsilon \in \mathcal{H}_\gamma$ is chosen so that $\|F_\epsilon - F\|_\infty < \epsilon$. This can be done by the universality of Gaussian kernels [49].

The RKHS \mathcal{H}_γ corresponding to the Gaussian kernel $k(x, y) = \exp(-\gamma\|x - y\|_2^2)$, where $x, y \in \mathbb{R}^D$ and $\gamma > 0$ is a fixed parameter, is large enough to guarantee that the predictor function F can be approximated arbitrarily close by functions in \mathcal{H}_γ . Steinwart [49] has shown that the Gaussian kernel is universal, meaning that continuous functions defined over compact sets can be approximated arbitrarily well. In particular, since the predictor function F is continuous, for any $\epsilon > 0$, there exists some $F_\epsilon \in \mathcal{H}_\gamma$ such that $\|F - F_\epsilon\|_\infty < \epsilon$.

For a fixed $\epsilon > 0$, choose a function $F_\epsilon \in \mathcal{H}_\gamma$ such that $\|F_\epsilon - F\|_\infty < \epsilon$ in a compact domain containing the invariant set X . Set $\Lambda = \epsilon^2 / \|F_\epsilon\|_{\mathcal{H}_\gamma}^2$. Divide X using boxes of dimension 2^{-l} , where $l > 0$ is large enough to ensure that F_ϵ varies by less than $\sqrt{\epsilon}/2$. Then T is chosen so that all boxes are adequately sampled (with respect to the ergodic measure). The bound Δ is chosen small enough so that $B_1 \Delta^{1/2} / \Lambda < \epsilon^{1/2}$. This determines how large n is. With this setup, we can show that the resulting predictor $F_{n,T}$ satisfies

$$\mu \{x \in X \mid |F_{n,T}(x) - F(x)| > 3\sqrt{\epsilon}\} < 8\epsilon / (1 - \epsilon)$$

for initial points $\mathbf{x} \in X$ of measure greater than $1 - \epsilon$ and with probability $1 - p(n, T, \Delta, \mathbf{x})$, where $p \rightarrow 0$ as $\epsilon \rightarrow 0$. Hence, as $n, T \rightarrow \infty$, our estimated predictor

$F_{n,T}$ converges to the exact predictor F .

CHAPTER II

Delay embeddings of periodic orbits using a fixed observation function

2.1 Introduction

Suppose a physical system is described by the differential equation $\frac{dx}{dt} = f(x)$, where $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$. Often the state vector x is unobservable in its entirety, and that is especially true if d is large. Thus, reconstructing the flow from observations is not straightforward. The technique of delay coordinates makes it possible to look at a single scalar observation and reconstruct the dynamics. We denote the scalar that is observed by πx . The observation function π could be a projection to a single coordinate, for example, when the velocity of a fluid flow is recorded at a single point and in a single direction. It could be some other linear function of x . More generally, the observation function πx could be nonlinear.

If $\phi_t(x)$ is the time- t flow map, the idea behind delay coordinates [39, 47, 54] is to use the delay vector

$$\xi(x; \tau, n) = (\pi x, \pi \phi_{-\tau}(x), \dots, \pi \phi_{-(n-1)\tau}(x)),$$

which is observable, as a surrogate for the point x in phase space. For a suitable choice of delay τ and embedding dimension n , delay coordinates yield a faithful representation of the phase space in a sense we will explain. Delay coordinates have

been employed in many applications [3, 53]. Current theory for delay coordinates [47] applies perturbations to the observation function π . We consider the situation where the observation function is fixed as a linear projection and only the dynamical system $\frac{dx}{dt} = f(x)$ is perturbed.

Packard et al [39] demonstrated that coordinate vectors such as $(\pi\phi_t(x), \frac{d}{dt}\pi\phi_t(x))$ give good representations of strange attractors. They noted that delay coordinate vectors would be equivalent to coordinate vectors formed using derivatives of the observed quantity.

A mathematical analysis of delay coordinates was undertaken in a famous paper by Takens [54] and independently by Aeyels [2]. In particular, Takens considered when $x \rightarrow \xi(x; \tau, n)$ is an embedding. Suppose M is a manifold of dimension m , $A \subset M$ a submanifold of M of dimension d , and $f : M \rightarrow N$ a continuous map from M to the manifold N . The restriction $f|_A$ is an embedding of A in N if the tangent map df has full rank at every point of A , $f|_A$ is injective, and $f|_A$ maps open sets in A to open sets in its range in the subspace topology [17, 25]. For the definition to make sense, the manifolds and f must be at least C^1 . More generally, the manifolds M, N and the map f may be assumed to be C^r with $r \geq 1$ or with $r = \infty$. Takens concluded that delay coordinates yield an embedding of compact manifolds without boundary if $n \geq 2m + 1$, for *generic* observation functions π and *generic* vector fields f . A property is *generic* in the C^r topology if it holds for functions f or π belonging to a countable intersection of open and dense sets [41]. Because the C^r spaces are Baire spaces [25], a countable intersection of open and dense sets is dense as well as uncountable.

The paper by Sauer et al [47] marked a major advance in the theory of delay coordinates. The approach to embedding theorems outlined by Takens relied on

parametric transversality. Parametric transversality arguments typically have a local part and a global part, and the transition from local arguments to a global theorem is made using partitions of unity [25].

Sauer et al [47] sidestepped transversality theory almost entirely. Unlike in transversality theory, there is no explicitly local part in the arguments of Sauer et al [47]. The local part of the argument comes down to a verification of Lipschitz continuity. The set being embedded is only assumed to have finite box counting dimension. The arguments are mostly probabilistic and the globalization step relies only on the finiteness of the box counting dimension. The only real analogy to differential topology appears to be to the proof of Sard's theorem [25], which too is proved using probabilistic arguments. Sauer et al prove prevalence [26], which goes beyond genericity. A property is prevalent with respect to the observation function π , if the property holds when any given π is replaced by $\pi + \sum_{\alpha \in I_\alpha} c_\alpha p_\alpha$, with p_α being monomials indexed by the finite set I_α , for almost every choice of the coefficients c_α .

The embedding theorem of Sauer et al [47] fixes the dynamical system and allows only the observation function π to be perturbed. The statements of genericity and prevalence are with regard to π , not the original dynamical system. If consideration is restricted to subsets A of box counting dimension d , Sauer et al only require $n > 2d$. Thus, we could even have $n < m$.

As mentioned, we investigate embedding theorems in which the observation function is fixed. For example, π could be fixed as a linear projection that extracts some component of the state vector. We allow perturbations of the dynamical system only.

The motivation for considering such embedding theorems is as follows. First, on purely aesthetic grounds, it appears desirable to have an embedding theory that depends upon the dynamics and not the observation function. Second, in many

applications the observation function is fixed, whereas the dynamical system itself is parametrized [3, 6, 16, 42, 53]. If π extracts a single component at a single point in the velocity field of a fluid, it is more pertinent to make the embedding theory depend upon the dynamics rather than upon the observation function.

Aeyels [2] stated that delay coordinates are injective for generic flows and a fixed observation function. In the context of applications, stronger theorems would be desirable as argued by Sauer et al [47]. First, an open and dense set can have arbitrarily small measure implying that prevalence, which is stronger than genericity, is a more appropriate concept. Second, the dynamics may be confined to an attractor of dimension much smaller than that of the state vector of the flow. In such a situation, we would like the embedding dimension to be determined by the dimension of the attractor and not the dimension of the state vector of the flow.

In this chapter, we consider the second of these two directions. Obtaining an embedding dimension that depends on the dimension of the attractor and not the flow introduces new difficulties when the observation is fixed and the flow is parametrized. Current proofs [47, 54] rely on perturbing the observation function to produce an embedding. When the observation function is fixed, the additional step of propagating perturbations to the flow to the observed delay coordinates will need to be handled. We need to understand how perturbing the flow perturbs the invariant set or attractor, which is assumed to persist, and how the perturbations to the invariant set or attractor propagate to delay coordinates. When the flow is fixed and the observation function is perturbed, the attractor to be embedded, which depends only upon the flow, is unchanged by the perturbations. In contrast, when the observation function is fixed and the flow is perturbed, the set to be embedded is altered by the perturbations.



Figure 2.1: A periodic signal (only a single period is shown) and its delay embedding in \mathbb{R}^3 with delay τ . The points a, b, c map to A, B, C with delay coordinates.

To get a handle on such difficulties, we limit ourselves to hyperbolic periodic orbits and prove that they embed generically in \mathbb{R}^3 . The techniques we use are those of transversality theory. Although periodic orbits are only a special case, they are an important special case and arise frequently in applications, for example [5, 15].

To conclude this introduction, we mention some other extensions of delay coordinate embedding theory. Embedding theory has been considered for endomorphisms [55] as well as delay differential equations [9], for continuous but not necessarily smooth observation functions [18, 19], and in concert with Kalman filtering [22]. The concept of determining modes and points in fluid mechanics and PDE is related to embedding theory [28, 42, 43]. Delay coordinates have been used for noise reduction [44, 57]. The embedding theory of Sauer et al [47] has been generalized to PDE by Robinson [42, 43]. The current embedding theory for PDE also relies on perturbing the observation function.

2.2 Embedding periodic signals in \mathbb{R}^3

In the next section, we consider periodic solutions of differential equations. In this section, we begin by considering periodic signals. A periodic signal is any function

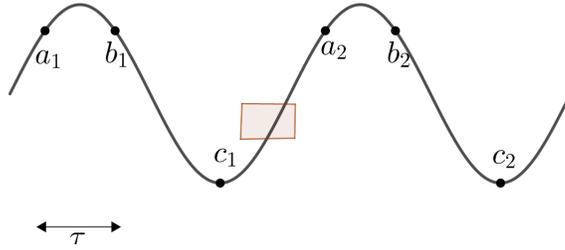


Figure 2.2: The points a_1 and a_2 , and likewise b_1, b_2 and c_1, c_2 , map to the same point in \mathbb{R}^3 under delay embedding with the delay shown. The fundamental period of this signal is half of what is shown. However, by modifying the signal in the box shown, its fundamental period becomes equal to the interval shown and the delay map still fails to be injective because c_1 and c_2 map to the same point in \mathbb{R}^3 .

$o : \mathbb{R} \rightarrow \mathbb{R}$ with a period $T > 0$. Figure 2.1 shows a periodic signal and its delay embedding in \mathbb{R}^3 .

To make the definition of periodic signals more precise, let \mathcal{O}^r be the set of C^r functions $o : [0, T] \rightarrow \mathbb{R}$ with period $T > 0$. Periodicity requires r derivatives of $o(t)$ to match at $t = 0$ and $t = T$. The domain of functions in \mathcal{O}^r , which we will write as $[0, T)$ for signals o of period T , is compact and homeomorphic to S^1 . More precisely, the domain is the identification space obtained by identifying 0 and T in $[0, T]$. For convenience, we shall refer to it as $[0, T)$, with the understanding that when we refer to an interval (α, β) it can wrap around. The elements of \mathcal{O}^r will be referred to as periodic signals. Even if $o \in \mathcal{O}^r$ is constant, it must be equipped with a period $T > 0$, and if T is chosen differently, we get a different element of \mathcal{O}^r .

For the periodic signal shown in Figure 2.1, the map $t \rightarrow (o(t), o(t - \tau), o(t - 2\tau))$ for $0 \leq t < T$ results in an embedding of the circle $[0, T)$ maps to a distinct point in \mathbb{R}^3 so that the delay map is injective. The delay is also immersive because a small movement along the periodic signal maps to a small and nonzero movement in the embedding space \mathbb{R}^3 . Because the delay map is both injective and immersive, it is an embedding.

Figure 2.2 shows a situation in which the delay map is not injective. This example is in fact the same as in Figure 2.1 but the period is taken to be double of what it is in Figure 2.1. As a result, points which are separated by the fundamental period map to the same point in \mathbb{R}^3 . As shown in Figure 2.2, the signal may be modified so that the fundamental interval is not repeated and the delay map still fails to be injective. Later in this section, we will prove that signals whose delay maps embed the circle in \mathbb{R}^3 are more typical.

Local argument for periodic signals

If $r \in \mathbb{Z}^+$ and $o, o' \in \mathcal{O}^r$ are two periodic signals, define

$$(2.1) \quad d_r(o, o') = \sup_{k=0, \dots, r} \sup_{0 \leq s < 1} |o^{(k)}(sT) - o'^{(k)}(sT')| + |T - T'|.$$

The C^r topology on \mathcal{O}^r is defined by this metric. The \mathcal{O}^r norm of a periodic signal is $\|o\|_r = \sup_{k=0, \dots, r} \sup_{0 \leq t < T} |o^{(k)}(t)|$. By our definition, \mathcal{O}^r is not a vector space because signals with different periods cannot be added. However, signals of a fixed period are a vector space and $\|\cdot\|_r$ is a norm over it. The C^∞ topology is the union of C^r topologies over $r \in \mathbb{Z}^+$ as explained in [25]. For concepts and results of differentiable topology, such as critical points, regular values, and Sard's theorem, our main reference is Hirsch [25]. The same topics are discussed from a dynamical point of view in [40, 41].

Figure 2.2 shows a signal which does not embed the circle in \mathbb{R}^3 under delay mapping. However, it is clear from observation that points that are nearby such as a_1 and b_1 map to distinct points in \mathbb{R}^3 . In fact, quite generally, if the number of critical points in $[0, T)$ is finite, nearby points in the signal will map to distinct points in \mathbb{R}^3 , as we later prove. We begin by considering whether any periodic signal may be perturbed slightly so that it has only finitely many critical points.

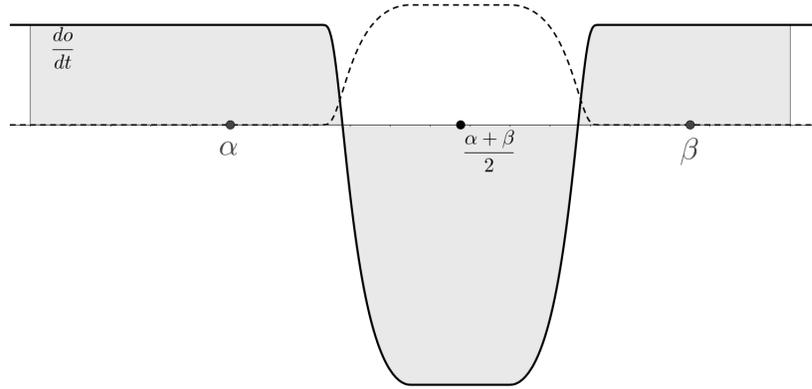


Figure 2.3: An infinitely differentiable (bump) function (dashed line), which is zero outside (α, β) and 1 near the middle of that interval, subtracted from a constant value of $\frac{do}{dt}$. If the amount subtracted is adjusted, the integral of $\frac{do}{dt}$ over one full period becomes zero as shown.

Lemma II.1. *Let $o \in \mathcal{O}^r$, $r \geq 2$, be a periodic signal of period $T > 0$. If 0 is a regular value of do/dt , then the periodic signal $o(t)$ has finitely many critical points in $[0, T)$.*

Proof. Suppose $do/dt = 0$ at infinitely many points on the compact circle $[0, T)$. Let $p \in [0, T)$ be an accumulation point of the set of zeros. Then $d^2o(p)/dt^2 = 0$ and $do(p)/dt = 0$ implying that 0 is not a regular value of do/dt . \square

The following lemma generates a periodic signal of period T whose derivative is $\frac{do}{dt} = \epsilon$ everywhere except over a given interval (α, β) . Any function whose derivative is $\frac{do}{dt} = \epsilon$, $\epsilon \neq 0$, everywhere cannot be periodic. Therefore, the proof of the lemma comes down to modifying the derivative carefully in the interval (α, β) .

Lemma II.2. *Given $(\alpha, \beta) \subset [0, T)$ and $\delta > 0$, for all sufficiently small ϵ there exists an infinitely differentiable periodic signal o of period T such that $do(t)/dt = \epsilon$ for $t \notin (\alpha, \beta)$ and $|do(t)/dt| < \delta$ for $t \in (\alpha, \beta)$. In addition, for $r \in \mathbb{Z}^+$, $\|o\|_r \rightarrow 0$ as $\epsilon \rightarrow 0$.*

Proof. Let $\lambda(x)$ be an infinitely differentiable bump function with $\lambda(x) \in [0, 1]$ for $x \in [0, 1]$, $\lambda(x) = 1$ for $x \in [1/4, 3/4]$, and $\lambda(x) = 0$ for $x \in [0, 1/8]$ and $x \in [7/8, 1]$. If $\int_0^1 \lambda(x) dx = c$ then $1/2 < c < 1$. The bump function $\lambda(x)$ is used to modify do/dt in the interval (α, β) .

Define $do(t)/dt = \epsilon$ for $t \notin (\alpha, \beta)$ and more generally

$$\frac{do(t)}{dt} = \epsilon - k\lambda((t - \alpha)/(\beta - \alpha))$$

for $t \in [0, T)$. The idea behind the construction is shown in Figure 2.3: if the bump function is shifted to the interval (α, β) and a suitable multiple is subtracted, $\frac{do}{dt}$ may then be integrated to obtain a periodic function.

More precisely, it follows that $\int_0^T (do(t)/dt) dt = \epsilon T - k(\beta - \alpha)c$. The integral is zero if $k = \epsilon T/(\beta - \alpha)c$. For ϵ small, k is small as well. We may obtain $o(t)$ by integrating $do(t)/dt$, with $\|o\|_r$ proportional to ϵ . \square

The following lemma proves that any sufficiently smooth periodic signal can be perturbed to a nearby periodic signal with finitely many critical points.

Lemma II.3. *If $o' \in \mathcal{O}^r$, $r \geq 2$, is a periodic signal, there exists another periodic signal o of the same period with $d_r(o, o')$ arbitrarily small and such that o has only finitely many critical points (including local maxima and minima) and 0 is a regular value of do/dt .*

Proof. If $o'(t)$ is constant we can perturb to $\epsilon \sin(tT/2\pi)$ for arbitrarily small ϵ and verify the theorem. We will assume that o' is not constant.

Consider $\frac{do'}{dt}(t)$ as a map from the circle $[0, T')$ to \mathbb{R} . If 0 is a regular value of this map, we are done by Lemma II.1.

If not, there exists a regular value ϵ of do'/dt arbitrarily close to 0 by Sard's theorem (here $r \geq 2$ is needed). Suppose we look at $do'(t)/dt - \epsilon$. This function has

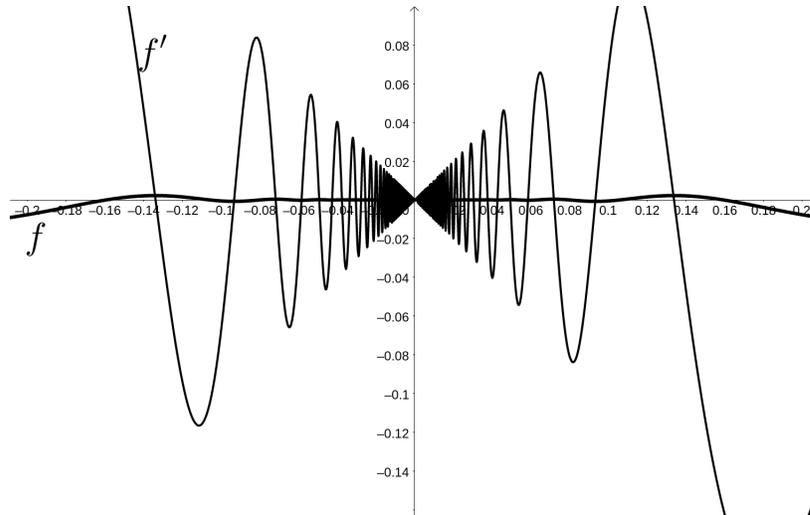


Figure 2.4: The function $f(x) = x^3 \sin\left(\frac{1}{x}\right)$ and its derivative.

a regular value at 0. However, the corresponding perturbation of o' is $o'(t) - t\epsilon$ and is not periodic.

Because $o'(t)$ is not constant, there exists an interval (α, β) in the circle $[0, T)$ over which $do'(t)/dt$ is nonzero. Without loss of generality, we assume $do'(t)/dt > \delta > 0$ in the interval (α, β) (consider $-o'(t)$ for the case where the derivative is negative). Using Lemma II.2, we may find a periodic signal $p(t)$ such that $dp/dt = \epsilon$ for $t \notin (\alpha, \beta)$ and $|dp/dt| < \delta$ for $t \in (\alpha, \beta)$. Set $o(t) = o'(t) - p(t)$ to obtain a periodic signal with 0 being a regular value of do/dt to complete the proof. \square

Remark. Lemma II.1 is evidently true if we only assume the second derivative of the periodic signal $o(t)$ to exist and not necessarily continuous. In fact, Lemma II.3 is also true under the same weaker assumption because, in one dimension, Sard's theorem requires only the existence of the derivative (see Exercise 1 of Section 3.1 of [25]).

The proof of Lemma II.3 may be illustrated using Figure 2.4. The figure shows a part of the graph of $f(x) = x^3 \sin(1/x)$ and its derivative $f'(x)$. It is evident that the critical points of f' , where $f''(x) = 0$, accumulate at the origin. In fact, a small

perturbation cannot eliminate the accumulation of critical points because $f(x)$ does not have a second derivative at $x = 0$. However, if $f(x) = x^5 \sin(1/x)$, a function whose second derivative looks like the derivative show in Figure 2.4, Sard's theorem may be used to obtain a small perturbation such that 0 is a regular value of the derivative of the perturbed function.

If o is a periodic signal with finitely many critical points, then its circular domain $[0, T)$ may be decomposed into finitely many intervals with local minima and maxima at either end. Let μ denote the minimum width among such intervals. Because $o(t)$ is monotonic in each interval, we refer to each such interval as the minimum interval of strict monotonicity. If the delay is τ , we denote the point $(o(t), o(t - \tau), o(t - 2\tau))$ by $o(t; \tau)$.

Lemma II.4. *If $0 < |t_1 - t_2| \leq \mu/3$, where μ is the minimum interval of strict monotonicity, and if the delay τ satisfies $0 < \tau \leq \mu/3$, then $o(t_1; \tau) \neq o(t_2; \tau)$. If 0 is a regular value of $\frac{do(t)}{dt}$, we also have $\frac{do(t; \tau)}{dt} \neq 0$ for all $t \in [0, T)$.*

Proof. Because $|t_1 - t_2| \leq \mu/3$, t_1 and t_2 lie in either the same interval of strict monotonicity of the periodic signal $o(t)$ or in neighboring intervals. If they lie in the same interval, we must have either $o(t_1) < o(t_2)$ or $o(t_2) < o(t_1)$ proving the lemma.

If t_1 and t_2 lie in neighboring intervals, we may assume $t_1 < t_2$ without loss of generality. If $o(t_1) \neq o(t_2)$, there is nothing to prove. So we assume $o(t_1) = o(t_2)$ in addition. Again without loss of generality, we assume that $o(t)$ first increases and then decreases as t increases from t_1 to t_2 .

With these assumptions, t_1 and $t_1 - \tau$ must lie in the same interval of monotonicity because $\tau \leq \mu/3$, and therefore $o(t_1 - \tau) < o(t_1)$. Further $t_2 - \tau \in (t_1 - \tau, t_2)$ and the unique minimum of $o(t)$ for $t \in [t_1 - \tau, t_2]$ is attained when $t = t_1 - \tau$. Therefore $o(t_1 - \tau) < o(t_2 - \tau)$, and we once again have $o(t_1; \tau) \neq o(t_2; \tau)$.

For the claim about $\frac{do(t;\tau)}{dt} \neq 0$, we note that $\frac{do}{dt}$ cannot equal zero at both t and $t - \tau$, because $\tau < \mu$. \square

With Lemma II.4, the local argument for embedding periodic signals is partly complete. Globalizing the argument will involve additional perturbations, which we now define.

Let λ be a C^∞ bump function with $\lambda(x) = 1$ for $|x| \leq 1/2$, $\lambda(x) = 0$ for $|x| \geq 1$, and $\lambda(x) \in [0, 1]$ for all $x \in \mathbb{R}$. Let $h = \tau/2$ and $j \in \mathbb{Z}$. Define

$$\lambda_j(t) = \lambda\left(\frac{t - jh}{h}\right)$$

for $j = 0, 1, \dots, n$ and $n = \lfloor T/h \rfloor$. We interpret t modulo T and regard $\lambda_j(t)$ as a periodic signal with the circular domain $[0, T)$: a pulse of period T and width h centered at jh which is equal to 1 for $|t - jh| \leq h/2$. We now consider the perturbation

$$(2.2) \quad o_\epsilon(t) = o(t) + \epsilon_0 \lambda_0(t) + \epsilon_1 \lambda_1(t) + \dots + \epsilon_n \lambda_n(t),$$

where $\epsilon = (\epsilon_0, \dots, \epsilon_n) \in \mathbb{R}^{n+1}$. For any $t_0 \in [0, T)$, there exists a bump function $\lambda_j(t)$ with $0 \leq j \leq n$ such that $\lambda_j(t_0) = 1$ and therefore $\lambda_j(t) = 0$ if $|t - t_0| \geq \tau = 2h$.

Before we turn to the global argument, we must prove that the local structure asserted by Lemma II.4 is preserved when o is perturbed to o_ϵ as in (2.2). The lemma below guarantees $o_\epsilon(t_1; \tau) \neq o_\epsilon(t_2; \tau)$ for $|t_1 - t_2| \leq 3\tau$. The bound 3τ ensures that $o_\epsilon(t_1; \tau) = o_\epsilon(t_2; \tau)$ can happen only when the intervals $[t_1 - 2\tau, t_1]$ and $[t_2 - 2\tau, t_2]$ do not overlap.

Lemma II.5. *Let $o \in \mathcal{O}^r$, $r \geq 2$, be a periodic signal defined over the domain $[0, T)$ and with minimum interval of strict monotonicity equal to μ . Assume that 0 is a regular value of do/dt . There exists ϵ_0 such that if $\|\epsilon\| \leq \epsilon_0$, then for the perturbation*

defined by (2.2) and delay τ satisfying $0 < \tau < \mu/12$, we have $o_\epsilon(t_1; \tau) \neq o_\epsilon(t_2; \tau)$ for all (t_1, t_2) with $|t_1 - t_2| \leq 3\tau$. In addition, 0 remains a regular value of $\frac{do_\epsilon}{dt}$.

Proof. By assumption the periodic signal $o(t)$ has finitely many critical points. Let $t_1 < t_2 < \dots < t_k$ be the critical points in the circular interval $[0, T)$; at these points and only at these, we have $do/dt = 0$. Since 0 is a regular value of do/dt , we have $\frac{d^2o(t_j)}{dt^2} \neq 0$ for $j = 1, \dots, k$.

In the circle $[0, T)$, choose compact intervals $K_i = [t_i - \delta, t_i + \delta]$, $i = 1, \dots, k$, such that $\delta < \mu/4$ and $\frac{d^2o(t)}{dt^2} \neq 0$ for any $t \in K_i$. By continuity in the perturbing parameters ϵ_i , for sufficiently small $\|\epsilon\|$ the perturbed periodic signal (2.2) also has nonzero second derivative on $\cup K_i$.

Define the interval K'_i to be $[t_i + \delta/2, t_{i+1} - \delta/2]$ (K'_k wraps around the circle). Each K'_k is an interval of strict monotonicity. By compactness, $|do/dt|$ attains a minimum strictly greater than 0 over $\cup K'_i$. Again by continuity, any perturbation of the form (2.2) with $\|\epsilon\|$ sufficiently small also has nonzero derivative over $\cup K'_i$.

Thus, for $\|\epsilon\|$ sufficiently small, K'_i remain intervals of strict monotonicity for the perturbed periodic signal, and each K_i can contain at most one critical point of the perturbed periodic signal. The minimum interval of strict monotonicity is at least $\mu - \delta \geq 3\mu/4$. We now apply Lemma II.4 to infer that $0 < \tau \leq \mu/4$ implies $o_\epsilon(t_1; \tau) \neq o_\epsilon(t_2; \tau)$ for $0 < |t_1 - t_2| \leq \mu/4$. We limit τ to the interval $(0, \mu/12)$ to complete the proof. \square

Global argument for periodic signals

The global argument relies on the parametric transversality theorem [25, 41].

Lemma II.6. *Let $o \in \mathcal{O}^r$, $r \geq 2$, be a periodic signal defined over the circle $[0, T)$.*

There exists an arbitrarily small perturbation of the periodic signal o to o' , with the

same period, and a delay $\tau > 0$, such that $t \rightarrow o'(t; \tau)$ is an embedding, with 0 a regular value of do'/dt .

Proof. By Lemma II.3, we may make an initial perturbation to o if necessary and assume that o has finitely many critical points, that 0 is a regular value of do/dt , and that $\mu > 0$ is the minimum width of an interval of strict monotonicity.

Now consider perturbations of o to o_ϵ of the form (2.2). By Lemma II.5, we may assume $o_\epsilon(t_1; \tau) \neq o_\epsilon(t_2; \tau)$ for $t_1 \neq t_2$ and $|t_1 - t_2| \leq 3\tau$ for $\tau < \mu/12$, provided $\|\epsilon\|$ is sufficiently small.

Consider the set

$$\mathcal{T} = \left\{ (t_1, t_2) \mid |t_1 - t_2| > 3\tau, t_1 \in [0, T), t_2 \in [0, T) \right\},$$

where $[0, T)$ is interpreted as the circle, as before. For the applicability of the parametric transversality theorem later in the proof, it is important to note that \mathcal{T} is a manifold of dimension 2 *without* a boundary.

Consider $(o_\epsilon(t_1; \tau), o_\epsilon(t_2; \tau))$ as a function from the domain $\{(\epsilon_1, \dots, \epsilon_n)\} \times \mathcal{T}$ to $\mathbb{R}^6 = \mathbb{R}^3 \times \mathbb{R}^3$. We will now verify that this function is transverse to the diagonal in $\mathbb{R}^3 \times \mathbb{R}^3$. If $o_\epsilon(t_1; \tau) \neq o_\epsilon(t_2; \tau)$ there is nothing to prove. Suppose $o_\epsilon(t_1; \tau) = o_\epsilon(t_2; \tau)$ and consider the point in \mathbb{R}^6 given by

$$(o_\epsilon(t_1), o_\epsilon(t_1 - \tau), o_\epsilon(t_1 - 2\tau), o_\epsilon(t_2), o_\epsilon(t_2 - \tau), o_\epsilon(t_2 - 2\tau))$$

The intervals $[t_1 - 2\tau, t_1]$ and $[t_2 - 2\tau, t_2]$ are disjoint because $|t_1 - t_2| > 3\tau$. By construction, there exist $i_1, i_2, i_3, i_4, i_5, i_6$ such that $\lambda_{i_1}, \lambda_{i_2}, \lambda_{i_3}, \lambda_{i_4}, \lambda_{i_5}, \lambda_{i_6}$ are each equal to 1 at exactly one of the six points $t_1, t_1 - \tau, t_1 - 2\tau, t_2, t_2 - \tau, t_2 - 2\tau$ and zero at the others. If the tangent direction in the domain is taken to perturb ϵ_{i_j} for $j \in \{1, \dots, 6\}$, it maps to a perturbation of the j -th coordinate in \mathbb{R}^6 , more precisely

the elementary vector \mathbf{e}_j . Therefore, the tangent map is surjective and transversality is verified.

By the parametric transversality theorem [Hirsch, Chapter 3, Theorem 2.7], we may choose ϵ arbitrarily small such that $(o_\epsilon(t_1; \tau), o_\epsilon(t_2; \tau))$ considered as a function from \mathcal{T} to \mathbb{R}^6 is transverse to the diagonal of $\mathbb{R}^3 \times \mathbb{R}^3$. Since \mathcal{T} is of dimension 2, that can only happen if $o_\epsilon(t_1; \tau) \neq o_\epsilon(t_2; \tau)$ for $(t_1, t_2) \in \mathcal{T}$.

To complete the proof, we only need to check the smoothness/dimension condition in the parametric transversality theorem. The dimension of \mathcal{T} is 2 and the codimension of the diagonal in \mathbb{R}^6 is 3. Thus, it is sufficient if the map from $\{(\epsilon_1, \dots, \epsilon_n)\} \times \mathcal{T}$ to \mathbb{R}^6 is C^1 which it is. \square

Lemma II.7. *Let $o \in \mathcal{O}^r$, $r \geq 2$, be a periodic signal such that $t \rightarrow o(t; \tau)$ is an embedding of the circle $[0, T)$ in \mathbb{R}^3 for delay $\tau > 0$. There exists $\epsilon_0 > 0$ such that $d_r(o, o') < \epsilon_0$ and $T = T'$ (perturbation has same period) imply that $t \rightarrow o'(t; \tau)$ is also an embedding of the circle $[0, T)$.*

Proof. By the inverse function theorem (see [25, Appendix]), there exists $\epsilon_0 > 0$ such that for every $\tilde{t} \in [0, T)$ there exists a neighborhood of \tilde{t} over which $t \rightarrow o'(t; \tau)$ is an injection if $d_r(o', o) < \epsilon_0$ and $T = T'$. Using a Lebesgue- δ argument we may assume that $o'(t_1; \tau) \neq o'(t_2; \tau)$ for $0 < |t_1 - t_2| < \epsilon_0$, making ϵ_0 smaller if necessary.

Although arguments like the one above are common in differential topology, we state the version of the inverse function theorem invoked for clarity. The version used is as follows. Suppose f is a C^r map from U , an open subset of \mathbb{R}^m to V , an open subset of \mathbb{R}^n with $m < n$. Suppose $f(x) = y$ and that the tangent map $\frac{\partial f}{\partial x}$ is injective at x . Then there exists a neighborhood \mathcal{N} of f in the weak C^r topology ($r \geq 1$), a neighborhood U' of x , V' of y , and W' of $0 \in \mathbb{R}^{n-m}$, such that for every $g \in \mathcal{N}$ there exists a diffeomorphism $G : V' \rightarrow U' \times W'$ with G^{-1} restricted to $U' \times 0$

coinciding with g . This theorem is applied with $m = 1$ and $n = 3$.

The rest of the proof is a standard compactness argument. Let

$$\min_{|t_1 - t_2| \geq \epsilon_0} |o(t_1; \tau) - o(t_2; \tau)| = \delta > 0,$$

where the minimum exists because of compactness and is greater than 0 because $t \rightarrow o(t; \tau)$ is an embedding. By continuity, the minimum must be positive for o' sufficiently close to o . Similarly, immersivity of o' sufficiently close to o is a direct consequence of compactness of the circle. Thus, $t \rightarrow o'(t; \tau)$ is also an embedding. \square

Theorem II.8. *The set of periodic signals $o \in \mathcal{O}^r$, $r \geq 2$, for which there exists a delay $\tau > 0$ such that $t \rightarrow o(t; \tau)$ is an embedding of the circle $[0, T)$ in \mathbb{R}^3 is open and dense in \mathcal{O}^r .*

Proof. By Lemma II.6, there exists an arbitrarily small perturbation to o' such that $t \rightarrow o'(t; \tau)$ is an embedding for $0 < \tau < \tau_0$ and with 0 a regular value of do'/dt . Thus the set of periodic signals with a delay embedding and with 0 a regular value of do/dt is dense. We only have to prove that the set is open.

Given periodic signal o with $t \rightarrow o(t; \tau)$ an embedding, Lemma II.7 shows that $t \rightarrow o'(t; \tau)$ remains an embedding for $d_r(o, o')$ sufficiently small if $T = T'$. If $T \neq T'$, we may still apply Lemma II.7, by defining $o''(t) = o'(tT'/T)$ which is a periodic signal of period T . If $d_r(o, o') \rightarrow 0$, then $d_r(o, o'') \rightarrow 0$. Finally, $t \rightarrow o''(t; \tau)$ is an embedding implies that $t \rightarrow o'(t; \tilde{\tau})$ is an embedding with $\tilde{\tau} = \tau T'/T$. \square

Theorem II.9. *Suppose that $o \in \mathcal{O}^r$, $r \geq 2$, and that $t \rightarrow o(t; \tau)$ is an embedding of the circle for some delay $\tau > 0$. Then $t \rightarrow o(t; \tau')$ remains an embedding if τ' is close enough to τ .*

Proof. The arguments used in Lemma II.7 and Theorem II.8 apply with little change. \square

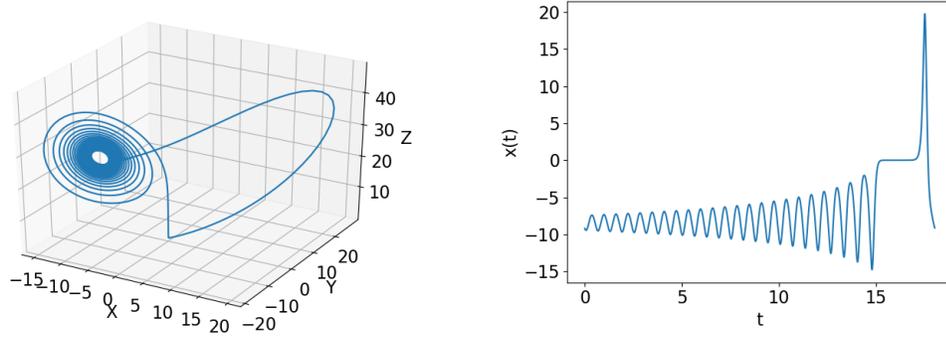


Figure 2.5: A periodic orbit of the classical Lorenz system and its x -coordinate as a function of time (over a single period). The periodic orbit shown is $A^{24}B$ in the nomenclature of [58].

2.3 Embedding periodic orbits in \mathbb{R}^3

Figure 2.5 shows a periodic orbit of the classical Lorenz system given by $dx/dt = 10(y - x)$, $dy/dt = -y - xz + 28x$, $dz/dt = -8z/3 + xy$.¹ The signal extracted from that orbit is nearly flat for a significant duration when the origin is approached.

In this section, we will prove that “typical” periodic orbits (in a sense that will be made precise) yield signals that result in embeddings of the circle. The following proposition proves that an embedding using delay coordinates persists when the vector field is perturbed slightly. It is the easier half of the argument.

Proposition II.10. *Let $\frac{dx}{dt} = f(x)$, where $x \in \mathbb{R}^d$, $f : U \rightarrow \mathbb{R}^d$, and U an open subset of \mathbb{R}^d , be a dynamical system with f a C^{r-1} vector field, $r \geq 2$. Let $\mathbf{p} : [0, T) \rightarrow U$ be a hyperbolic periodic solution of period $T > 0$. Let $\mathbf{a} \in \mathbb{R}^d$ and $\mathbf{a} \neq 0$. Assume that $t \rightarrow (\mathbf{a} \cdot \mathbf{p}(t), \mathbf{a} \cdot \mathbf{p}(t - \tau), \mathbf{a} \cdot \mathbf{p}(t - 2\tau))$ be an embedding of the circle $[0, T)$ in \mathbb{R}^3 . There exists an open neighborhood of f in the C^{r-1} topology such that for each g in that neighborhood, there exists a C^r -close hyperbolic periodic solution $\mathbf{p}'(t)$ of period T' of $\frac{dx}{dt} = g(x)$ and a τ' close to τ such that $t \rightarrow (\mathbf{a} \cdot \mathbf{p}'(t), \mathbf{a} \cdot \mathbf{p}'(t - \tau'), \mathbf{a} \cdot \mathbf{p}'(t - 2\tau'))$*

¹The periodic orbit of Figure 2.5 in [58] could not be computed using the techniques of [58]. It was computed some years later using an initial guess that was constructed from the periodic orbit $A^{25}B^{25}$.

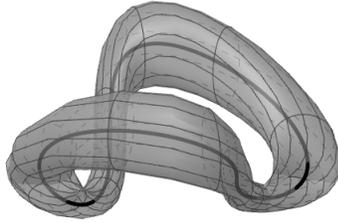


Figure 2.6: A periodic orbit with a tube around it.

is an embedding of the circle $[0, T')$ in \mathbb{R}^3 .

Proof. The fact that a hyperbolic periodic solution such as \mathbf{p} perturbs to a nearby hyperbolic solution \mathbf{p}' in a small enough open neighborhood of f is a standard result [42, Chapter 5]. If the signal $o(t) = \mathbf{a} \cdot \mathbf{p}(t)$ is such that $t \rightarrow o(t; \tau)$ is an embedding of the circle, then $t \rightarrow o'(t; \tau')$ is also an embedding for $o'(t) = \mathbf{a} \cdot \mathbf{p}'(t; \tau')$ by Theorem II.8. The proof of Theorem II.8 uses the choice $\tau' = \tau T'/T$. \square

Suppose that the delay map of a signal obtained by projecting the first component of a periodic orbit does not embed in \mathbb{R}^3 . We will show that the differential equation $\frac{dx}{dt} = f(x)$, $x \in \mathbb{R}^d$, can be perturbed ever so slightly such that a nearby periodic orbit of the perturbed equation results in an embedding of the circle. The proof relies on constructing a tube around the periodic orbit. A tube around a periodic orbit is illustrated in Figure 2.6.

To construct a tube around any periodic orbit in \mathbb{R}^d , we begin by defining \mathcal{P}^r in analogy to \mathcal{O}^r . Let \mathcal{P}^r be the set of periodic orbits $\mathbf{p} : [0, T) \rightarrow \mathbb{R}^d$ that are r times continuously differentiable. As before, we assume that $[0, T)$ is a parametrization of S^1 and $T > 0$ for the period. As a part of the definition of \mathcal{P} , we require $\frac{d\mathbf{p}}{dt} \neq 0$ for $t \in [0, T)$. The set \mathcal{P}^r is endowed with a topology by defining the metric d_r in

analogy with (2.1):

$$d_r(\mathbf{p}, \mathbf{p}') = \sup_{k=0, \dots, r} \sup_{0 \leq s < 1} \left| \|\mathbf{p}^{(k)}(sT) - \mathbf{p}'^{(k)}(sT')\| \right| + |T - T'|.$$

The norm over \mathbb{R}^d is the 2-norm. The k th derivative of \mathbf{p} is denoted by $\mathbf{p}^{(k)}$. For convenience, $\frac{d\mathbf{p}}{dt}$ and $\frac{d^2\mathbf{p}}{dt^2}$ are also denoted as $\dot{\mathbf{p}}$ and $\ddot{\mathbf{p}}$, respectively. The tangent vector at t is defined as $\mathbf{s}(t) = \dot{\mathbf{p}}(t) / \|\dot{\mathbf{p}}(t)\|$.

We denote the projection from \mathbb{R}^d to the first coordinate by π_1 . If \mathbf{p} is a solution of the dynamical system $\frac{dx}{dt} = f(x)$, we wish to show that either $o(t) = \pi_1 \mathbf{p}(t)$ is such that $t \rightarrow o(t; \tau)$ is an embedding of the circle $[0, T)$ for some delay $\tau > 0$, or that there exists an arbitrarily close perturbed dynamical system $\frac{dx}{dt} = f'(x)$ with a nearby periodic orbit \mathbf{p}' such that $t \rightarrow o'(t; \tau)$ is an embedding of the circle, if $o' = \pi_1 \circ \mathbf{p}'$.

To begin with, the signal $o(t)$ may even be identically zero. In our proof, we use the results of the previous section to perturb it to $o'(t)$ such that $t \rightarrow o'(t; \tau)$ is an embedding and then show how to perturb the flow to realize $o'(t)$ as $\pi_1 \circ \mathbf{p}'$.

The next lemma constructs a tube around the periodic orbit \mathbf{p} in \mathbb{R}^d (see Figure 2.6). That tube will be used to perturb f to f' . Known results in differential geometry [14, 27] may be used to assert the existence of a tube. However, uniformity and smoothness guarantees that we need could not be found in the literature. Therefore, an elementary proof of the lemma is included. The proof will later be modified to deduce the existence of a tube whose radius is uniform in a neighborhood of \mathbf{p} . In the following lemma, δ may be thought of as the radius of a tube around \mathbf{p} .

Lemma II.11. *Suppose $\mathbf{p} \in \mathcal{P}^r$, $r \geq 2$, and that its period is $T > 0$. Then there exists $\delta > 0$ such that*

- $\|\dot{\mathbf{p}}(t)\|^2 - \delta \|\ddot{\mathbf{p}}\| > \delta$ for $t \in [0, T)$,

- if $x \in \mathbb{R}^d$ and $\text{dist}(x, \mathbf{p}) \leq \delta$, there exists a unique $t \in [0, T)$ such that $\text{dist}(x, \mathbf{p}) = \|x - \mathbf{p}(t)\|$.

Proof. The proof is organized so as to be easy to uniformize in the next lemma.

1. *Choice of \mathbf{m} and \mathbf{m}^* .* Let $2\mathbf{m} = \min_{t \in [0, T)} \|\dot{\mathbf{p}}(t)\| > 0$ and $\mathbf{m}^* = \max_{t \in [0, T)} \|\ddot{\mathbf{p}}(t)\|$.

The first part of the lemma would be satisfied if $4\mathbf{m}^2 - \delta\mathbf{m}^* > \delta$, or if $\delta < \frac{4\mathbf{m}^2}{1+\mathbf{m}^*}$.

2. *Choice of \mathfrak{M} and \mathfrak{r} .* First, we introduce the notation

$$\left. \frac{d\mathbf{p}}{dt} \right|_{[t_1, t_2]}$$

for a vector each of whose components is the corresponding component of $\dot{\mathbf{p}}$ evaluate at some $t \in [t_1, t_2]$. Crucially, each component may chose a different t . This notation will facilitate application of the mean value theorem. The interval $[t_1, t_2]$ may wrap around $[0, T)$, in which case the interval width must be taken to be $T + t_2 - t_1$ and not $t_2 - t_1$. We ignore such wrap-arounds from this point onwards.

Suppose $t_1 < t_2$ and $t_m = \frac{t_1+t_2}{2}$. Then

$$\left\| \dot{\mathbf{p}}(t_m) - \left. \frac{d\mathbf{p}}{dt} \right|_{[t_1, t_2]} \right\| \leq \max_{t \in [0, T)} \|\ddot{\mathbf{p}}\|_{\infty} \sqrt{d}(t_2 - t_1).$$

The \sqrt{d} factor here arises in converting a componentwise bound using the ∞ -norm to a bound on the 2-norm. Evidently, if we choose $\mathfrak{M} = \max_{t \in [0, T)} \|\ddot{\mathbf{p}}\|_{\infty} \times \sqrt{d}$ and $\mathfrak{r} = \frac{\mathbf{m}}{\mathfrak{M}}$, we may assert that

$$(2.3) \quad \left\| \dot{\mathbf{p}}(t_m) - \left. \frac{d\mathbf{p}}{dt} \right|_{[t_1, t_2]} \right\| \leq \mathbf{m}$$

for $t_1 < t_2$ and $t_2 - t_1 \leq \mathfrak{r}$.

If $\mathbf{s}(t_m)$ is the unit tangent vector to \mathbf{p} at t_m , we have

$$\begin{aligned} \mathbf{s}(t_m) \cdot (\mathbf{p}(t_2) - \mathbf{p}(t_1)) &= \mathbf{s}(t_m) \cdot \left(\frac{d\mathbf{p}}{dt} \Big|_{[t_1, t_2]} (t_2 - t_1) \right) \\ &= \mathbf{s}(t_m) \cdot \dot{\mathbf{p}}(t_m)(t_2 - t_1) + \mathbf{s}(t_m) \cdot \left(\frac{d\mathbf{p}}{dt} \Big|_{[t_1, t_2]} - \dot{\mathbf{p}}(t_m) \right) (t_2 - t_1), \end{aligned}$$

where the first equality is obtained by applying the mean value theorem to each component of $\mathbf{p}(t_2) - \mathbf{p}(t_1)$. Now, $\mathbf{s}(t_m) \cdot \dot{\mathbf{p}}(t_m) = \|\dot{\mathbf{p}}(t_m)\| \geq 2\mathbf{m}$ by choice of \mathbf{m} . By (2.3), the second term in the display above is at most $\mathbf{m}(t_2 - t_1)$ in magnitude. Therefore,

$$|\mathbf{s}(t_m) \cdot (\mathbf{p}(t_2) - \mathbf{p}(t_1))| \geq \mathbf{m}(t_2 - t_1)$$

for $t_1 < t_2$ and $t_2 - t_1 \leq \mathbf{r}$.

3. *Choice of \mathfrak{M}^* .* Suppose \mathbf{w}_1 is a vector orthogonal to $\mathbf{s}(t_1)$ and $t_1 < t_2$ with $t_m = \frac{t_1 + t_2}{2}$ as before. Then, we have $\mathbf{s}(t_m) \cdot \mathbf{w}_1 = (\mathbf{s}(t_m) - \mathbf{s}(t_1)) \cdot \mathbf{w}_1$, which implies

$$\begin{aligned} |\mathbf{s}(t_m) \cdot \mathbf{w}_1| &\leq \|\mathbf{s}(t_m) - \mathbf{s}(t_1)\| \|\mathbf{w}_1\| \\ &\leq \sqrt{d} \max_{t \in [0, T]} \|\dot{\mathbf{s}}(t)\|_{\infty} (t_m - t_1) \|\mathbf{w}_1\|, \end{aligned}$$

where the \sqrt{d} factor arises in converting a componentwise bound to a bound on the 2-norm. An explicit formula for $\dot{\mathbf{s}}$, the time derivative of the unit tangent, will be given in the next proof. If we choose $\mathfrak{M}^* = \sqrt{d} \max_{t \in [0, T]} \|\dot{\mathbf{s}}(t)\|_{\infty}$, we may replicate the argument given using \mathbf{w}_1, t_1 with \mathbf{w}_2, t_2 and assert

$$|\mathbf{s}(t_m) \cdot \mathbf{w}_1| < \mathfrak{M}^* \|\mathbf{w}_1\| (t_2 - t_1) \quad \text{and} \quad |\mathbf{s}(t_m) \cdot \mathbf{w}_2| < \mathfrak{M}^* \|\mathbf{w}_2\| (t_2 - t_1).$$

4. *Choice of Δ .* We define $\Delta = \min_{|t_2 - t_1| \geq \mathbf{r}} \|\mathbf{p}(t_2) - \mathbf{p}(t_1)\|$. Because a periodic orbit cannot self-intersect, we must have $\Delta > 0$.

We will choose δ to be smaller than the least of

$$\frac{4\mathfrak{m}^2}{1 + \mathfrak{m}^*}, \frac{\mathfrak{m}}{2\mathfrak{M}^*}, \frac{\Delta}{2}.$$

The first part of the lemma follows immediately. Now suppose $x \in \mathbb{R}^d$ and $\text{dist}(x, \mathbf{p}) \leq \delta$. Suppose $\text{dist}(x, \mathbf{p})$ is equal to $\|x - \mathbf{p}(t_1)\|$ as well as $\|x - \mathbf{p}(t_2)\|$ for $t_1 < t_2$. By item 4 above, we must have $t_2 - t_1 < \mathfrak{r}$, which we will now assume.

Because $t = t_1$ minimizes $(x - \mathbf{p}(t)) \cdot (x - \mathbf{p}(t))$, we may differentiate and deduce $(x - \mathbf{p}(t_1)) \cdot \dot{\mathbf{p}}(t_1) = 0$. Equivalently $(x - \mathbf{p}(t_1)) \cdot \mathbf{s}(t_1) = 0$. Thus, we may write $x = \mathbf{p}(t_1) + \mathbf{w}_1$, with \mathbf{w}_1 orthogonal to the tangent $\mathbf{s}(t_1)$ and $\text{dist}(x, \mathbf{p}) = \|\mathbf{w}_1\|$. Likewise, we may write $x = \mathbf{p}(t_2) + \mathbf{w}_2$, with \mathbf{w}_2 orthogonal to the tangent $\mathbf{s}(t_2)$ and $\text{dist}(x, \mathbf{p}) = \|\mathbf{w}_2\|$.

From $\mathbf{p}(t_1) + \mathbf{w}_1 = \mathbf{p}(t_2) + \mathbf{w}_2$, we obtain

$$\mathbf{s}(t_2) \cdot (\mathbf{p}(t_2) - \mathbf{p}(t_1)) = \mathbf{s}(t_2) \cdot (\mathbf{w}_1 - \mathbf{w}_2).$$

Taking absolute values, applying item 2 above to the left hand side, and item 3 above to the right hand side, we get

$$\mathfrak{m}(t_2 - t_1) < \mathfrak{M}^* (\|\mathbf{w}_1\| + \|\mathbf{w}_2\|) (t_2 - t_1),$$

or $\text{dist}(x, \mathbf{p}) > \frac{\mathfrak{m}}{2\mathfrak{M}^*} \geq \delta$, contradicting our hypothesis about x . Thus, the assumption $t_1 < t_2$ is mistaken, and we can only have $t_1 = t_2$ proving the second part of the lemma. \square

The following lemma is a uniform version of the preceding Lemma II.11. The lemma allows us to construct a tube of radius δ around all periodic orbits of period T that are within a distance ϵ of \mathbf{p} . Its proof is a minor modification of the preceding proof.

Lemma II.12. *Suppose $\mathbf{p} \in \mathcal{P}^r$, $r \geq 2$, and that its period is $T > 0$. Then there exist $\epsilon > 0$ and $\delta > 0$ such that $\mathbf{p}' \in \mathcal{P}^r$, with the same period as \mathbf{p} , and $d_r(\mathbf{p}, \mathbf{p}') \leq \epsilon$ imply that*

- $\left| \|\dot{\mathbf{p}}'(t)\|^2 - \delta \|\ddot{\mathbf{p}}'\| \right| > \delta$ for $t \in [0, T)$,
- if $x \in \mathbb{R}^d$ and $\text{dist}(x, \mathbf{p}') \leq \delta$, then there exists a unique $t \in [0, T)$ such that $\text{dist}(x, \mathbf{p}') = \|x - \mathbf{p}'(t)\|$.

Proof. In the previous proof, we demonstrated the existence of a δ that works for \mathbf{p} . This proof comes down to choosing ϵ so that $\mathbf{m}, \mathbf{m}^*, \mathfrak{M}, \mathfrak{r}, \mathfrak{M}^*$, and Δ work for all \mathbf{p}' with the same period as \mathbf{p} and satisfying $d_r(\mathbf{p}, \mathbf{p}') \leq \epsilon$.

The quantity \mathbf{m} is a lower bound on $\|\dot{\mathbf{p}}(t)\|$. Because ϵ controls $\|\dot{\mathbf{p}}(t) - \dot{\mathbf{p}}'(t)\|$ over $t \in [0, T)$, we may assume ϵ small enough and replace \mathbf{m} by $\mathbf{m}/2$ to make it work for \mathbf{p}' .

The quantity \mathbf{m}^* is an upper bound on $\|\ddot{\mathbf{p}}(t)\|$. Because ϵ controls $\|\ddot{\mathbf{p}}(t) - \ddot{\mathbf{p}}'(t)\|$ over $t \in [0, T)$, we may assume ϵ small enough and replace \mathbf{m}^* by $2\mathbf{m}^*$ to make it work for \mathbf{p}' .

The quantity \mathfrak{M} is essentially an upper bound on $\|\ddot{\mathbf{p}}(t)\|_\infty$. Because ϵ controls $\|\ddot{\mathbf{p}}(t) - \ddot{\mathbf{p}}'(t)\|$ over $t \in [0, T)$, we may assume ϵ small enough and replace \mathfrak{M} by $2\mathfrak{M}$ to make it work for \mathbf{p}' .

We may use the same definition of $\mathfrak{r} = \frac{\mathbf{m}}{\mathfrak{M}}$ after modifying \mathbf{m} and \mathfrak{M} as above.

The quantity \mathfrak{M}^* is essentially an upper bound on $\|\dot{\mathbf{s}}(t)\|_\infty$. The unit tangent vector \mathbf{s} is given by $\mathbf{s} = \dot{\mathbf{p}} / (\dot{\mathbf{p}} \cdot \dot{\mathbf{p}})^{1/2}$. Differentiating, we obtain

$$\dot{\mathbf{s}} = \frac{\ddot{\mathbf{p}}}{(\dot{\mathbf{p}} \cdot \dot{\mathbf{p}})^{1/2}} - \frac{\dot{\mathbf{p}} (\ddot{\mathbf{p}} \cdot \dot{\mathbf{p}})}{(\dot{\mathbf{p}} \cdot \dot{\mathbf{p}})^{3/2}}.$$

Because $r \geq 2$, we may control the variation in $\mathbf{p}, \dot{\mathbf{p}}$, and $\ddot{\mathbf{p}}$ by making ϵ small. Thus, we may assume ϵ small enough and replace \mathfrak{M}^* by $2\mathfrak{M}^*$ to make it work for \mathbf{p}' .

We begin by defining $\Delta = \min_{|t_2 - t_1| \geq \epsilon} \|\mathbf{p}(t_2) - \mathbf{p}(t_1)\|$ as before. By assuming ϵ small enough and replacing Δ by $\Delta/2$, we may assume Δ to work for all \mathbf{p}' .

The rest of the proof of the previous lemma works without change. \square

Half of the smoothness lemma that follows is a special case of the main theorem in [14]. Given a periodic orbit and a tube around it, the lemma shows that each point in the tube can be expressed as a sum of a point on the periodic orbit and a vector orthogonal to the tangent at that point. Additionally, the lemma provides smoothness and uniformity guarantees.

Lemma II.13. *Assume the same setting as in the previous Lemma II.12. Given \mathbf{p}' with $d_r(\mathbf{p}, \mathbf{p}') \leq \epsilon$ and a point $x_0 \in \mathbb{R}^d$ with $\text{dist}(x_0, \mathbf{p}') \leq \delta$, we may send $x_0 \rightarrow t_0$, where $\mathbf{p}'(t_0)$ is the unique point on \mathbf{p}' closest to x_0 , and $x_0 \rightarrow \mathbf{w}_0$, where $\mathbf{w}_0 = x_0 - \mathbf{p}'(t_0)$. The functions $t_0(x_0)$ and $\mathbf{w}_0(x_0)$ are C^{r-1} . In addition, the magnitudes of all derivatives of order $r - 1$ or less have upper bounds that depend only on \mathbf{p} and δ .*

Proof. It is sufficient to prove the lemma for $t_0(x_0)$. The assertions about $\mathbf{w}_0(x_0)$ follow easily from that point.

The function $(x_0 - \mathbf{p}'(t)) \cdot (x_0 - \mathbf{p}'(t))$ has a unique minimum at $t = t_0$. By differentiating, we get the equation $(x_0 - \mathbf{p}'(t_0)) \cdot \frac{d\mathbf{p}'(t_0)}{dt} = 0$. If we define

$$\mathfrak{f}(x_0, t_0) = (x_0 - \mathbf{p}'(t_0)) \cdot \frac{d\mathbf{p}'(t_0)}{dt}$$

We may think of the equation $\mathfrak{f}(x_0, t_0) = 0$ as implicitly defining $t_0(x_0)$ as a function of x_0 . We have

$$\frac{\partial \mathfrak{f}}{\partial t_0} = \ddot{\mathbf{p}}'(t_0) \cdot (x_0 - \mathbf{p}'(t_0)) - \frac{d\mathbf{p}'(t_0)}{dt} \cdot \frac{d\mathbf{p}'(t_0)}{dt}.$$

Here $\|x_0 - \mathbf{p}'(t_0)\| = \text{dist}(x_0, \mathbf{p}') \leq \delta$. We may use the first part of Lemma II.12 and conclude that the partial derivative $\partial \mathfrak{f} / \partial t_0$ is greater than δ in magnitude.

Thus, the C^{r-1} smoothness of $t_0(x_0)$ follows by the implicit function theorem. To upper bound the magnitudes of the derivatives, we simply have to use chain rule and implicit differentiation. For example, if $x_0 = (\xi_1, \dots, \xi_d)$, we have

$$(2.4) \quad \frac{\partial t_0}{\partial \xi_1} = - \frac{\mathbf{e}_1 \cdot \frac{d\mathbf{p}'(t_0)}{dt}}{\frac{\partial f}{\partial t_0}},$$

where $\mathbf{e}_1 = (1, 0, \dots, 0)$. Now the denominator is δ or more in magnitude and the magnitude of the numerator has an upper bound that depends only on \mathbf{p}' .

To obtain bounds for derivatives of $t_0(x_0)$ of order $r-1$ or less, we may repeatedly differentiate (2.4). The bounds on the derivatives obtained in this manner depend only on the first r derivatives of \mathbf{p}' and δ . If we assume $\epsilon < 1$, we may bound the first r derivatives of \mathbf{p}' in terms of the derivatives of \mathbf{p} . Thus, the magnitudes of all derivatives of order $r-1$ or less have upper bounds that depend only on \mathbf{p} and δ . \square

Theorem II.14. *Let $\mathbf{p}(t)$ be a periodic solution of the dynamical system $d\mathbf{x}/dt = f(\mathbf{x})$, where f is C^{r-1} . If $o(t) = \pi_1 \mathbf{p}(t)$ is a periodic signal, there exists either a delay $\tau > 0$ such that $t \rightarrow o(t; \tau)$, $0 \leq t < T$, is an embedding of the circle $[0, T)$ or another vector field f' , arbitrarily close to f in the C^{r-1} topology, with a periodic solution $\mathbf{p}'(t)$ arbitrarily close to $\mathbf{p}(t)$ in \mathcal{P}^r and of the same period such that $t \rightarrow \pi_1 \mathbf{p}'(t; \tau)$ is an embedding of the circle $[0, T)$ for some $\tau > 0$.*

Proof. Let $o(t) = \pi_1 \mathbf{p}(t)$ and assume that there is no delay $\tau > 0$ such that $t \rightarrow o(t; \tau)$ is an embedding. By Lemma II.6, we can find a periodic signal $o'(t)$ of period T , and arbitrarily close to $o(t)$ in \mathcal{O}^r , such that $t \rightarrow o'(t; \tau)$ for some $\tau > 0$. Define

$$(2.5) \quad \mathbf{p}'(t) = \mathbf{p}(t) + \begin{pmatrix} o'(t) - o(t) \\ 0 \\ \vdots \end{pmatrix}.$$

It suffices to construct a vector field f' such that $\mathbf{p}'(t)$ is a periodic solution of $\frac{dx}{dt} = f'(x)$ and $f' \rightarrow f$ as $\mathbf{p}' \rightarrow \mathbf{p}$.

Using Lemmas II.11 and II.12, find an $\epsilon > 0$ and a $\delta > 0$, such that a δ -tube may be constructed as in the lemma for all periodic orbits \mathbf{p}' of the same period as \mathbf{p} satisfying $d_r(\mathbf{p}, \mathbf{p}') < \epsilon$. In addition, by taking o' close enough to o , we may assume that $d_r(\mathbf{p}, \mathbf{p}') < \epsilon$.

The following calculation is the heart of the proof:

$$\begin{aligned} \frac{d\mathbf{p}'(t)}{dt} &= \frac{d\mathbf{p}(t)}{dt} + \epsilon_1(t) \\ &= f(\mathbf{p}(t)) + \epsilon_1(t) \\ &= f(\mathbf{p}'(t)) + \epsilon_1(t) + \epsilon_2(t), \end{aligned}$$

where

$$\epsilon_1(t) = \begin{pmatrix} \frac{d(o'(t) - o(t))}{dt} \\ 0 \\ \vdots \end{pmatrix}$$

and $\epsilon_2(t) = f(\mathbf{p}(t)) - f(\mathbf{p}'(t))$. Evidently, as $o' \rightarrow o$ in \mathcal{O}^r , the periodic signals $\epsilon_1(t)$ and $\epsilon_2(t)$ go to 0 in \mathcal{O}^{r-1} .

Let $\lambda : \mathbb{R} \rightarrow \mathbb{R}$ be a C^∞ bump function with $\lambda(x) = 1$ for $|x| \leq 1/2$ and $\lambda(x) = 0$ for $|x| \geq 3/4$. Suppose x_0 is a point in the δ -tube around \mathbf{p}' . Then Lemma II.13, allows us to write x_0 as $x_0 = \mathbf{p}'(t_0(x_0)) + \mathbf{w}_0(x_0)$. The perturbation $\delta f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is defined as

$$\delta f(x_0) = (\epsilon_1(t_0(x_0)) + \epsilon_2(t_0(x_0)))\lambda \left(\frac{\mathbf{w}_0(x_0) \cdot \mathbf{w}_0(x_0)}{\delta^2} \right)$$

for x_0 in the δ -tube around \mathbf{p}' , and zero otherwise. As a consequence of Lemma II.13, $\delta f \rightarrow 0$ in the C^{r-1} sense as $o' \rightarrow o$.

By construction, $\mathbf{p}'(t)$ is a periodic solution of the dynamical system $dx/dt = f'(x)$, with $f' = f + \delta f$. □

Finally, as a consequence of Proposition II.10 and Theorem II.14, we have the following theorem.

Theorem II.15. *Let $\frac{dx}{dt} = f(x)$, where $x \in \mathbb{R}^d$, $f : U \rightarrow \mathbb{R}^d$, and U an open subset of \mathbb{R}^d , be a C^r , $r \geq 2$, dynamical system. Let $\mathbf{a} \in \mathbb{R}$ be a nonzero vector. Let $\mathbf{p} : [0, T) \rightarrow U$ be a hyperbolic periodic solution of period $T > 0$. There exists an open neighborhood of f in the C^{r-1} topology such that for an open and dense set of g in that neighborhood admit a nearby hyperbolic periodic solution $\mathbf{p}'(t)$ of $dx'/dt = g(x')$ of period T' and a delay $\tau' > 0$ such that the delay map $t \rightarrow (\mathbf{a} \cdot \mathbf{p}'(t), \mathbf{a} \cdot \mathbf{p}'(t - \tau'), \mathbf{a} \cdot \mathbf{p}'(t - 2\tau'))$ is an embedding of the circle $[0, T')$ in \mathbb{R}^3 .*

Proof. Proposition II.10 and Theorem II.14 imply Theorem II.15 with $\mathbf{a} = (1, 0, \dots, 0)$. The theorem may be reduced to that case for any $\mathbf{a} \neq 0$ by a linear change of variables. □

The theorem does not assert that periodic orbits can be embedded in \mathbb{R}^3 for an open and dense set of C^r vector fields g . Instead, the theorem limits itself to a neighborhood of a vector field f which is known to admit a hyperbolic periodic orbit. Such a restriction is essential because there exist open sets of vector fields none of which admit any periodic solution.

2.4 Discussion

In this chapter, we have considered an extension of the delay coordinate embedding theory. The current embedding theory of Sauer et al [47] is based on fixing the dynamical system and perturbing the observation function. We have obtained an embedding theorem for periodic orbits that fixes the observation function but perturbs the dynamical system.

Periodic solutions are a special case that arise in applications [5, 15]. However, a generalization to a broader setting is desirable both from the theoretical point of view as well as for wider applicability.

Our approach in this chapter relies heavily on the periodicity of signals. Yet some differences between our approach and that of Sauer et al may be pertinent to more general settings. The approach of Sauer et al is able to handle aspects of the embedding result, such as injectivity, immersivity, and distinct points on the same periodic orbit, relatively independently. Our argument is more layered. A global argument is structured above a local argument, and the argument for periodic orbits relies on the argument for periodic signals.

CHAPTER III

Prevalence of delay embeddigs with a fixed observation function

3.1 Introduction

Let $x_{j+1} = \phi(x_j)$ be a dynamical system. If o is a scalar valued observation function, the delay map is given by

$$F_0(x_1) = (o(x_1), \dots, o(x_D)).$$

The question of when F_0 is an embedding was considered by Aeyels [2] and Takens [54]. Suppose that $x_j \in \mathbb{R}^n$ but with the dynamics confined to an invariant submanifold of dimension $d \leq n$. Alternatively, we may assume $x_j \in \mathbf{m}$, where \mathbf{m} is a manifold of dimension d . Based on an analogy to Whitney embedding [24], we may expect F_0 to be an embedding for generic o for embedding dimension $D \geq 2d + 1$. Here genericity is with respect to the space of functions o under a C^r topology with $r \geq 2$ [24].

Sauer et al [47] introduced a new point of view, supported by deep ideas, into the theory of delay embeddings. If $x \in \mathbb{R}^n$ and $\alpha \in \mathbb{Z}_0^n$ is a multi-index, denote the monomial x^α by $p_\alpha(x)$. Instead of assuming the observation function o to be any C^r function, Sauer et al take the observation function to be the sum of some fixed function o^* and a finite linear combination of the monomials $p_\alpha(x)$. Proofs

of genericity rely on “bump” functions or C^∞ functions with compact support. Although the device of bump functions is of much utility in differential topology [24], bump functions hardly ever arise in applications. In contrast, physical models often use polynomials. Thus, limiting the perturbations to a finite linear combination of polynomials is a welcome shift in point of view.

A property is generic in a Baire space if it holds for a countable intersection of open and dense sets. A generic set is always dense but it may be of probability zero (in a reasonable sense). For example, generic subsets of $[0, 1]$ of probability zero may be constructed easily. Thus, it may be questioned whether the concept of genericity captures the notion of what is typical in applications.

Sauer et al [47] answered that question by introducing the notion of prevalence. To say that delay embeddings are prevalent is equivalent to saying that the delay map is an embedding for almost every linear combination of polynomials. If probabilities are defined by normalizing the Lebesgue measure, we may say that the delay map is an embedding with probability one.

Suppose $x_k = \phi^k(x_1)$ and $y_k = \phi^k(y_1)$. For F_0 to be an injection, we must have $F_0(x_1) \neq F_0(y_1)$ whenever $x_1 \neq y_1$. A major difficulty in the proof of injectivity arises in handling points $x_1 \neq y_1$ but with overlapping orbits. For example, we may have $y_1 = x_2$ or $y_1 = x_3$. Related difficulties arise in handling periodic points and in the proof of immersivity (an embedding must be injective as well as immersive). Sauer et al [47] introduced several key ideas for handling these difficulties. However, there is a minor gap in their proof. In section 4, we fix that gap and show that earlier mathematical treatments have serious deficiencies. Therefore, proofs prior to Sauer et al cannot be accepted.

The proof of Sauer et al [47] is quite informal. We give a more formally precise

development of their ideas in sections 2 and 3. Later, we consider the case where the observation map is fixed at $o = \pi_1$, with π_1 being the projection to the first coordinate and with polynomial perturbations applied directly to ϕ . Ideas essential for the new developments are interspersed in sections 2 and 3. Sauer et al include a filtering step applied to the delay map in their main theorems. In addition to mathematical informality, the filtering step makes the essential ideas difficult to grasp and verify. Thus, the filtering step is omitted in section 4, where we derive their main results in a modified form.

From section 5 onwards, we treat the case where $o = \pi_1$ and ϕ itself is perturbed by polynomials. There are two main motivations for considering this case. First, from a purely aesthetic point of view, it is desirable to make the theory of delay embeddings depend upon the dynamics and not the observation function. Second, the setting with $o = \pi_1$ is pertinent to applications. For example, the most natural way to extract a time series from a fluid flow is to simply record the fluid velocity at a fixed point [54].

The main technical novelty in our approach is related to the concept of Lebesgue points. Our delay embedding theorem for the $o = \pi_1$ case requires $D \geq 4d + 2$, although the work in chapter 2 of this thesis suggests $D \geq 2d + 1$. In the concluding section, we express the hope that the technique of Lebesgue points may prove useful in obtaining prevalence versions of some classical results in dynamical systems theory. In that regard, we mention the extensions of delay embedding theory to PDE by Robinson [42, 43]. A more complete account of other mathematical investigations in embedding theory may be found in the introduction to chapter 2 in this thesis.

3.2 Transfer of volume

A key idea in the work of Sauer et al [47] is to transfer volumes from embedding space to parameter space. For an example of what we mean by transfer of volume, suppose A is a square matrix. Then a volume equal to \mathbf{v} in the range is transferred to $\mathbf{v}/\det A$ in the domain.

Suppose $G : \mathbb{R}^{D_\alpha} \times \mathbb{R}^D \rightarrow \mathbb{R}^D$ is a C^r function with $r \geq 2$. Here \mathbb{R}^{D_α} is the space of parameters and we will denote a point in parameter space by (c_α) or c_α , with the understanding that (c_α) (or c_α) is a column vector. The transfer of volume is carried out with fixed $\mathfrak{z} \in \mathbb{R}^D$. Thus, the dependence of $G(c_\alpha, \mathfrak{z})$ on \mathfrak{z} , which will be nonlinear, does not come up in the transfer of volume argument. When the map ϕ is fixed and only the observation function is parametrized, G is linear in the parameters c_α . The embedding space is \mathbb{R}^D and the dimension D of this space is of much importance. The rank of G is mainly constrained by D because $D_\alpha \gg D$, and the rank determines how much volume (or how little, with lesser the better) is transferred from embedding space to parameter space.

In the following lemma and later we refer to $\mu(B_1 \cap B_2)/\mu(B_2)$, where $\mu(\cdot)$ is the Lebesgue measure, as the probability of B_1 relative to B_2 (both sets are assumed to be measurable). Measure will always refer to Lebesgue measure. The following lemma transfers the volume of a ball of radius $L\epsilon$ in \mathbb{R}^D to parameter space. All norms in this chapter are spectral or L^2 norms.

Lemma III.1 ([47]). *Let $\mathfrak{g}(c_\alpha) = A(c_\alpha) + \mathfrak{g}_0$ be a linear (affine) map from \mathbb{R}^{D_α} to \mathbb{R}^D , with A being a $D \times D_\alpha$ matrix. Suppose the first \mathfrak{r} singular values of A are at least as great as $\sigma > 0$. Then the measure of the set*

$$(3.1) \quad \left\{ c_\alpha \mid \|A(c_\alpha) + \mathfrak{g}_0\| \leq L\epsilon \right\} \cap \left\{ c_\alpha \mid \|c_\alpha\| \leq a \right\}$$

is less than or equal to

$$(3.2) \quad 2^{D_\alpha} L^\tau \epsilon^\tau a^{D_\alpha - \tau} / \sigma^\tau,$$

and the probability of $\|A(c_\alpha) + \mathbf{g}_0\| \leq L\epsilon$ relative to $\|c_\alpha\| \leq a$ is less than or equal to

$$D_\alpha! L^\tau \epsilon^\tau / \sigma^\tau a^\tau.$$

Proof. Suppose $\mathbf{u}_1, \dots, \mathbf{u}_{D_\alpha}$ are the right singular vectors, $\mathbf{v}_1, \dots, \mathbf{v}_D$ are the left singular vectors, and $\sigma_1, \dots, \sigma_{D_\alpha}$ the singular values of A . (see [56]). Let $(c_\alpha) = \sum_{i=1}^{D_\alpha} \mathbf{c}_i \mathbf{u}_i$ and $\mathbf{g}_0 = \sum_{i=1}^D \mathbf{g}_i \mathbf{v}_i$.

For $i = 1, \dots, \tau$, $\|A(c_\alpha) + \mathbf{g}_0\| \leq L\epsilon$ implies that $|\sigma_i \mathbf{c}_i + \mathbf{g}_i| \leq L\epsilon$ and therefore $|\mathbf{c}_i + \mathbf{g}_i/\sigma_i| \leq L\epsilon/\sigma_i \leq L\epsilon/\sigma$. Thus, the coefficient \mathbf{c}_i must lie in an interval of measure less than $2L\epsilon/\sigma$ for $i = 1, \dots, \tau$.

For $i = \tau + 1, \dots, D_\alpha$, $\|c_\alpha\| \leq a$ implies that \mathbf{c}_i must vary inside the interval $[-a, a]$, whose length is $2a$.

Therefore, the volume of the set (3.1) is bounded above by $(2L\epsilon/\sigma)^\tau (2a)^{D_\alpha - \tau}$, which simplifies to (3.2).

For the statement about the probability of $\|A(c_\alpha) + \mathbf{g}_0\| \leq L\epsilon$ relative to $\|c_\alpha\| \leq a$, we divide (3.2) by γa^{D_α} , where γ is the volume of the unit sphere in \mathbb{R}^{D_α} , to obtain

$$2^{D_\alpha} L^\tau \epsilon^\tau / \gamma \sigma^\tau a^\tau.$$

The proof is completed using $\gamma = \pi^{D_\alpha/2} / \Gamma(D_\alpha/2 + 1) \geq 2^{D_\alpha} / D_\alpha!$. \square

Lemma III.1 shows how a volume $\|\mathbf{g}(c_\alpha)\| \leq L\epsilon$ in embedding space is transferred to a probability relative to $\|c_\alpha\| \leq a$ in parameter space. The transferred probability

is proportional to ϵ^τ , and therefore, as the rank τ increases, the probability becomes smaller.

To obtain prevalence with the observation function fixed and the map parametrized, we will rely on the following nonlinear transfer of volume lemma. When the previous Lemma III.1 is applied, L will be a Lipschitz constant. When the following lemma is applied, L will be a Lipschitz constant as well as a bound on the quadratic remainder term in a Taylor series.

Lemma III.2. *Suppose $\mathbf{g} : \mathbb{R}^{D_\alpha} \rightarrow \mathbb{R}^D$ is a C^2 function, with the Taylor series $\mathbf{g}(c_\alpha) = \mathbf{g}_0 + A(c_\alpha) + \mathbf{h}(c_\alpha)$. We assume that both $\mathbf{g}(\cdot)$ and $\mathbf{h}(\cdot)$ are defined for $\|c_\alpha\| \leq a$ and that $\|\mathbf{h}(c_\alpha)\| \leq L\|c_\alpha\|^2$. We also assume that the first τ singular values of A are at least as great as $\sigma > 0$. Then the probability of $\|\mathbf{g}(c_\alpha)\| \leq L\epsilon$ relative to $\|c_\alpha\| \leq \epsilon^{1/2}$ is less than or equal to*

$$D_\alpha! 2^\tau L^\tau \epsilon^{\tau/2} / \sigma^\tau$$

for $0 < \epsilon^{1/2} \leq a$.

Proof. If $\epsilon^{1/2} \leq a$ and $\|c_\alpha\| \leq \epsilon^{1/2}$, then $\|\mathbf{h}(c_\alpha)\| \leq L\epsilon$. Therefore, $\|A(c_\alpha) + \mathbf{g}_0\| \leq 2L\epsilon$. The proof is completed by applying the previous lemma with $L \leftarrow 2L$ and $a \leftarrow \epsilon^{1/2}$. \square

Applications of Lemmas III.1 and III.2 will require us to get a handle on singular values. We will turn to that in the next section. Before doing so, we recapitulate an elegant argument of Sauer et al [47]. This argument, although elementary, gives a good idea of the general approach when the observation function is parametrized.

Suppose K is a smooth sub-manifold or even a fractal set of box counting dimension d and with compact closure that is a subset of \mathbb{R}^n . Let the embedding dimension

be $D > 2d$. If $d \in \mathbb{Z}^+$, we can take $D = 2d + 1$ as in Whitney's embedding theorem [24]. The following assumptions are made about the constant C_K :

Assumption about C_K (1): The set K can be covered with C_K/ϵ^d ϵ -balls for any $\epsilon > 0$.

Assumption about C_K (2): The set $K \times K$ can be covered with C_K/ϵ^{2d} ϵ -balls for any $\epsilon > 0$.

All balls are spherical.

A linear map from \mathbb{R}^n to \mathbb{R}^D can be written as $F_\alpha(x) = \sum_{\alpha \in \mathcal{I}} c_\alpha \mathbf{m}_\alpha x$, where \mathcal{I} is the index set (i, j) , $1 \leq i \leq D$ and $1 \leq j \leq n$, and \mathbf{m}_α is the matrix with 1 in the i, j th position if $\alpha = (i, j)$ and zero everywhere else. Here $D_\alpha = nD$. We use c_α both to refer to an entry of the vector (c_α) as in the definition of F_α and to the vector as a whole as in $\|c_\alpha\|$. The slight ambiguity, which is resolved from context, is highly convenient. In most instances, c_α refers to the vector as a whole.

Define $G_\alpha(x, y) = F_\alpha(x) - F_\alpha(y)$. Assume $\|c_\alpha\| \leq a_0$. By compactness of the ball $\|c_\alpha\| \leq a_0$, we may assume the Lipschitz constant of $G_\alpha(x, y)$ (with respect to x, y) to be bounded above by L . Define $\mathcal{K}(\delta)$ to be the set of all points $(x, y) \in K \times K$ satisfying $\|x - y\| \geq \delta > 0$. Cover $\mathcal{K}(\delta)$ using C_K/ϵ^{2d} balls. Suppose $G_\alpha(x, y) = 0$ for some $(x, y) \in \mathcal{K}(\delta)$. Then by the Lipschitz bound, we must have $\|G_\alpha(x, y)\| \leq L\epsilon$ for (x, y) that is a center of one of the C_K/ϵ^{2d} covering $\mathcal{K}(\delta)$.

The rest of the argument hinges on transferring the volume $\|G_\alpha(x, y)\| \leq L\epsilon$ to parameter space. To do so, write $G_\alpha(x, y)$ in the form

$$\left(\mathbf{m}_{1,1}(x - y), \mathbf{m}_{1,2}(x - y), \dots \right) (c_\alpha)$$

and observe that every column in the resulting matrix is in \mathbb{R}^D and is all zeros except for a single entry equal to $\pi_i x - \pi_i y$, where π_i denotes the projection to the

i th coordinate, for some $i \in \{1, \dots, n\}$. The first D singular values of that matrix are all equal to $\|x - y\| \geq \delta$. Thus, we may transfer volumes using Lemma III.1 and assert that the probability of $G_\alpha(x, y) = 0$ for some $(x, y) \in \mathcal{K}(\delta)$ relative to $\|c_\alpha\| \leq a_0$ is at most

$$\frac{C_K}{\epsilon^{2d}} \times \frac{(nD)! L^D \epsilon^D}{\delta^D a_0^D}.$$

By taking the limit $\epsilon \rightarrow 0$ and because $D > 2d$, it follows that $G_\alpha(x, y) = 0$ for some $(x, y) \in \mathcal{K}(\delta)$ only for a set of c_α of probability zero relative to the ball $\|c_\alpha\| \leq a_0$. By taking the union of the probability zero sets with $\delta = 1, \frac{1}{2}, \frac{1}{2^2}, \dots$, we may conclude that $G_\alpha(x, y) = 0$ for some $(x, y) \in K \times K$, $x \neq y$, only for a set of c_α of probability zero relative to $\|c_\alpha\| \leq a_0$. Equivalently, $x \rightarrow F_\alpha(x)$ is injective for $x \in K$ with probability one relative to the ball $\|c_\alpha\| \leq a_0$ in parameter space.

The argument derives its power by simply refining the cover of $\mathcal{K}(\delta)$ by using smaller and smaller ϵ -balls. If $dF_\alpha(x, v)$ is the tangent map at x applied to the tangent vector v , then $dF_\alpha(x, v) = F_\alpha(v)$ because of the linearity of $F_\alpha(x)$ in x . If K is a submanifold then T_1K is the unit tangent bundle consisting of points (x, v) with $\|v\| = 1$. Injectivity may be proved by considering $dF_\alpha(x, v)$ instead of $G_\alpha(x, y)$, with Lemma III.1 invoked with $\sigma \leftarrow 1$.

3.3 Rank lemmas

In proving a version of the Whitney embedding theorem, the argument of Sauer et al [47] reviewed above writes $G_\alpha(x, y) = \mathcal{M}.c_\alpha$ and relies on explicit knowledge of singular values of \mathcal{M} . In general, singular values of \mathcal{M} cannot be obtained so explicitly. Instead, the approach is to first argue that \mathcal{M} has rank D or greater for every $(x, y) \in \mathcal{K}(\delta)$ and then observe that

$$\sigma_\delta = \min_{(x, y) \in \mathcal{K}(\delta)} \sigma_D(\mathcal{M}) > 0$$

because the D th singular value $\sigma_D(\mathcal{M})$ is continuous in x, y and $\mathcal{K}(\delta)$ is compact. The argument may then be completed by applying Lemma III.1 with $\sigma \leftarrow \sigma_\delta$.

To support such an argument, we give a few rank lemmas in this section. The first two lemmas are from [47]. Rank lemmas of this type are known in multivariate approximation theory [32], although they are buried inside more sophisticated results.

Suppose $z \in \mathbb{R}^d$. As noted already, the projection to the i th coordinate is denoted by π_i . If $\alpha = (\alpha_1, \dots, \alpha_d)$, $\alpha_i \in \mathbb{Z}^+ \cup \{0\}$, is a multi-index, then $z^\alpha = \prod_{i=1}^d (\pi_i z)^{\alpha_i}$ as usual and $|\alpha| = \sum_{i=1}^d |\alpha_i|$. In later arguments, it is essential to take the gradient of z^α with respect to z . For notational convenience, we always denote z^α by $p_\alpha(z)$. The index set \mathcal{I}_{D^+} is the set of all α such that $|\alpha| \leq D^+$. By elementary combinatorics, the cardinality of \mathcal{I}_{D^+} is $\binom{d+D^+}{D^+}$.

Suppose $z_1, z_2, \dots, z_{D'}$ are distinct points in \mathbb{R}^d . Then

$$(3.3) \quad \begin{pmatrix} p_\alpha(z_1) \\ \vdots \\ p_\alpha(z_{D'}) \end{pmatrix}$$

denotes the multivariate Vandermonde matrix with the column index $\alpha \in \mathcal{I}_{D^+}$ for some D^+ . The dimension of the matrix is $D' \times |\mathcal{I}_{D^+}|$, where $|\mathcal{I}_{D^+}|$ is the cardinality of \mathcal{I}_{D^+} .

Lemma III.3 ([47]). *For $\alpha \in \mathcal{I}_{D^+}$ and $D^+ \geq D' - 1$, the rank of the Vandermonde matrix (3.3) is equal to the number of its rows.*

Proof. Following [47], let Q be a $d \times d$ orthogonal matrix drawn from the Haar measure. If z_1 and z_2 are distinct, then $\pi_i z_1 \neq \pi_i z_2$ for any i for Q outside a set of measure 0. Therefore, we can find a Q such that $\pi_1 Q z_1, \dots, \pi_1 Q z_{D'}$ are distinct. We may interpolate arbitrary values at z_j using a univariate polynomial $\mathbf{p}(\pi_1 Q z)$ of

degree $D' - 1$. Because we can write

$$\mathbf{p}(\pi_1 Q z) = \begin{pmatrix} p_\alpha(z_1) \\ \vdots \\ p_\alpha(z_{D'}) \end{pmatrix} (c_\alpha)$$

for a suitable choice of c_α , it follows that the rank of (3.3) is equal to the number of its rows. \square

Let

$$(3.4) \quad \begin{pmatrix} \nabla p_\alpha(z_1) \\ \vdots \\ \nabla p_\alpha(z_{D'}) \end{pmatrix}$$

be the multivariate incomplete Hermite matrix at $z_1, \dots, z_{D'}$ and with $\alpha \in \mathcal{I}_{D^+}$.

Lemma III.4 ([47]). *The rank of the incomplete Hermite matrix (3.4) is equal to the number of its rows if $D^+ \geq D'$.*

Proof. Arguing as in the previous lemma, we may assume $\pi_i Q z_1, \dots, \pi_i Q z_{D'}$ to be distinct for $i = 1, \dots, d$. Following [47] and assuming Q to be the identity without loss of generality, we may then find a polynomial $\mathbf{p}_i(\pi_i z)$ of degree $D' - 1$ that interpolates the i th component of the prescribed gradients at $z_1, \dots, z_{D'}$. We may then obtain the prescribed gradients from $\mathbf{p}(z) = \int \mathbf{p}_1 d\pi_1 z + \dots + \int \mathbf{p}_d d\pi_d z$. Thus, a linear combination of the columns of (3.4) can produce any prescribed gradients. \square

To obtain prevalence results with a fixed observation function, Lemmas III.3 and

III.4 need to be combined into another lemma. Therefore, let

$$(3.5) \quad \begin{pmatrix} p_\alpha(z_1) \\ \vdots \\ p_\alpha(z_{D'}) \\ \nabla p_\alpha(z_1) \\ \vdots \\ \nabla p_\alpha(z_{D'}) \end{pmatrix}$$

be the multivariate Hermite matrix at $z_1, \dots, z_{D'}$ and with $\alpha \in \mathcal{I}_{D^+}$.

Lemma III.5. *The rank of the Hermite matrix (3.5) is equal to the number of its rows if $D^+ \geq 2D' - 1$.*

Proof. Suppose function values as well as gradients are prescribed at $z_1, \dots, z_{D'}$. We may obtain the prescribed gradients at $z_2, \dots, z_{D'}$ as in the previous proof in the form $\mathbf{p}(z) = \int \mathbf{p}_2 d\pi_2 z + \dots + \int \mathbf{p}_d d\pi_d z$. To obtain suitable function values as well as the π_1 component of the gradients, we may take the polynomial $\mathbf{p}(z) + \mathbf{q}(\pi_1 z)$ with \mathbf{q} being a suitable univariate Hermite interpolant of degree $2D' - 1$. \square

A matrix M is said to be circulant if its subsequent rows are obtained by rotating the first row. If the number of columns is \mathbf{n} and the first row is $\mathbf{a}_1, \dots, \mathbf{a}_n$, the second row must be $\mathbf{a}_n, \mathbf{a}_1, \dots, \mathbf{a}_{n-1}$. The following lemma about circulant matrices will be used in the next section to refine the discussion of [47].

Lemma III.6. *Let M be a $\mathbf{m} \times D'$ circulant matrix whose first row is $1, 0^{j_1}, -1, 0^{j_2}$, where 0^{j_1} is 0 repeated j_1 times. The rank of M is equal to \mathbf{m} if $\mathbf{m} \leq \lceil D'/2 \rceil$.*

Proof. We must have $j_1 + j_2 = D' - 2$. Either j_1 or j_2 must be less than or equal to $(D' - 2)/2$. Because they are both integers, either j_1 or j_2 must be $\leq \lfloor \frac{D'-2}{2} \rfloor$.

Without loss of generality, we assume $j_1 \leq \lfloor D'/2 \rfloor - 1$. As the rows are rotated, the -1 appears in column $j_1 + k + 1$ for $k = 1, \dots, \mathbf{m}$. The columns do not wrap around because

$$j_1 + \mathbf{m} + 1 \leq \lfloor D'/2 \rfloor - 1 + \lceil D'/2 \rceil + 1 \leq D'.$$

All those columns are linearly independent. \square

The final rank lemma is obvious from elementary linear algebra. We state it explicitly because it is invoked often and has a key position in the framework of [47]. For the most part, the lemma is invoked silently.

Lemma III.7. *If the rank of the matrix B is equal to the number of its rows, the rank of the product AB is equal to the rank of A .*

3.4 Review of Sauer et al [47]

In this section, we review the main results and proofs of [47]. Our aim is two-fold. The review helps us prepare the ground for our results about prevalence with a fixed observation map. Second, we point out and fix an error in [47], while presenting the proof with greater formal precision and completeness. The error in [47] is a minor one relative to the depth of ideas found in that paper. We also point out errors and gaps in earlier mathematical treatments that are much more serious.

Let $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a diffeomorphism that is at least C^2 . We adopt the following convention:

Convention about x, y : If x_1 is a point in \mathbb{R}^n , then $x_2 = \phi(x_1)$, $x_3 = \phi(x_2)$, and so on. Similarly, $y_2 = \phi(y_1)$, $y_3 = \phi(y_2)$, and so on. It must be noted that this convention does not apply to z . For example, $z_1, \dots, z_{D'}$ are any distinct points in Lemma III.3.

The observation function is assumed to be the (at least twice continuously differentiable) function $o : \mathbb{R}^n \rightarrow \mathbb{R}$, which maps every state vector to a real number. If the state vector is x_1 , the corresponding delay vector is

$$F_0(x_1) = \begin{pmatrix} o(x_1) \\ \vdots \\ o(x_D) \end{pmatrix},$$

where D will be referred to as the embedding dimension.

Let $K \subset \mathbb{R}^n$ be a possibly fractal set of box counting dimension d . The set K is assumed to be compact. The delay mapping F_0 restricted to K may not be injective. To examine the injectivity more generally, we perturb the observation function to

$$o(x) + \sum_{\alpha \in \mathcal{I}_{2D-1}} c_\alpha p_\alpha(x)$$

and examine injectivity in the ball $\|c_\alpha\| \leq a_0$ with $a_0 > 0$ and fixed. The perturbed delay vector becomes

$$F_\alpha(x) = F_0(x) + \begin{pmatrix} p_\alpha(x_1) \\ \vdots \\ p_\alpha(x_D) \end{pmatrix} (c_\alpha),$$

with α ranging over \mathcal{I}_{2D-1} . We use F_α instead of F_{c_α} to denote the delay vector for simplicity and without risk of confusion. The two assumptions about C_K made in the previous section are carried forward.

Theorem III.8 ([47]). *If $D > 2d$ and ϕ has finitely many periodic points x of periods less than $2D$, the delay mapping $x \rightarrow F_\alpha(x)$ is injective for $x \in K$ for a set of c_α of probability 1 relative to $\|c_\alpha\| \leq a_0$.*

Theorem III.8 is less general than corresponding statements in [47]. Our aim is

to exhibit techniques while forsaking generality. The manner in which more general statements can be obtained is discussed later.

Proof. Define $G_\alpha(x_1, y_1) = F_\alpha(x_1) - F_\alpha(y_1)$. We then have $G_\alpha(x_1, y_1) = F_0(x_1) - F_0(y_1) + \mathcal{M}(c_\alpha)$, where

$$\mathcal{M} = J\mathcal{V}, \quad J = \begin{pmatrix} 1 & & -1 & & \\ & \ddots & & \ddots & \\ & & 1 & & -1 \end{pmatrix}, \quad \mathcal{V} = \begin{pmatrix} p_\alpha(x_1) \\ \vdots \\ p_\alpha(x_D) \\ p_\alpha(y_1) \\ \vdots \\ p_\alpha(y_D) \end{pmatrix}.$$

Here J is $D \times 2D$ and \mathcal{V} is $2D \times D_\alpha$, where D_α is the cardinality of \mathcal{I}_{2D-1} . The proof turns on the determination of the rank of \mathcal{M} . If x_i and y_i , $1 \leq i \leq D$, are $2D$ distinct points, we may apply Lemmas III.3 and III.7 and immediately conclude that the rank of \mathcal{M} is D . However, if not all points are distinct, the rank of \mathcal{V} is obviously not equal to the number of rows. Several cases need to be considered to determine the rank of \mathcal{M} .

Case 1: both x_1 and y_1 are periodic of period less than $2D$ with $x_1 \neq y_1$. The set of such pairs (x_1, y_1) is finite (by assumption) and will be denoted by \mathcal{K}_1 . There are two subcases.

Case 1.1: x_1 and y_1 lie on distinct orbits. If so \mathcal{M} can be written in a compressed

form as $\mathcal{M} = J_c \mathcal{V}_c$ with

$$J_c = \begin{pmatrix} C_1 & C_2 \end{pmatrix}, \quad \mathcal{V}_c = \begin{pmatrix} p_\alpha(x_1) \\ \vdots \\ p_\alpha(x_{\mathbf{p}}) \\ p_\alpha(y_1) \\ \vdots \\ p_\alpha(y_{\mathbf{q}}) \end{pmatrix},$$

where \mathbf{p}, \mathbf{q} are the periods of x_1, y_1 (or D if the periods are greater than D), respectively. Further, C_1 is a $D \times \mathbf{p}$ circulant matrix with first row $1, 0, \dots$ and C_2 is a $D \times \mathbf{q}$ circulant matrix with first row $-1, 0, \dots$. The rank of \mathcal{V}_c is equal to the number of its rows by Lemma III.3 and J_c is nonzero. Therefore, we may assert that the rank of \mathcal{M} is 1 or greater.

Case 1.2: x_1 and y_1 lie on the same periodic orbit. In this case, we may write

$$J_c = C_1, \quad \mathcal{V}_c = \begin{pmatrix} p_\alpha(x_1) \\ \vdots \\ p_\alpha(x_{\mathbf{p}}) \end{pmatrix},$$

where \mathbf{p} is the period of x_1 , C_1 is a $D \times \mathbf{p}$ circulant matrix whose first row is of the form $1, 0, \dots, 0, -1, 0, \dots, 0$. Again, we conclude that the rank of \mathcal{M} is greater than 1.

Suppose $G_\alpha(x_1, y_1) = 0$ for some $(x_1, y_1) \in \mathcal{K}_1$. Then $\mathcal{M}c_\alpha = 0$ and c_α must lie on a hyperplane of co-dimension 1 or greater. Because \mathcal{K}_1 is finite, we may assert $G_\alpha(x_1, y_1) \neq 0$ for all $(x_1, y_1) \in \mathcal{K}_1$ with probability 1 relative to the ball $\|c_\alpha\| \leq a_0$. Case 1 is now complete.

Case 2: Define $\mathcal{K}_2(\delta)$ to be the set of all $(x_1, y_1) \in K \times K$ satisfying

1. $\|x_1 - y_1\| \geq \delta$,

2. $\text{dist}((x_1, y_1), \mathcal{K}_1) \geq \delta$. All distances in this chapter use the L_2 or spectral norm.

The matrix \mathcal{M} has a rank equal to D for each point in $\mathcal{K}_2(\delta)$, as we will prove by breaking up case 2 into subcases.

Case 2.1: $x_1, \dots, x_D, y_1, \dots, y_D$ are $2D$ distinct points. In this case, $\mathcal{M} = J\mathcal{V}$ has rank equal to D as noted at the beginning of the proof.

Case 2.2: x_1, \dots, x_D are distinct, y_1, \dots, y_D are distinct, and neither x_1 nor y_1 is a periodic point of period less than $2D$, but $y_1 = x_j$ or $x_1 = y_j$ for $j \in \{2, \dots, D\}$. Without loss of generality, we assume $y_1 = x_j$.

In this case, the compressed form is $\mathcal{M} = J_c \mathcal{V}_c$ with

$$J_c = (C_1), \quad \mathcal{V}_c = \begin{pmatrix} p_\alpha(x_1) \\ \vdots \\ p_\alpha(x_{D+j-1}) \end{pmatrix},$$

where C_1 is $D \times (D+j-1)$ circulant matrix with first row equal to $1, 0^{j-2}, -1, 0^{D-1}$.

The -1 does not wrap around and the rank of C_1 and therefore of \mathcal{M} is D .

Case 2.3: x_1 periodic of period less than $2D$ and y_1 not so (or vice versa, which may be ignored without loss of generality). In this case, the compressed form is $\mathcal{M} = J_c \mathcal{V}_c$ with

$$J_c = (C_1, C_2), \quad \mathcal{V}_c = \begin{pmatrix} p_\alpha(x_1) \\ \vdots \\ p_\alpha(x_{\mathbf{p}}) \\ p_\alpha(y_1) \\ \vdots \\ p_\alpha(y_D) \end{pmatrix},$$

where \mathbf{p} is the period of x_1 , C_1 is a $D \times \mathbf{p}$ circulant matrix with first row $1, 0, \dots$, and C_2 is a $D \times D$ circulant matrix with first row $-1, 0, \dots$. The column rank of C_2

is equal to D and therefore the rank of \mathcal{M} is also D .

We can now complete case 2 as follows. Suppose $G_\alpha(x_1, y_1) = 0$ for some $(x_1, y_1) \in \mathcal{K}_2(\delta)$. By assumption (2) about C_K , cover $\mathcal{K}_2(\delta)$ with C_K/ϵ^{2d} or fewer ϵ -balls for $\epsilon > 0$. At this point, we introduce an assumption about L :

Assumption about L (1): The Lipschitz constant of $G_\alpha(x_1, y_1)$ with respect to $(x_1, y_1) \in K \times K$ and with $\|c_\alpha\| \leq a_0$ is bounded by L . The existence of L is a consequence of the compactness of $K \times K$, the compactness of $\|c_\alpha\| \leq a_0$, and the differentiability assumption about the observation function o and the diffeomorphism ϕ .

It then follows that if $G_\alpha(x_1, y_1) = 0$ at some point $(x_1, y_1) \in \mathcal{K}_2(\delta)$, then $\|G_\alpha(x_1, y_1)\| = \|\mathcal{M}(c_\alpha)\| \leq L\epsilon$ at the center of one of the ϵ -balls covering $\mathcal{K}_2(\delta)$. Define

$$\sigma_\delta = \min_{(x_1, y_1) \in \mathcal{K}_2(\delta)} \sigma_D(\mathcal{M}).$$

By compactness of $\mathcal{K}_2(\delta)$, σ_δ exists and is positive. By the transfer of volume Lemma III.1, which is applied with $\mathfrak{r} \leftarrow D$, the probability of $\|G_\alpha(x_1, y_1)\| \leq L\epsilon$ relative to the ball $\|c_\alpha\| \leq a_0$ at a point $(x_1, y_1) \in \mathcal{K}_2(\delta)$ is upper bounded by

$$\frac{D_\alpha! L^D \epsilon^D}{\sigma_\delta^D a_0^D}.$$

Because $\mathcal{K}_2(\delta)$ can be covered with C_K/ϵ^{2d} or fewer ϵ -balls, the probability that $G_\alpha(x_1, y_1) = 0$ for some $(x_1, y_1) \in \mathcal{K}_2(\delta)$ is upper bounded by

$$\frac{C_K}{\epsilon^{2d}} \times \frac{D_\alpha! L^D \epsilon^D}{\sigma_\delta^D a_0^D}.$$

Because $D > 2d$ and by taking $\epsilon \rightarrow 0$, we conclude that the probability of $G_\alpha(x_1, y_1) = 0$ for some $(x_1, y_1) \in \mathcal{K}_2(\delta)$ relative to $\|c_\alpha\| \leq a_0$ is one. Case 2 is now complete.

To complete the proof of injectivity, take the union of the measure zero sets in case 2 with $\delta = 1, \frac{1}{2}, \frac{1}{2^2}, \dots$ and the measure zero set in case 1. Outside of that

measure 0 subset of the ball $\|c_\alpha\| \leq a_0$, we have $G_\alpha(x_1, y_1) \neq 0$ for $(x_1, y_1) \in K \times K$ and $x_1 \neq y_1$. \square

The ideas in the proof presented above are from [47], although our presentation is more precise and formally complete. Theorem III.8 makes an assumption on periodic points of period $< 2D$ and not $\leq D$ as in [47]. To see why the more stringent assumption is needed, we turn to [47, p. 611, case 3]. The case “ x and y are not both periodic with period $\leq w$ ” is considered (w is D in our notation) and it is stated that J_{xy} (which is J_c in our notation) is triangular of rank D . Unfortunately, that statement is not correct.

To understand why that statement is not true, assume $D = 6$. Suppose x_1 is a periodic point of period $8 > D$ and that $y_1 = x_5$. Then J_c will be a 6×8 circulant matrix which looks as follows:

$$\begin{pmatrix} 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}.$$

Evidently, the rank of this matrix is $4 < D$.

The easiest way to fix the minor error is to assume the number of periodic points of period $< 2D$ to be finite as we have done. However, Sauer et al [47] place conditions on the box counting dimension of the set of periodic points of period p . The conditions involving quantities such as $\text{rank}(BC_{pq}^w)$ are not easy to interpret and it is unclear what they mean. The basic idea of assuming a bound on the box counting dimension of periodic points of a certain period is a sound one. It can be developed

fully using Lemma III.6 about the rank of circulant matrices and variations of that lemma. We have not done so for two reasons. The proof becomes a great deal more complicated, and at this point having a clear and complete account of the main ideas appears more important than a slightly more general theorem. Additionally, if the box counting dimension of the set of periodic points is greater than 1, then 1 will be a characteristic multiplier that is repeated more than once, which is excluded in the immersivity theorem.

The gaps in [54] and [2] are much less minor. In [54], it is assumed that the delay map is an embedding in some neighborhood of the periodic points. The proof of that assumption is unlikely to be as straightforward as assumed. Even granting that assumption, the argument for transversality [54, p. 371] appears incomplete. In particular, it does not consider the possibility that perturbing the delay map of x may also perturb the delay map of x' , for example, when $x' = \phi(x)$ and the orbits of x, x' overlap. There are yet other aspects of the proof we were not able to verify. For example, [54, p. 370, case iii] seems to require x to be close to a periodic point and x' to be away from a periodic point. It is then asserted that $x, \dots, \phi^{2m}(x), x', \dots, \phi^{2m}(x')$ are distinct. How could that be true if x is a fixed point? How is the possibility $x' = \phi(x)$ handled?

The gaps in [2] also occur in handling overlaps of orbits and periodic points. The main argument [2, p. 598] entirely ignores the possibility that orbits of x^* and y^* may overlap. Further, it is suggested that difficulties associated with fixed points can be handled by adjusting the delays but no details are provided about carrying out that suggestion.

Going back to the work of Sauer et al [47], a point in our proof of Theorem III.8 is worth calling to attention. In the proof, $\mathcal{K}_2(\delta)$ is covered with ϵ -balls and it is

assumed that every ball center is in $\mathcal{K}_2(\delta)$. It is not sufficient to start with any cover of $K \times K$ because a ball center can be arbitrarily close to the diagonal or to a pair of periodic points and $\sigma_D(\mathcal{M})$ may become arbitrarily small.

If we say that a certain compact set \mathfrak{s} is covered by a certain number of ϵ -balls, it is assumed that each ball has a center that lies in \mathfrak{s} . That assumption comes up repeatedly in the proof of immersivity, which we now turn to. Once again all the ideas are from [47]. Here K is assumed to be a smooth, closed, and compact submanifold of dimension d and T_1K denotes its unit tangent bundle. If $x \in K$ and v is tangent to K at x , then $(x, v) \in T_1K$ if and only if $\|v\| = 1$.

Theorem III.9. [47] *If $D > 2d - 1$, K is invariant under ϕ , ϕ has finitely many points $x \in K$ of period less than D , and all characteristic multipliers of each of those points are distinct, then $x \rightarrow F_\alpha(x)$ is immersive over K with probability 1 relative to the ball $\|c_\alpha\| \leq a_0$.*

Proof. If $x \rightarrow F_\alpha(x)$ and v is a tangent vector to K at x , then we denote the vector that v is mapped to by $dF_\alpha(x, v)$. The following convention about v is an extension of the convention about x, y explained earlier.

Convention about v : If v_1 is tangent to K at x_1 , then $v_2 = \frac{\partial \phi}{\partial x} \Big|_{x_1} v_1$, $v_3 = \frac{\partial \phi}{\partial x} \Big|_{x_2} v_2$, and so on. Because ϕ is a diffeomorphism, v_i are all nonzero like v_1 .

We write $dF_\alpha(x_1, v_1) = dF_0(x_1, v_1) + \mathcal{N}(c_\alpha)$, where

$$\mathcal{N} = J\mathcal{H}, \quad J = \begin{pmatrix} v_1^T & & \\ & \ddots & \\ & & v_D^T \end{pmatrix}, \quad \mathcal{H} = \begin{pmatrix} \nabla p_\alpha(x_1) \\ \vdots \\ \nabla p_\alpha(x_D) \end{pmatrix}.$$

The proof will turn on the rank of $\mathcal{N} = J\mathcal{H}$. If x_1, \dots, x_D are distinct, the rank of \mathcal{N} is D because the rank of \mathcal{H} is equal to the number of its rows by Lemma III.4

and the rank of J is obviously D .

To study the rank of \mathcal{N} , it is useful to define the following disjoint sets of T_1K .

- \mathcal{K}_1 is the set of all (x_1, v_1) such that x_1 is a periodic point of period less than D and v_1 is an eigenvector of the periodic point x_1 . By eigenvector of a periodic point, we mean an eigenvector of the corresponding monodromy matrix.
- $\mathcal{K}_2(\delta)$ is the set of all (x_1, v_1) such that x_1 is a periodic point of period less than D and v_1 is a linear combination of two eigenvectors of x_1 . It is also required that

$$\text{dist}((x_1, v_1), \mathcal{K}_1) \geq \delta.$$

We will denote $\mathcal{K}_2(0)$, where this last condition is not operative, by \mathcal{K}_2 . Evidently, \mathcal{K}_1 is a subset of \mathcal{K}_2 .

- In general, $\mathcal{K}_r(\delta)$, where $r = 2, \dots, d$, is defined as the set of $(x_1, v_1) \in T_1K$ such that x_1 is a periodic point of period D or less and v_1 is a linear combination of r eigenvectors of the periodic point x_1 . It is also required that

$$\text{dist}((x_1, v_1), \mathcal{K}_{r-1}) \geq \delta.$$

We will denote $\mathcal{K}_r(0)$, where this last condition is not operative, by \mathcal{K}_r . Evidently, \mathcal{K}_{r-1} is a subset of \mathcal{K}_r .

This sequence of cases stops at $r = d$ and does not go up to $r = n$ because we are only interested in those eigenvectors of the periodic point x_1 that are also tangent to K . The assumption about the invariance of K is used here.

- The final case is $\mathcal{K}_D(\delta)$ which consists of all points $(x_1, v_1) \in T_1K$ such that x_1 is not periodic of period less than D and the distance to \mathcal{K}_d is $\geq \delta$.

is $\min(rp, D)$ because the Vandermonde matrix

$$\begin{pmatrix} 1 & 1 & \dots & 1 \\ \lambda_1 & \lambda_2 & \dots & \lambda_r \\ & \vdots & & \\ \lambda_1^{r-1} & \lambda_2^{r-1} & & \lambda_r^{r-1} \end{pmatrix}$$

has full rank, the λ_i being distinct by assumption. Therefore, the rank of \mathcal{N} is r or greater for each $(x_1, y_1) \in \mathcal{K}_r(\delta)$.

To complete the proof, we note that \mathcal{K}_r is of dimension $r - 1$ for $r = 1, \dots, d$ and that T_1K is of dimension $2d - 1$. A new assumption about C_K is useful.

Assumption about C_K (3): It is assumed that \mathcal{K}_r can be covered with C_K/ϵ^{r-1} ϵ -balls for $r = 1, \dots, d$. It is assumed that T_1K and therefore $\mathcal{K}_D(\delta)$ can be covered with C_K/ϵ^{2d-1} ϵ -balls.

We also extend the assumption about the Lipschitz bound L .

Assumption about L (2): It is assumed that the Lipschitz constant of dF_α with respect to $(x_1, v_1) \in T_1K$ for $\|c_\alpha\| \leq a_0$ is upper bounded by L . This assumption too may be verified using compactness like the first assumption about L .

The proof may now be completed easily. Suppose $dF_\alpha(x_1, v_1) = 0$ for some $(x_1, v_1) \in \mathcal{K}_r(\delta)$. Then $\|dF_\alpha(x_1, v_1)\| \leq L\epsilon$ at the center of one of the C_K/ϵ^{r-1} balls covering $\mathcal{K}_r(\delta)$. By the transfer of volume Lemma III.1, the probability of such an event is upper bounded by

$$\frac{C_K}{\epsilon^{r-1}} \times \frac{D_\alpha! L^r \epsilon^r}{\sigma_\delta^r a_0^r},$$

where $\sigma_\delta = \min \sigma_r(\mathcal{N})$ over $(x_1, v_1) \in \mathcal{K}_r(\delta)$. The probability evidently goes to 0 as $\epsilon \rightarrow 0$ leaving us with a measure zero set of c_α where F_α is not immersive at some point in $\mathcal{K}_r(\delta)$ for $r = 2, \dots, d$. The sets \mathcal{K}_1 and $\mathcal{K}_D(\delta)$ are handled similarly. \square

Theorem III.9 assumes K to be a closed and compact submanifold. That assumption implies $\mathcal{K}_D(\delta)$ to be compact. If $\mathcal{K}_D(\delta)$ is compact, we may conclude that $\min \sigma_D(\mathcal{N})$ over $(x_1, v_1) \in \mathcal{K}_D(\delta)$ exists and is positive. The assumptions on K can be reduced. However, the technicalities that arise (see [24]) are extraneous to the main ideas in this chapter.

3.5 Perturbing the dynamical system

Let $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a diffeomorphism, which is as before but with $n = d$. Let $\psi(x)$ denote $\frac{\partial \phi}{\partial x}$. The vector in \mathbb{R}^d with first component 1 and the others zero is denoted by \mathbf{e}_1 . The perturbed dynamical system is

$$\phi_\alpha(x) = \phi(x) + \mathbf{e}_1 (p_\alpha(x)) (c_\alpha)$$

with $\alpha \in \mathcal{I}_{2D-1}$, where D is the embedding dimension. It may be noted we are only perturbing the first coordinate of ϕ . Because the observation function will be assumed to be $o = \pi_1$, it is enough to perturb only the first coordinate.

The delay vector under ϕ is

$$F_0(x_1) = \begin{pmatrix} \pi_1 x_1 \\ \vdots \\ \pi_1 x_D \end{pmatrix}.$$

Convention about \tilde{x} : It is assumed that $\tilde{x}_1 = x_1$. Thereafter, it is assumed that

$$\tilde{x}_2 = \phi_\alpha(\tilde{x}_1), \tilde{x}_3 = \phi_\alpha(\tilde{x}_2), \text{ and so on.}$$

The delay vector under ϕ_α is therefore

$$F_\alpha(x_1) = \begin{pmatrix} \pi_1 \tilde{x}_1 \\ \vdots \\ \pi_1 \tilde{x}_D \end{pmatrix}.$$

It is worthy of notice that ϕ_α perturbs only the first component of ϕ . Because the delay vector is built up using π_1 , ϕ_α must perturb the first component. If not, the perturbation may not propagate to the delay vector at all. It turns out that perturbing only the first component is also sufficient to obtain a prevalence theorem.

Our first task is to express F_α as a perturbation of F_0 . That can be done by simply iterating the definition of ϕ_α :

$$\begin{aligned}\tilde{x}_1 &= x_1 \\ \tilde{x}_2 &= x_2 + \mathbf{e}_1(p_\alpha(x_1))(c_\alpha) \\ \tilde{x}_3 &= x_3 + \mathbf{e}_1(p_\alpha(x_2))(c_\alpha) + \psi(x_2)\mathbf{e}_1(p_\alpha(x_1))(c_\alpha) + \mathcal{O}(c_\alpha^2).\end{aligned}$$

Above and later, $\mathcal{O}(c_\alpha^2)$ is the same as $\mathcal{O}(\|c_\alpha\|^2)$. By following the pattern, we obtain

(3.6)

$$\tilde{x}_j = x_j + \mathbf{e}_1(p_\alpha(x_{j-1}))(c_\alpha) + \rho_{j-1}(x_2, \dots, x_{j-1}, p_\alpha(x_1), \dots, p_\alpha(x_{j-2}))(c_\alpha) + \mathcal{O}(c_\alpha^2)$$

for $j = 2, \dots, D$. Here it is important to note that ρ_{j-1} is linear in $p_\alpha(x_1), \dots, p_\alpha(x_{j-2})$.

For brevity, we will rewrite (3.6) as

$$(3.7) \quad \tilde{x}_j = x_j + \mathbf{e}_1(p_\alpha(x_{j-1}))(c_\alpha) + \rho_{j-1}(c_\alpha) + \mathcal{O}(c_\alpha^2).$$

We then get

$$F_\alpha(x_1) = F_0(x_1) + \begin{pmatrix} 0 \\ V(x_1) \end{pmatrix} (c_\alpha) + \mathcal{O}(c_\alpha^2),$$

with the matrix $V(x_1)$ defined by

$$V(x_1) = \begin{pmatrix} p_\alpha(x_1) \\ p_\alpha(x_2) + \pi_1 \rho_2 \\ \vdots \\ p_\alpha(x_{D-1}) + \pi_1 \rho_{D-1} \end{pmatrix}.$$

The next lemma is about the rank of $V(x_1)$.

Lemma III.10. *If x_1, \dots, x_{D-1} are distinct, the rank of $V(x_1)$ is equal to the number of its rows.*

Proof. Suppose we consider

$$\mathcal{V} = \begin{pmatrix} p_\alpha(x_1) \\ p_\alpha(x_2) \\ \vdots \\ p_\alpha(x_{D-1}) \end{pmatrix}.$$

The rank lemma III.3 tells us that the rank of \mathcal{V} is equal to the number of its rows.

Now to produce a vector $(\mathbf{a}_1, \dots, \mathbf{a}_{D-1})^T$ in the range of $V(x_1)$, we proceed as follows. Define

$$\mathbf{a}'_1 = \mathbf{a}_1$$

$$\mathbf{a}'_2 = \mathbf{a}_2 - \rho(x_1, \mathbf{a}'_1)$$

$$\mathbf{a}'_3 = \mathbf{a}_3 - \rho(x_1, x_2, \mathbf{a}'_1, \mathbf{a}'_2)$$

and so on. Because of the linearity of ρ in \mathbf{a}_i , the vector (c_α) that satisfies $\mathcal{V}(c_\alpha) = (\mathbf{a}'_1, \dots, \mathbf{a}'_{D-1})^T$ also satisfies $V(x_1)(c_\alpha) = (\mathbf{a}_1, \dots, \mathbf{a}_{D-1})^T$. \square

The next lemma is similar. Part (c) of the following lemma is more general than Lemma III.10 because we allow $D^+ > D$.

Lemma III.11. *The following matrices have rank equal to the number of rows:*

1. $\begin{pmatrix} V(x_1) \\ V(y_1) \end{pmatrix}$ assuming $x_1, \dots, x_{D-1}, y_1, \dots, y_{D-1}$ to be distinct.
2. $\begin{pmatrix} V(x_1) \\ \mathbf{m}_k \end{pmatrix}$, where \mathbf{m}_k is the first k rows of $V(y_1)$, assuming $x_1, \dots, x_{D-1}, y_1, \dots, y_k$ to be distinct.

$$3. \left(\begin{array}{c} p_\alpha(x_1) \\ p_\alpha(x_2) + \pi_1 \rho_2 \\ \vdots \\ p_\alpha(x_{D^+-1}) + \pi_1 \rho_{D^+-1} \end{array} \right) \text{ assuming } x_1, \dots, x_{D^+-1} \text{ are distinct and } D^+ \leq 2D.$$

Proof. Similar to the previous proof. \square

Our second task in this section is to obtain $dF_\alpha(x_1, v_1)$ as a perturbation of $dF_0(x_1, v_1)$.

It is helpful to introduce another convention:

Convention about w : $w_1 = v_1$, w_2 is obtained as $\left. \frac{\partial \phi_\alpha}{\partial x} \right|_{\tilde{x}_1} w_1$, w_3 is obtained as

$$\left. \frac{\partial \phi_\alpha}{\partial x} \right|_{\tilde{x}_2} w_2, \text{ and so on.}$$

Thus, in effect we need to obtain perturbative expansions of w_i . To do so, let us first note that

$$\frac{\partial \phi_\alpha}{\partial x} = \psi(x) + \mathbf{e}_1 (\nabla p_\alpha(x)^T) (c_\alpha).$$

We substitute the above equation into the iteration that defines w_i and obtain

$$w_1 = v_1$$

$$w_2 = \psi(\tilde{x}_1)w_1 + \mathbf{e}_1 (v_1^T \nabla p_\alpha(x_1)) (c_\alpha)$$

$$w_3 = \psi(\tilde{x}_2)w_2 + \mathbf{e}_1 (v_2^T \nabla p_\alpha(x_2)) (c_\alpha) + \varrho_2 + \mathcal{O}(c_\alpha^2)$$

and so on. If we now use (3.6) to substitute for \tilde{x}_j , we obtain

$$w_j = v_j + \mathbf{e}_1 (v_{j-1}^T \nabla p_\alpha(x_{j-1})) (c_\alpha) + \varrho_{j-1}(c_\alpha) + \mathcal{O}(c_\alpha^2),$$

where ρ_{j-1} is linear in

$$p_\alpha(x_1), \dots, p_\alpha(x_{j-2}), \nabla p_\alpha(x_1), \dots, \nabla p_\alpha(x_{j-2}).$$

We may then write

$$\begin{aligned} dF_\alpha(x_1, v_1) &= \begin{pmatrix} \pi_1 w_1 \\ \vdots \\ \pi_1 w_D \end{pmatrix} \\ &= dF_0(x_1, v_1) + \begin{pmatrix} 0 \\ H(x_1, v_1) \end{pmatrix} (c_\alpha) + \mathcal{O}(c_\alpha^2), \end{aligned}$$

where

$$H(x_1, v_1) = \begin{pmatrix} v_1^T \nabla p_\alpha(x_1) \\ v_2^T \nabla p_\alpha(x_2) + \pi_1 \varrho_2 \\ v_3^T \nabla p_\alpha(x_3) + \pi_1 \varrho_3 \\ \vdots \\ v_{D-1}^T \nabla p_\alpha(x_{D-1}) + \pi_1 \varrho_{D-1} \end{pmatrix}.$$

The second task for this section concludes with a lemma about the rank of $H(x_1, v_1)$.

Lemma III.12. *If x_1, \dots, x_{D-1} are distinct, the rank of $H(x_1, v_1)$ is equal to the number of its rows.*

Proof. The proof is similar to that of Lemma III.10. First consider

$$\begin{pmatrix} p_\alpha(x_1) \\ \vdots \\ p_\alpha(x_{D-1}) \\ v_1^T \nabla p_\alpha(x_1) \\ \vdots \\ v_{D-1}^T \nabla p_\alpha(x_{D-1}) \end{pmatrix}.$$

By Lemma III.5, the rank of this matrix is equal to the number of its rows. Suppose we want to find (c_α) such that $H(x_1, v_1)(c_\alpha)$ equals a specified vector $(\mathbf{a}_1, \dots, \mathbf{a}_{D-1})^T$.

To do so, we find a vector (c_α) such that the matrix displayed above applied to c_α is equal to

$$\begin{pmatrix} 0 \\ \vdots \\ 0 \\ \mathbf{a}'_1 \\ \vdots \\ \mathbf{a}'_{D-1} \end{pmatrix},$$

where $\mathbf{a}'_1 = \mathbf{a}_1$, $\mathbf{a}'_2 = \mathbf{a}_2 - \mathbf{r}_2$, where \mathbf{r}_2 is $\pi \varrho_2$ evaluated by replacing $p_\alpha(x_1)$ by 0 and $v_1^T \nabla p_\alpha(x_1)$ by \mathbf{a}'_1 , and so on. \square

The third and final task of this section is to track the perturbation of fixed points when the map ϕ is perturbed to ϕ_α .

Lemma III.13. *Suppose $z_0 = \phi(z_0)$ and $\psi(z_0)$ has no eigenvalue equal to 1. Under $\phi \rightarrow \phi_\alpha$, the fixed point z_0 perturbs to*

$$z_0(c_\alpha) = z_0 + (I - \psi(z_0))^{-1} \mathbf{e}_1(p_\alpha(z_0))(c_\alpha) + \mathcal{O}(c_\alpha^2).$$

Proof. The function $z_0(c_\alpha)$ exists by the implicit function theorem. To obtain the expansion given in the lemma, start with

$$\phi(z_0) + \mathbf{e}_1(p_\alpha(z_0))(c_\alpha) = z_0$$

differentiate with respect to c_α and obtain $\frac{\partial z_0}{\partial c_\alpha}$ at $c_\alpha = 0$ using implicit differentiation. \square

3.6 The setting for injectivity and immersivity theorems

In the case where ϕ is fixed and only the observation function o is perturbed, injectivity and immersivity are proved with respect to the ball $\|c_\alpha\| \leq a_0$, where $a_0 > 0$ can be anything. Such a thing is plainly impossible when ϕ is perturbed to ϕ_α . Under a perturbation, the map may even fail to be well defined or might blow-up in finite time. Therefore, we have to specify the setting for injectivity and immersivity theorems more carefully.

We will assume that K is a compact sphere in \mathbb{R}^d centered at the origin. The map ϕ_α will be proved to be injective and immersive over K . It is assumed that K^+ is a compact sphere bigger than K and containing K . If $x_1 \in K$, it is assumed that x_1, \dots, x_D all remain in K^+ . In addition, a_0 is assumed to be so small that $\tilde{x}_1, \dots, \tilde{x}_D$ all remain in K^+ for all $\|c_\alpha\| \leq a_0$. Further assumptions are enumerated below:

1. $\phi_\alpha : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is assumed to be a diffeomorphism (for $\|c_\alpha\| \leq a_0$), that is C^3 or better.
2. The map ϕ_α has exactly m fixed points, denoted by $\xi_1(c_\alpha), \dots, \xi_m(c_\alpha)$.
3. The map ϕ_α has no other periodic points of period less than $2D$.
4. All the fixed points are hyperbolic and $\pi_1 \xi_i(c_\alpha) \neq \pi_1 \xi_j(c_\alpha)$ if $i \neq j$. This assumption is made with the intention of simplifying the proof so as to bring out the main techniques with greater clarity. Here we are essentially assuming injectivity between fixed points.
5. We will also assume that dF_α is immersive at each fixed point for the same reason.

Now we will recall a few basic facts about Lebesgue points. A point $\mathbf{a} \in \mathbb{R}^n$ is a Lebesgue point of a measurable set $A \subset \mathbb{R}^n$ if

$$\lim_{\epsilon \rightarrow 0} \frac{\mu(A \cap \{u \mid \|u - \mathbf{a}\| < \epsilon\})}{\mu(\{u \mid \|u - \mathbf{a}\| < \epsilon\})} = 1.$$

We will need the following basic lemma.

Lemma III.14. *If every point of the measurable set B is a Lebesgue point of the measurable set A , then $\mu(B - A) = 0$.*

Proof. Almost every point of A is a Lebesgue point of A [45]. Similarly, almost every point of A^c , the complement of A , is a Lebesgue point of A^c . If \mathbf{a} is a Lebesgue point of A^c ,

$$\lim_{\epsilon \rightarrow 0} \frac{\mu(A \cap \{u \mid \|u - \mathbf{a}\| < \epsilon\})}{\mu(\{u \mid \|u - \mathbf{a}\| < \epsilon\})} = 0.$$

The lemma follows from these observations. \square

Lemma III.14 will be crucial to our proof that ϕ_α is an embedding with probability 1 relative to $\|c_\alpha\| < a_0$. In the case where ϕ is fixed and only the observation function is perturbed, the proofs of injectivity and immersivity consider the ball $\|c_\alpha\| \leq a_0$ all at once. Such a thing is not possible here. Instead, we have to pick c_α^* satisfying $\|c_\alpha^*\| < a_0$ and localize around it and that is where Lemma III.14 comes in.

In order to localize around c_α^* , we adopt new notation that is centered at c_α^* . The re-centered diffeomorphism $\phi(x) + \mathbf{e}_1(p_\alpha(x))(c_\alpha^*)$ is denoted by $\Phi(x)$. Similarly, Ψ denotes $\psi(x) + \mathbf{e}_1(\nabla p_\alpha(x))(c_\alpha^*)$. When we localize around c_α^* , $\Phi_\alpha(x)$ will denote $\Phi(x) + \mathbf{e}_1(p_\alpha(x))(c_\alpha)$. The fixed point $\xi_j(c_\alpha^*)$ is denoted Σ_j . The fixed point $\xi_j(c_\alpha^* + c_\alpha)$ is denoted $\Sigma_j(c_\alpha)$.

Convention about x, y updated: x_1, x_2, \dots are iterates of x_1 under Φ . Similarly,

y_1, y_2, \dots are iterates of y_1 under Φ .

Convention about \tilde{x} updated: $\tilde{x}_1 = x_1$ and $\tilde{x}_1, \tilde{x}_2, \dots$ are iterates of x_1 under Φ_α .

Convention about v updated: we assume $(x_1, v_1) \in T_1K$ and v_2, v_3, \dots are obtained by iterating $d\Phi$.

All the lemmas of the previous section continue to hold after re-centering. The delay vector $F_\alpha(x)$ defined in the previous section will be denoted by $\mathbb{F}_0(x)$ if c_α is replaced by c_α^* . Similarly, if c_α is replaced by $c_\alpha^* + c_\alpha$ in the definition of $F_\alpha(x)$, we will denote the re-centered delay vector by $\mathbb{F}_\alpha(x)$.

We may write

$$\mathbb{F}_\alpha(x_1) = \mathbb{F}_0(x_1) + \begin{pmatrix} 0 \\ \mathbb{V}(x_1) \end{pmatrix} (c_\alpha) + \mathcal{O}(c_\alpha^2),$$

with the definition of $\mathbb{V}(x_1)$ being the same as that of $V(x_1)$ but with ψ replaced by Ψ . Likewise,

$$d\mathbb{F}_\alpha(x_1, v_1) = d\mathbb{F}_0(x_1, v_1) + \begin{pmatrix} 0 \\ \mathbb{H}(x_1, v_1) \end{pmatrix} (c_\alpha) + \mathcal{O}(c_\alpha^2),$$

with a similar alteration of the definition of $H(x_1, v_1)$ to get $\mathbb{H}(x_1, v_1)$.

Finally, we note that the centered analogue of $G_\alpha(x_1, y_1) = F_\alpha(x_1) - F_\alpha(y_1)$ is $\mathbb{G}_\alpha(x_1, y_1) = \mathbb{F}_\alpha(x_1) - \mathbb{F}_\alpha(y_1)$.

3.7 Proof of injectivity

In this section, our purpose is to prove that $F_\alpha(x_1)$, defined in section 5, is injective for $x_1 \in K$. The assumptions about C_K and L are carried forward from earlier sections, although the third assumption about C_K is not necessary in its entirety. Further assumptions will be stated as the need arises. Let us define Δ is the minimum distance between fixed points of F_α in K for $\|c_\alpha\| \leq a_0$.

Let us define $\mathcal{A}_{1,\delta}$ to be the set of c_α satisfying

Define $\sigma_\delta = \min \sigma_{D-1}(\mathcal{M})$, where the minimum is over $x_1 \in \mathcal{K}_{1,\delta}$ and $\|c_\alpha\| \leq a_1$. Cover $\mathcal{K}_1(\delta)$ with C_K/ϵ^d ϵ -balls.

Assumption about L (3): In (3.8), the $\mathcal{O}(c_\alpha^2)$ term is upper bounded by $L\|c_\alpha\|^2$.

Like the earlier assumptions about L , this assumption too is a direct consequence of compactness. The earlier assumptions used L as a bound on Lipschitz constants. Here L is used as a bound on the Taylor series remainder.

Now suppose $\mathbb{G}_\alpha(\Sigma_j(c_\alpha), x_1) = 0$ for some $j \in \{1, \dots, m\}$ and some $x_1 \in \mathcal{K}_{1,\delta}$. Because the Lipschitz constant of $\mathbb{G}_\alpha(\Sigma_j(c_\alpha), x_1)$ with respect to x_1 is bounded by L , we must have $\|\mathbb{G}_\alpha(\Sigma_j(c_\alpha), x_1)\| \leq L\epsilon$ at an x_1 that is at the center of one of the balls covering $\mathcal{K}_{1,\delta}$.

Applying the nonlinear transfer of volume Lemma III.2 with $\mathfrak{r} \leftarrow D-1$ and $\sigma \leftarrow \sigma_\delta$, we find that the probability of $\|\mathbb{G}_\alpha(\Sigma_j(c_\alpha), x_1)\| \leq L\epsilon$ relative to $\|c_\alpha\| \leq \epsilon^{1/2} < a_1$ is upper bounded by

$$D_\alpha! 2^{D-1} L^{D-1} \epsilon^{(D-1)/2} / \sigma_\delta^{D-1}.$$

Because the number of fixed points is m and the number of balls covering $\mathcal{K}_{1,\delta}$ is C_K/ϵ^d , the probability of $\mathbb{G}_\alpha(\Sigma_j(c_\alpha), x_1) = 0$ for some $j \in \{1, \dots, m\}$ and some $x_1 \in \mathcal{K}_{1,\delta}$ relative to $\|c_\alpha\| \leq \epsilon^{1/2}$ is upper bounded by

$$m \times \frac{C_K}{\epsilon^d} \times \frac{D_\alpha! 2^{D-1} L^{D-1} \epsilon^{(D-1)/2}}{\sigma_\delta^{D-1}}.$$

Evidently, the probability goes to zero as $\epsilon \rightarrow 0$ if $D \geq 2d+2$. Thus, we have shown that c_α^* is a Lebesgue point of $\mathcal{A}_{1,\delta}$ proving the lemma. \square

Now define $\mathcal{A}_{2,\delta}$ to be the set of c_α satisfying

1. $\|c_\alpha\| < a_0$
2. $G_\alpha(x_1, \phi_\alpha(x_1)) \neq 0$ for $x_1 \in K$ with $\|x_1 - \xi_j(\alpha)\| \geq 3\delta$ for each $j \in \{1, \dots, m\}$.

Lemma III.16. *If $D \geq 2d + 1$, every point of $\|c_\alpha\| < a_0$ is a Lebesgue point of $\mathcal{A}_{2,\delta}$ and therefore the probability of $\mathcal{A}_{2,\delta}$ relative to $\|c_\alpha\| < a_0$ is 1.*

Proof. As before, we pick c_α^* satisfying $\|c_\alpha^*\| < a_0$ and will give an argument centered at c_α^* to show that c_α^* is a Lebesgue point of $\mathcal{A}_{1,\delta}$. As before, pick $a_1 > 0$ so small that $\|\Sigma_j(\alpha) - \Sigma_j\| < \delta$ for $\|c_\alpha\| \leq a_1$. As before, define $\mathcal{K}_{1,\delta}$ as the set of $x_1 \in K$ such that $\|x_1 - \Sigma_j\| \geq 2\delta$ for each $j \in \{1, \dots, m\}$.

Using (3.7), we get

$$(3.9) \quad \mathbb{G}_\alpha(\tilde{x}_1, \tilde{x}_2) = \begin{pmatrix} \pi_1 x_1 - \pi_1 x_2 \\ \vdots \\ \pi_1 x_D - \pi_1 x_{D+1} \end{pmatrix} + \mathcal{M}(c_\alpha) + \mathcal{O}(c_\alpha^2)$$

with $\mathcal{M} = J\mathcal{V}$ and

$$J = \begin{pmatrix} -1 & & & \\ & 1 & -1 & \\ & & 1 & -1 \\ & & & 1 & -1 \end{pmatrix}, \quad \mathcal{V} = \begin{pmatrix} p_\alpha(x_1) \\ p_\alpha(x_2) + \pi_1 \rho_2 \\ \vdots \\ p_\alpha(x_D) + \pi_1 \rho_D \end{pmatrix}.$$

By Lemma III.11 (c), the rank of \mathcal{V} is equal to the number of its rows. Therefore, the rank of $\mathcal{M} = J\mathcal{V}$ is equal to D .

Define $\sigma_\delta = \min \sigma_D(\mathcal{M})$, where the minimum is over $x_1 \in \mathcal{K}_{1,\delta}$ and $\|c_\alpha\| \leq a_1$. Cover $\mathcal{K}_1(\delta)$ with C_K/ϵ^d ϵ -balls.

Assumption about L (4): In (3.9), the $\mathcal{O}(c_\alpha^2)$ term is upper bounded by $L\|c_\alpha^2\|$.

The first two assumptions about L are both obtained from upper bounds on the derivative of $F_\alpha(x)$ or $\mathbb{F}_\alpha(x)$ with respect to x . This assumption as well as the preceding one are obtained from upper bounds on the second derivative. In all cases, the assumptions are direct consequences of the compactness of K and the ball $\|c_\alpha\| \leq a_0$.

If $\mathbb{G}_\alpha(\tilde{x}_1, \tilde{x}_2) = 0$ for some $x_1 \in \mathcal{K}_{1,\delta}$, we must have $\|\mathbb{G}_\alpha(\tilde{x}_1, \tilde{x}_2)\| \leq L\epsilon$ for some x_1 that is the center of one of the balls covering $\mathcal{K}_{1,\delta}$. Using the nonlinear transfer of volume Lemma III.2, we find the probability of $\mathbb{G}_\alpha(\tilde{x}_1, \tilde{x}_2) = 0$ for some $x_1 \in \mathcal{K}_{1,\delta}$ relative to the ball $\|c_\alpha\| \leq \epsilon^{1/2} < a_1$ to be upper bounded by

$$\frac{C_K}{\epsilon^d} \times \frac{D_\alpha! 2^D L^D \epsilon^{D/2}}{\sigma_\delta^D}.$$

The limit of this probability as $\epsilon \rightarrow 0$ is zero. It follows that c_α^* is a Lebesgue point of $\mathcal{A}_{2,\delta}$ completing the proof of this lemma. \square

Lemma III.16 allows us to conclude that the delay vectors of x_1 and $\phi_\alpha(x_1)$ do not coincide typically if x_1 is a little removed from the fixed points of ϕ_α . More generally, we need to argue that the delay vectors of x_1 and $\phi_\alpha^{k-1}(x)$ do not coincide for $k = 3, \dots, D$. To make that argument, we define $\mathcal{A}_{k,\delta}$ to be the set of c_α satisfying

1. $\|c_\alpha\| < a_0$
2. $G_\alpha(x_1, \phi_\alpha^{k-1}(x_1)) \neq 0$ for $x_1 \in K$ with $\|x_1 - \xi_j(\alpha)\| \geq 3\delta$ for each $j \in \{1, \dots, m\}$

for $k = 2, \dots, D$.

Lemma III.17. *For $D \geq 2d + 1$ and $k = 2, \dots, D$, every point of $\|c_\alpha\| < a_0$ is a Lebesgue point of $\mathcal{A}_{k,\delta}$ and therefore the probability of $\mathcal{A}_{k,\delta}$ relative to the ball $\|c_\alpha\| < a_0$ is 1.*

Proof. The proof is almost identical to that of the previous lemma, which is a special case. The only significant difference occurs in the definition of \mathcal{V} . In the general case,

$$\mathcal{V} = \begin{pmatrix} p_\alpha(x_1) \\ p_\alpha(x_2) + \pi_1 \rho_2 \\ \vdots \\ p_\alpha(x_{D+k-2}) + \pi_1 \rho_{D+k-2} \end{pmatrix}.$$

Note that Lemma III.11 (c) still applies, implying the rank of \mathcal{V} to be equal to the number of its rows, because $D + k - 2 \leq 2D$. \square

The final lemma of this section pertains to the set $\mathcal{A}_{xy}(\delta)$. It is defined as the set of all c_α such that $\|c_\alpha\| < a_0$ and $G_\alpha(x_1, y_1) \neq 0$ provided

1. $x_1, y_1 \in K$
2. $\|x_1 - y_1\| \geq \delta$ (which excludes the diagonal of $K \times K$)
3. $\|x_1 - \xi_j(\alpha)\| \geq 3\delta$ and $\|y_1 - \xi_j(\alpha)\| \geq 3\delta$ for $j \in \{1, \dots, m\}$ (so that both x_1 and y_1 stay away from fixed points)
4. $\|x_1 - \phi^{k-1}(y_1)\| \geq 2\delta$ and $\|y_1 - \phi^{k-1}(x_1)\| \geq 2\delta$ for $k = 2, \dots, D$ (so that x_1 does not come too close to the iterates of y_1 and vice versa).

Lemma III.18. *For $D \geq 4d + 2$, every point of $\|c_\alpha\| < a_0$ is a Lebesgue point of $\mathcal{A}_{xy,\delta}$ and therefore the probability of $\mathcal{A}_{xy,\delta}$ relative to the ball $\|c_\alpha\| < a_0$ is 1.*

Proof. Again the argument begins by centering at some c_α^* satisfying $\|c_\alpha^*\| < a_0$. However, the conditions on a_1 this time are different. The radius a_1 must be so small that for $\|c_\alpha\| \leq a_1$ the following conditions are satisfied:

1. $\|\Sigma_j(\alpha) - \Sigma_j\| < \delta$
2. For any $x_1 \in K$, $\|\tilde{x}_j - x_j\| \leq \delta$ for $j = 1, \dots, D$.

The set $\mathcal{K}_{xy,\delta}$ is defined as the set of $(x_1, y_1) \in K \times K$ satisfying the following conditions:

1. $\|x_1 - \Sigma_j\| \geq 2\delta$ and $\|y_1 - \Sigma_j\| \geq 2\delta$ for $j \in \{1, \dots, m\}$
2. $\|x_1 - y_1\| \geq \delta$

3. $\|x_1 - y_j\| \geq \delta$ and $\|y_1 - x_j\| \geq \delta$ for $j \in \{2, \dots, m\}$.

We have

$$(3.10) \quad \mathbb{G}_\alpha(x_1, y_1) = \mathbb{G}_0(x_1, y_1) + \mathcal{M}(c_\alpha) + \mathcal{O}(c_\alpha^2).$$

The top row of \mathcal{M} is zero. The rest of the $D - 1$ rows below are given by $J\mathcal{V}$

$$J = \begin{pmatrix} 1 & & -1 & \\ & \ddots & & \\ & & 1 & -1 \end{pmatrix}, \quad \mathcal{V} = \begin{pmatrix} \mathbb{V}(x_1) \\ \mathbb{V}(y_1) \end{pmatrix}.$$

By Lemma III.11, the rank of \mathcal{V} is equal to the number of its rows. Therefore the ranks of $J\mathcal{V}$ and \mathcal{M} are both equal to $D - 1$.

Define $\sigma_\delta = \min \sigma_{D-1}(\mathcal{M})$, where the minimum is over $(x_1, y_1) \in \mathcal{K}_{xy,\delta}$ and $\|c_\alpha\| \leq a_1$. Cover \mathcal{K}_{xy} with C_K/ϵ^{2d} balls.

Assumption about L (5): The $\mathcal{O}(c_\alpha^2)$ term in (3.10) is upper bounded by $L\|c_\alpha\|^2$.

Suppose $\mathbb{G}_\alpha(x_1, y_1) = 0$ for some $(x_1, y_1) \in \mathcal{K}_{xy,\delta}$. Then we must have $\|\mathbb{G}_\alpha(x_1, y_1)\| \leq L\epsilon$ for an (x_1, y_1) that is at the center of one of the balls covering $\mathcal{K}_{xy,\delta}$. Applying the nonlinear transfer of volume Lemma III.2, we find the probability of $\mathbb{G}_\alpha(x_1, y_1) = 0$ for some $(x_1, y_1) \in \mathcal{K}_{xy,\delta}$ relative to the ball $\|c_\alpha\| \leq \epsilon^{1/2} < a_1$ to be upper bounded by

$$\frac{C_K}{\epsilon^{2d}} \times \frac{D_\alpha! 2^{D-1} L^{D-1} \epsilon^{\frac{D-1}{2}}}{\sigma_\delta^D}.$$

If $D \geq 4d+2$, the limit of this probability as $\epsilon \rightarrow 0$ is 0. Therefore, every c_α^* satisfying $\|c_\alpha^*\| < a_0$ is a Lebesgue point of $\mathcal{A}_{xy,\delta}$, which completes the proof of the lemma. \square

We are now prepared to state and prove the main theorem of this section.

Theorem III.19. *Assuming a_0 and ϕ_α satisfy the conditions laid down in section 6 and $D \geq 4d + 2$, the delay mapping F_α is injective on the set K with probability one relative to the ball $\|c_\alpha\| < a_0$.*

Proof. The proof follows from Lemmas III.15, III.17, and III.18 by taking the limit $\delta \rightarrow 0$ through a countable sequence. \square

3.8 Proof of immersivity

All the main techniques have been demonstrated in the proof of injectivity of the delay mapping F_α . The assumption in section 6 that dF_α is immersive at all fixed points in K simplifies the proof of immersivity considerably.

Define $\mathcal{A}_{T,\delta}$ as the set of all c_α satisfying $\|c_\alpha\| < a_0$ and F_α is immersive at all $x_1 \in K$ satisfying $\|x_1 - \xi_j(\alpha)\| \geq 3\delta$ for $j \in \{1, \dots, m\}$. In other words, we are requiring $dF_\alpha(x_1, v_1) \neq 0$ if $(x_1, v_1) \in T_1K$ and x_1 is removed from each periodic point by at least 3δ .

Lemma III.20. *For $D \geq 4d$, every point of $\|c_\alpha\| < a_0$ is a Lebesgue point of $\mathcal{A}_{T,\delta}$ and therefore the probability of $\mathcal{A}_{T,\delta}$ relative to $\|c_\alpha\| < a_0$ is 1.*

Proof. We center at c_α^* satisfying $\|c_\alpha^*\| < a_0$ as before. Again as before, we assume a_1 to be so small that $\|\Sigma_j(c_\alpha) - \Sigma_j\| < \delta$ for $\|c_\alpha\| \leq a_1$.

Define $\mathcal{K}_{T,\delta}$ to be the set of all $(x_1, v_1) \in T_1K$ satisfying $\|x_1 - \Sigma_j\| \geq 2\delta$ for $j \in \{1, \dots, m\}$. Then

$$(3.11) \quad dF_\alpha(x_1, v_1) = dF_0(x_1, v_1) + \mathcal{N}(c_\alpha) + \mathcal{O}(c_\alpha^2)$$

with

$$\mathcal{N} = \begin{pmatrix} 0 \\ \mathbb{H}(x_1, v_1) \end{pmatrix}.$$

By Lemma III.12, the rank of \mathcal{N} is $D - 1$.

Define $\sigma_\delta = \min \sigma_{D-1}(\mathcal{N})$, where the minimum is taken over $(x_1, v_1) \in \mathcal{K}_{T,\delta}$ and $\|c_\alpha\| \leq a_1$. Cover $\mathcal{K}_{T,\delta}$ with C_K/ϵ^{2d-1} ϵ -balls.

Assumption about L (5): In (3.11), the $\mathcal{O}(c_\alpha^2)$ term is upper bounded by $L\|c_\alpha\|^2$.

Here, we are effectively assuming a bound on the third derivative of $F_\alpha(x_1)$ with respect to x_1 over the compact sets $x_1 \in T_1K$ and $\|c_\alpha\| \leq a_0$.

If $d\mathbb{F}_\alpha(x_1, v_1) = 0$ for some $(x_1, v_1) \in \mathcal{K}_{T,\delta}$, then we must have $\|d\mathbb{F}_\alpha(x_1, v_1)\|$ for some (x_1, v_1) that is at the center of one of the ϵ -balls covering $\mathcal{K}_{T,\delta}$. The nonlinear transfer of volume lemma III.2 implies that the probability of $d\mathbb{F}_\alpha(x_1, v_1) = 0$ for some $(x_1, v_1) \in \mathcal{K}_{T,\delta}$ relative to $\|c_\alpha\| \leq \epsilon^{1/2} < a_1$ is upper bounded by

$$\frac{C_K}{\epsilon^{2d-1}} \times \frac{D_\alpha! 2^{D-1} L^{D-1} \epsilon^{\frac{D-1}{2}}}{\sigma_\delta^{D-1}}.$$

If $D \geq 4d$, this probability goes to zero as $\epsilon \rightarrow 0$. Therefore, every $\|c_\alpha^*\| < a_0$ is a Lebesgue point of $\mathcal{A}_{T,\delta}$, proving the lemma. \square

We are now prepared to state and prove the immersivity theorem.

Theorem III.21. *Suppose a_0 and ϕ_α satisfy the assumptions laid down in section 6 and suppose $D \geq 2d$. The delay map F_α is then immersive at every point of K with probability 1 relative to the ball $\|c_\alpha\| < a_0$.*

Proof. The proof follows by taking $\delta \rightarrow 0$ through a countable sequence in the previous Lemma III.20 and using the assumption made in section 6 about immersivity at fixed points. \square

3.9 Discussion

The delay map may be viewed in light of the Whitney embedding theorem [24]. However, it has some characteristics of its own. One of these is the possibility that

orbits of two distinct points can overlap. There are other distinctive characteristics related to periodic orbits and eigenvectors.

In this article, we showed how to prove that the delay map is an embedding using the concept of Lebesgue points. For the delay map $F_\alpha(x)$ with $o = \pi_1$ to be an embedding with probability 1 relative to the ball $\|c_\alpha\| < 1$, we require the embedding dimension to satisfy $D \geq 4d + 2$.

We conjecture that the delay mapping is an embedding for $D \geq 2d + 1$. The more restrictive $4d + 2$ requirement comes in when applying the nonlinear transfer of volume lemma. The extra dimensions are used to absorb the effect of the nonlinear term. Some evidence for this conjecture may be found in chapter 2 of this thesis.

In our opinion, it would be desirable to obtain prevalence versions of classical theorems such as the Kupka-Smale theorem [41]. The differential topology proofs rely heavily on the bump function and genericity is weaker than almost sureness in probability. It is hoped that the technique based on Lebesgue points introduced here will be useful in that regard.

CHAPTER IV

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

4.1 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 [13]. 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

$$(4.1) \quad (x(t), x(t - \tau), \dots, x(t - (D - 1)\tau))$$

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 [47].

Farmer and Sidorowich [13] 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 [33, 36, 37]. De-

tailed computations demonstrating the advantages of kernel based predictors were given by Mller et al [37] and are also discussed in the textbook of Schlkopf and Smola [48]. Kernel methods still appear to be the best, or among the best, for the prediction of stationary time series [34, 46].

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, [29, 31, 30] 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 [31] 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 [29], has been considered by Steinwart and Anghel [50] (also see [7]). 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 [29]. 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 [47] 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, 20, 51]. 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 [37] 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 [59]. 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 [34, 35] and optimal consistency estimates for stationary data [23].

4.2 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 [47], 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 [10, 11], 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 IV.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 Mller et al [37] 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 Mller et al computes f_m such that the functional

$$(4.2) \quad \frac{1}{Nn+1} \sum_{j=0}^{Nn} |f(u_\eta(jh; \tau; \omega)) - u_\eta(jh + \tau; \omega)|_\epsilon + \Lambda \|f\|_{K_\gamma}^2$$

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 Mller 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 [8] to compute a function $u_s(t; \omega)$, $t \in [-(D-1)\tau, N\tau + t_f]$ such that the functional

$$(4.3) \quad \frac{1}{(N + n_f + D - 1)n + 1} \sum_{j=-(D-1)n}^{(N+n_f)n} (u_\eta(jh; \omega) - \tilde{u}(jh))^2 + \lambda \int_{-(D-1)\tau}^{N\tau+t_f} \tilde{u}''(t)^2 dt$$

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 Mller et al. The predictor f_1 is computed as

$$(4.4) \quad f_1 = \operatorname{argmin}_{f \in \mathcal{H}_k} \frac{1}{Nn+1} \sum_{j=0}^{Nn} (f(u_s(jh; \tau; \omega)) - u_s(jh + t_f; \omega))^2 + \Lambda \|f\|_{K_\gamma}^2.$$

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 (4.4) differs from the algorithm of Mller 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 [7], 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 [10, 11] (see pages 132 and 133 of [11]). 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 IV.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 (4.3) and (4.4) follow. First, the lemma assumes a fixed choice of λ (the relevant theorem in [10, 11] in fact allows λ to lie in an interval). In our algorithm, λ is determined using cross-validation because of its practical effectiveness [60].

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 [52] 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 Gyrfi et al [21]. 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 Gyrfi 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 [49], 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 (4.4) 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 IV.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 IV.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 Stewart, Hush, and Scovel [51]. However, the algorithm of (4.3) and (4.4) does not fit into the framework of [51]. 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 [51]. The analysis summarized above and given in detail in the following section relies on ∞ -norm bounds.

4.3 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 [47]. 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 IV.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 [47] 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 [12]. It is stated and proved for completeness.

Lemma IV.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 IV.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

$$(4.5) \quad B = \sup_{\omega \in X} \|\omega\|_\infty + M$$

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 (4.4), 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 IV.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 IV.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 [61]), 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

$$(4.6) \quad B'_1 = \max(B, c_0 B, c_1 B),$$

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

$$(4.7) \quad |\mathcal{W}_1(f) - \mathcal{W}_2(f)| \leq \frac{1}{Nn+1} \sum_{j=0}^{Nn} \frac{4B'_1}{\Lambda^{1/2}} \left(\begin{array}{l} |f(u_s(jh; \tau; \omega)) - f(u(jh; \tau; \omega))| \\ + |u_s(jh; \tau; \omega) - u(jh; \tau; \omega)| \end{array} \right).$$

Here $\frac{4B'_1}{\Lambda^{1/2}}$ is used as 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 (4.7) and defining B_1 as $B_1 = \sqrt{8}B'_1$. \square

Lemma IV.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 IV.5, IV.6, and IV.4. Lemma IV.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 IV.5 and the choice of δ is justified by Lemma IV.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 IV.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 IV.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 IV.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 IV.9. *Assume $0 \leq \Lambda < 1$. Then $\|f_2 - f_3\| \leq \frac{B_1 h^{1/2}}{\Lambda}$.*

Proof. Follows from Lemmas IV.5, IV.8, and IV.4. Lemma IV.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 IV.5 and IV.8 and an additional argument as in the proof of Lemma IV.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 IV.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 $\|f_3\|_K^2 \leq \mathcal{W}_3(F_\epsilon)/\Lambda$. \square

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 IV.10 tells us that $\|f_3\|_K \leq \sqrt{2} \|F_\epsilon\|_K$, and therefore (by part (c) of Theorem 1 of [61]) $\|\partial f_3\|_\infty \leq \sqrt{2} D^{1/2} \|\partial^2 K\|_{2,\infty}^{1/2} \|F_\epsilon\|_K$. As a consequence of our choice of ℓ , $x, y \in A_j$ implies that

$$(4.8) \quad |f_3(x) - f_3(y)| < \delta/4,$$

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

$$(4.9) \quad |F(x) - F(y)| < \delta/4$$

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 IV.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 $T^* = \max(s_1, \dots, s_L)$. \square

Lemma IV.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*

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

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, $S_\delta \subset \bigcup_{j \in J} A_j$, and it is sufficient to bound the measure of $\bigcup_{j \in J} A_j$.

By (4.8) and (4.9), if $|f_3(x) - F(x)| \geq \delta$ for some $x \in A_j$ then for any $y \in A_j$, we have

$$\begin{aligned} |f_3(y) - F(y)| &\geq |f_3(x) - F(x)| - |f_3(x) - f_3(y)| - |F(x) - F(y)| \\ (4.10) \qquad &> \frac{\delta}{2}. \end{aligned}$$

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 (4.10) for $y = u(t; \tau; \omega) \in A_j$ with $j \in J$, and the final inequality is a consequence of Lemma IV.11 and $\omega \in S_{\epsilon, T^*}$.

Applying Lemma IV.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 IV.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 IV.7, IV.9, and IV.12. \square

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

4.4 Numerical illustrations

We compare three methods to compute an approximate predictor f . The first method is that of Mller et al [37] given in (4.2). The second method is exactly the same but with the least squares regression function. The third method is the convergent algorithm given by (4.3) and (4.4).

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 [37], 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 [37], 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 [37], the Mackey-Glass signal was down-sampled so that $nh = 1$

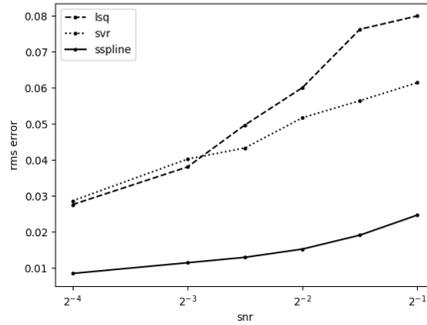


Figure 4.1: Root mean square errors in the prediction of the Mackey-Glass signal with $t_f = 1$ as a function of the signal to noise ratio. The superiority of the method using smooth splines is evident.

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 [37]. 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 [37].¹

Figure 4.1 demonstrates that (4.2) 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 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

¹The RMS error of 0.017 reported for $t_f = 1$ with SNR of 22.15% in [37] 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.

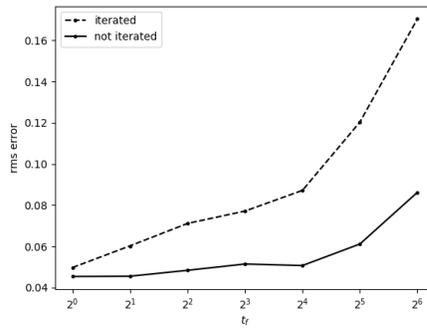


Figure 4.2: Comparison of the 1-step least squares predictor (without spline smoothing) iterated t_f times with the t_f -step predictor (without spline smoothing). The latter is seen to be superior.

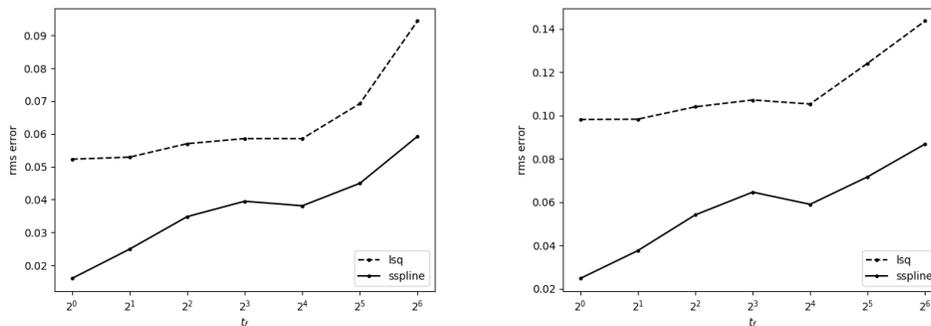


Figure 4.3: The plot on the left uses SNR of 0.2 and the plot on the right uses 0.4. The method using smooth splines does better in all instances.

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 [37]. This is not a good idea as explained in [59] 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.

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 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 $h = 0.01$.

4.5 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

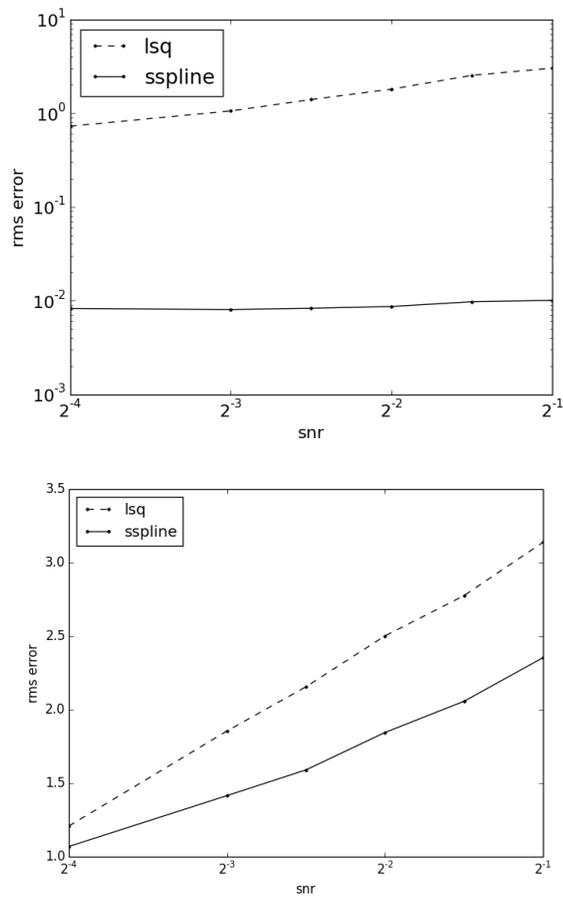


Figure 4.4: The advantage of spline smoothing for Lorenz is much less on the right with $h = 0.1$ than on the left with $h = 0.01$.

applying kernel-based regression to the denoised signal. The resulting predictor converges to the exact predictor under conditions described by Theorem IV.2.

As far as dynamical time series are concerned, the parameter estimation problem [34, 35] 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 [23, 50]. Rapid mixing results for flows may be found in [4] 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 [59], 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 chapter 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 [52]. 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 [57], 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.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] T.M. Adams and A.B. Nobel. On density estimation from ergodic processes. *The Annals of Probability*, 25:794–804, 1998.
- [2] D. Aeyels. Generic observability of differentiable systems. *SIAM Journal on Control and Optimization*, 19(5):595–603, 1981.
- [3] K.T. Alligood, T.D. Sauer, and J.A. Yorke. *Chaos: An Introduction to Dynamical Systems*. Springer, 2000.
- [4] 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.
- [5] C. Börgers. *An Introduction to Modeling Neuronal Dynamics*. Springer, 2017.
- [6] M. Casdagli. Nonlinear prediction of chaotic time series. *Physica D*, 35:335–356, 1989.
- [7] A. Christmann and I. Steinwart. Consistency and robustness of kernel-based regression in convex risk minimization. *Bernoulli*, 3:799–819, 2007.
- [8] C. de Boor. *A Practical Guide to Splines*. Springer, New York, revised edition, 2001.
- [9] M. Dellnitz, M. Hessel-Von Molo, and A. Ziessler. On the computation of attractors for delay differential equations. *Journal of Computational Dynamics*, 3:93–112, 2016.
- [10] 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.
- [11] P.P.B. Eggermont and V.N. LaRiccia. *Maximum Penalized Likelihood Estimation*, volume II. Springer, New York, 2009. Springer Series in Statistics.
- [12] I. Ekeland and R. Temam. *Convex Analysis and Variational Problems*. SIAM, Philadelphia, 1987.
- [13] J. Doyne Farmer and J.J. Sidorowich. Predicting chaotic time series. *Physical Review Letters*, 59:845–848, 1987.
- [14] R.L. Foote. Regularity of the distance function. *Proceedings of the American Mathematical Society*, 92:153–155, 1984.
- [15] D. Forger. *Biological Clocks, Rhythms, and Oscillations: The Theory of Biological Timekeeping*. MIT Press, 2017.
- [16] J. F. Gibson, J. D. Farmer, M. Casdagli, and S. Eubank. An analytic approach to practical state space reconstruction. *Physica D: Nonlinear Phenomena*, 57(1):1–30, 1992.
- [17] V. Guillemin and A. Pollack. *Differential topology*, volume 370. American Mathematical Society, 2010.

- [18] Y. Gutman. Takens embedding theorem with a continuous observable. In *In Ergodic Theory: Advances in Dynamical Systems*, pages 134–142. Walter de Gruyter GmbH and Co KG, 2016.
- [19] Y. Gutman, Y. Qiao, and G. Szabo. The embedding problem in topological dynamics and Takens’ theorem. *arxiv.org*, 2017.
- [20] 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.
- [21] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer, New York, 2002.
- [22] F. Hamilton, T. Berry, and T. Sauer. Kalman-takens filtering in the presence of dynamical noise. *European Journal of Physics*, to appear.
- [23] 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.
- [24] M. W. Hirsch. *Differential Topology*. Springer, 1976.
- [25] M. W. Hirsch. *Differential topology*, volume 33. Springer Science & Business Media, 2012.
- [26] B. R. Hunt, T. Sauer, and J. A. Yorke. Prevalence: a translation-invariant ”almost every” on infinite-dimensional spaces. *Bulletin of the American mathematical society*, 27(2):217–238, 1992.
- [27] S.G. Krantz and H.R. Parks. Distance to c^k hypersurfaces. *Journal of Differential Equations*, 40:116–120, 1981.
- [28] I. Kukavica and J. C. Robinson. Distinguishing smooth functions by a finite number of points values, and a version of the takens embedding theorem. *Physica D: Nonlinear Phenomena*, 196(1):45–66, 2004.
- [29] S. P. Lalley. Beneath the noise, chaos. *The Annals of Statistics*, 27:461–479, 1999.
- [30] S. P. Lalley and A. B. Nobel. Denoising deterministic time series. *Dynamics of PDE*, 3:259–279, 2006.
- [31] 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.
- [32] W.R. Madych. An estimate for multivariate interpolation II. *Journal of Approximation Theory*, 142:116–128, 2006.
- [33] 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.
- [34] K. McGoff, S. Mukherjee, and N. Pillai. Statistical inference for dynamical systems: A review. *Statistics Surveys*, 9:209–252, 2015.
- [35] K. McGoff and A.B. Nobel. Empirical risk minimization and complexity of dynamical models. *www.arxiv.org*, 2016.
- [36] 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.
- [37] 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.

- [38] A. Nobel. Consistent estimation of a dynamical map. In *Nonlinear dynamics and statistics*, pages 267–280. Springer, 2001.
- [39] N.H. Packard, J.P. Crutchfield, J.D. Farmer, and R.S. Shaw. Geometry from a time series. *Physical Review Letters*, 45:712–716, 1980.
- [40] J. Jr. Palis and W. De Melo. *Geometric theory of dynamical systems: an introduction*. Springer Science & Business Media, 2012.
- [41] C. Robinson. *Dynamical Systems: Stability, Symbolic Dynamics, and Chaos*. CRC Press, 1998.
- [42] J. C. Robinson. A topological delay embedding theorem for infinite-dimensional dynamical systems. *Nonlinearity*, 18(5):2135–2143, 2005.
- [43] J.C. Robinson. *Dimensions, Embeddings, and Attractors*. Cambridge, 2011.
- [44] M. Robinson. A topological low pass filter for quasiperiodic signals. *IEEE Sig. Proc. Lett.*, 23:1771–1775, 2016.
- [45] W. Rudin. *Real and Complex Analysis*. McGraw-Hill, 1986.
- [46] N. I. Sapankevych and R. Sankar. Time series prediction using support vector machines: a survey. *IEEE Computational Intelligence Magazine*, pages 24–38, May 2009.
- [47] T. Sauer, J. A. Yorke, and M. Casdagli. Embedology. *Journal of Statistical Physics*, 65(3):579–616, 1991.
- [48] B. Schölkopf and A.J. Smola. *Learning with Kernels*. MIT Press, MA, 2002.
- [49] I. Steinwart. On the influence of the kernel on the consistency of support vector machines. *Journal of Machine Learning Research*, 2:67–93, 2001.
- [50] 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.
- [51] I. Steinwart, D. Hush, and C. Scovel. Learning from dependent observations. *Journal of Multivariate Analysis*, 100:175–194, 2009.
- [52] C. Stone. Optimal rates of convergence for nonparametric regression. *Annals of Statistics*, 10:1040–1063, 1982.
- [53] S. Strogatz. *Nonlinear Dynamics: with Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, 2014.
- [54] F. Takens. Detecting strange attractors in turbulence. *Lecture Notes in Mathematics*, 898(1):366–381, 1981.
- [55] F. Takens. The reconstruction theorem for endomorphisms. *Bulletin of the Brazilian Mathematical Society*, 33(2):231–262, 2002.
- [56] L.N. Trefethen and D. Bau III. *Numerical Linear Algebra*. SIAM, 1997.
- [57] K. Urbanowicz and J.A. Holyst. Noise-level estimation of time series using coarse-grained entropy. *Physical Review E*, 67:046218, 2003.
- [58] D. Viswanath. Symbolic dynamics and periodic orbits of the Lorenz attractor. *Nonlinearity*, 16:1035–1056, 2003.
- [59] 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.

- [60] G. Wahba. *Spline Models for Observational Data*. SIAM, Philadelphia, 1990.
- [61] D.-X. Zhou. Derivative reproducing properties for kernel methods in learning theory. *Journal of Computational and Applied Mathematics*, 220:456–463, 2008.