

---

# Dimensionality Collapse: Optimal Measurement Selection for Low-Error Infinite-Horizon Forecasting

---

**Helmuth Naumer**

University of Illinois Urbana-Champaign

**Farzad Kamalabadi**

University of Illinois Urbana-Champaign

## Abstract

This work introduces a method to select linear functional measurements of a vector-valued time series optimized for forecasting distant time-horizons. By formulating and solving the problem of sequential linear measurement design as an infinite-horizon problem with the time-averaged trace of the Cramér–Rao lower bound (CRLB) for forecasting as the cost, the most informative data can be collected irrespective of the eventual forecasting algorithm. By introducing theoretical results regarding measurements under additive noise from natural exponential families, we construct an equivalent problem from which a local dimensionality reduction can be derived. This alternative formulation is based on the future collapse of dimensionality inherent in the limiting behavior of many differential equations and can be directly observed in the low-rank structure of the CRLB for forecasting. Implementations of both an approximate dynamic programming formulation and the proposed alternative are illustrated using an extended Kalman filter for state estimation, with results on simulated systems with limit cycles and chaotic behavior demonstrating a linear improvement in the CRLB as a function of the number of collapsing dimensions of the system.

## 1 INTRODUCTION

There is a long history of optimal experimental design in statistics. In such work, rather than directly design estimators, sequences of experiments are designed to optimize some information criteria. This criteria in turn implies the existence of accurate estimators. These optimization objectives often take the form of functions of the inverse Fisher

information matrix such as the alphabet optimality criteria including minimizing the trace, determinant, or maximum eigenvalue of the matrix (Chaloner and Verdinelli, 1995; Dette, 1997). In the past two decades, a significant motivation of optimal experimental design has been found in optimal sensor placement (Martínez and Bullo, 2006; Ranieri et al., 2014; Xygkis et al., 2018). In control theory, sequential experimental design appears in the design of informative inputs to a parameterized system for the purpose of identifying the unknown parameters (Fohring and Haber, 2016; Asprey and Macchietto, 2002; Huan and Marzouk, 2016; Titterington, 1980). Despite this history, it has recently been observed that optimizing for correct system parameters often produces poor forecasts of the future system state, and that predictive performance should be considered directly (Casey et al., 2007; Transtrum and Qiu, 2012; Letham et al., 2016). These works typically assume the ability to control the system itself, observing multiple full trajectories and driving the system into informative regions of the state space. Contrary to this, in many time-series forecasting problems, we only observe a single trajectory of which we have no control.

Time series forecasting is a long-studied problem often motivated by numerous applications such as weather prediction (Penland and Magorian, 1993; Ehrendorfer, 1994, 2006), finance (Abu-Mostafa and Atiya, 1996; Cao and Tay, 2001), and control (Bertsekas, 2017b). Oftentimes, algorithms focus only on short-term predictions, intrinsically coupled with state estimation as seen in classical techniques such as the Kalman filter and particle filter. In the extreme short-term, there have been significant developments in recent years applying machine learning tools such as generative adversarial networks to form the predictions (Shi et al., 2017; Ravuri et al., 2021).

On the other extreme timescale, there has been growing interest in reconstructing information about the limiting behavior of the time series from data. The approaches vary dramatically from the application of embeddings theorems to reconstruct the shape of attractors from lower-dimensional measurements (Gilpin, 2020), to regression techniques that learn sparse governing equations (Brunton et al., 2016; Kaiser et al., 2018; Zhang and Schaefer

fer, 2019; Messenger and Bortz, 2021), to Koopman operator and Dynamic Mode Decomposition (DMD) based techniques which form inherently linear approximations in higher dimensional spaces (Schmid, 2010; Williams et al., 2015; Proctor et al., 2016; Erichson et al., 2019). Furthermore, the scientific computing community has been recently experimenting with using neural networks to approximate the solutions of differential equations through physics-informed neural networks (Raissi et al., 2019) as a more computationally efficient alternative to Gaussian process regression (Raissi and Karniadakis, 2018).

Even when the underlying data is not governed by differential equations, many modern techniques seek to learn a differential equation which well-represents the phenomena of interest as a new general form of regression — enabling the application of dynamical systems tools to broader learning problems. While some of these techniques are focused on using the expressive nature of differential equations in lieu of neural networks (Chen et al., 2018; Dupont et al., 2019; Dandekar et al., 2021; Massaroli et al., 2020; Avelin and Nyström, 2021), others instead use the extensive understanding of differential equations to constrain the geometry of the estimates based on integrals of motion (Greydanus et al., 2019; Cranmer et al., 2020). Furthermore, there have been attempts to construct implicitly deep networks by fitting the fixed points of an ODE in deep equilibrium models (Bai et al., 2019, 2020; Pabbaraju et al., 2021; Gilton et al., 2021).

A fortunate consequence of this trend in modeling is that the generative assumptions involving differential equations often implicitly imbue the future with significant structure. Even Lyapunov stability, one of the weakest forms, immediately implies that the feasible set of future points monotonically shrinks over time (Khalil, 2002). In many cases though, this additional structure takes the form of the guarantee that the long-term behavior is actually confined to a lower dimensional space than the ambient dimensionality. In this work, we propose methods based on this dimensionality collapse for informed data collection in forecasting time series modeled by differential equations.

Our key observation is that, by incorporating the future structure of the dynamical system into the present, many linear measurements become entirely uninformative — no meaningful prediction regarding distant time-horizons can be made from the observation. By avoiding these uninformative measurements in favor of observing the informative subspace, we introduce a method for reducing the Cramér–Rao lower bound (CRLB) for forecasting. This reduction scales linearly with the dimensionality of the uninformative space, reducing the errors in estimation resulting from embedding a low-dimensional manifold of feasible data in a high-dimensional ambient observation space.

## Contributions

- We analyze the problem of sequential linear measurement design in systems governed by differential equations with low-dimensional limiting behavior, providing both theoretical results and practical algorithms.
- We prove that the time-averaged infinite horizon optimal sampling problem is equivalent to maximizing the sum of  $K$  squared inner products subject to a norm constraint, where  $K$  is the dimensionality of the limiting behavior, thereby reducing the dimensionality of the problem.
- We specialize our results to construct a computationally efficient solution to the problem when the system converges to an isolated limit cycle.
- We demonstrate the algorithm on simulated chaotic systems and limit cycles systems: illustrating the broad applicability of the model assumptions and the linear improvement of the CRLB as a function of the number of collapsing dimensions.

## 2 PROBLEM FORMULATION

In this work, we construct a method to optimize a policy for linear functional sampling which minimizes the CRLB for forecasting future values of the time series. There are many closely related variations of this problem, all built upon the same theme. We begin this section by defining the model assumptions, before explicitly defining the variations of the problem solved in the manuscript. The section concludes with a brief note on the dynamic programming approach to the problems, which represent a clear baseline.

**System Model** This work analyzes the state space model of the form

$$\dot{x} = f(x), \quad (1)$$

$$y_i = \langle u_i, x_{\tau_i} + \xi_i \rangle \quad (2)$$

where  $x \in \mathbb{R}^M$  is the state with temporal derivative  $\dot{x} = \frac{dx}{dt}$ ,  $f : \mathbb{R}^M \rightarrow \mathbb{R}^M$  is Lipschitz on some domain containing the system trajectory with a Lipschitz continuous derivative,  $y_i$  is the  $i$ 'th measurement,  $x_{\tau_i}$  is the state at time  $\tau_i$ ,  $\langle u_i, \cdot \rangle$  represents the  $i$ 'th chosen measurement functional based on the Euclidean inner product, and  $\xi_i$  is additive noise during observation  $i$ , independent between timesteps with further assumptions to be described in detail in the following paragraph. In general, we can combine unknown constant parameters of the dynamical system with the state. The Lipschitz dynamics constraint is common when working with differential equations because it enables the existence and uniqueness theorem, guaranteeing a unique solution to a differential equation for a given initial state (Khalil, 2002). The uniqueness of trajectories immediately implies the existence of a set of time-advancing

| SYSTEM           | PARAMETER CONDITIONS      | AMBIENT SPACE     | LIMITING SPACE                     |
|------------------|---------------------------|-------------------|------------------------------------|
| GRADIENT FLOW    | MORSE FUNCTION            | $\mathbb{R}^M$    | COUNTABLE SET                      |
| VAN DER POL      | KNOWN $\mu > 0$           | $\mathbb{R}^2$    | $\mathbb{S}^1$                     |
| HOPF BIFURCATION | UNKNOWN $\alpha, \lambda$ | $\mathbb{R}^4$    | $\mathbb{R}_+ \times \mathbb{S}^1$ |
| LORENZ           | KNOWN $\sigma, \rho$      | $\mathbb{R}^3$    | FRACTAL DIMENSION 2.06             |
| HEAT EQUATION    | KNOWN $\alpha > 0$        | $W^{1,2}(\Omega)$ | $\mathbb{R}$                       |

Table 1: Example dynamical systems with larger ambient space than limiting space.

functions  $\varphi^\tau : \mathbb{R}^M \rightarrow \mathbb{R}^M$  which satisfies the semigroup property  $\varphi^{\tau_1} \circ \varphi^{\tau_2} = \varphi^{\tau_1 + \tau_2}$ , where  $\circ$  represents function composition. Such a group of functions is known as the flow of the dynamical system. We denote the state at time  $\tau$  as  $x_\tau$ , and the flow is defined such that  $x_\tau = \varphi^\tau(x_0)$  for any given initial state  $x_0$ . If additionally  $f$  is differentiable in  $x$ , then  $\varphi^\tau$  can be directly shown to be continuously differentiable in  $x$  for all  $\tau$  through the construction of the sensitivity equation (Khalil, 2002). Differentiability is only a very mild additional assumption for Lipschitz dynamical systems, as Lipschitz functions from  $\mathbb{R}^N$  onto  $\mathbb{R}^N$  can always be uniformly approximated by smooth Lipschitz functions [see, for instance, (Hájek and Johaniš, 2010) for a more general case].

**Measurement Family** The family of measurements is defined through two steps. First, we assume that the state of the system represents some physically meaningful quantities for which we have some method of controlled linear functional observation under additive noise, as shown in Equation (2). Linear functional observations capture many common types of measurements. If our state space is extended, for instance, to almost everywhere bounded continuous functions, i.e.  $(C(\Omega), \|\cdot\|_\infty)$ , then line integrals representing tomography and Dirac delta functions representing discrete sensor placement both exist in our measurement family. Furthermore, in finite dimensional spaces, linear combinations of measurements represent common applications such as beamforming or discretized versions of the previous continuous examples.

We further assume that each  $\xi_i$  is in a natural exponential family (NEF). A random variable  $\xi$  is said to be in a NEF if the probability density function can be written as

$$p(\xi; \theta) = \exp\{\theta^\top \xi - A(\theta) + B(\xi)\} \quad (3)$$

where  $A$  is the log-partition function,  $\theta$  is an unknown parameter vector, and  $B$  is the log-base function. Many common distributions can be reparameterized into this form, such as Gaussian, Gamma, or Poisson distributions (Morris, 2006). *Note that the elements of the noise vector need not come from the same probability distribution, nor even the same family of distributions.* For instance, one element could represent Gaussian noise, while another could be Poisson noise.

**Forecasting Bound Objective** The CRLB is a well-known matrix inequality for the covariance of unbiased estimators. While a full description is provided in Appendix B, we briefly describe the key properties in this section. Through the well-known decomposition of mean squared error (MSE) into the sum of squared bias and variance,  $\text{MSE} = \mathbb{E}[\|\hat{x} - \mathbb{E}[\hat{x}]\|_2^2]$  for all  $\hat{x}$  such that  $\mathbb{E}[\hat{x} - x] = 0$ . By expressing the squared norm as  $\|\hat{x} - \mathbb{E}[\hat{x}]\|_2^2 = (\hat{x} - \mathbb{E}[\hat{x}])^\top (\hat{x} - \mathbb{E}[\hat{x}])$  and applying the cyclic permutation invariance and linearity of the trace, we observe that the MSE is lower bounded as

$$\mathbb{E}[\|\hat{x} - x\|_2^2] = \text{Tr}(\mathbb{E}[(\hat{x} - \mathbb{E}[\hat{x}])(\hat{x} - \mathbb{E}[\hat{x}])^\top]) \quad (4)$$

$$= \text{Tr}(\Sigma_{\hat{x}}) \geq \text{Tr}(\Sigma), \quad (5)$$

where  $\Sigma_{\hat{x}}$  is the covariance of  $\hat{x}$  and  $\Sigma$  is the CRLB, and  $\text{Tr}$  denotes the trace of a matrix. By considering the future state to be the unknown parameter, the same bound may be applied to forecasting. In forecasting deterministic Lipschitz dynamical systems based on measurements in exponential families of random variables, the CRLB undergoes the linear transformation

$$\Sigma_\tau = [d\varphi_{x_0}^\tau] \Sigma_0 [d\varphi_{x_0}^\tau]^\top, \quad (6)$$

where  $\Sigma_\tau$  denotes the CRLB for predicting the state at time  $\tau$ ,  $\Sigma_0$  denotes the CRLB for estimating the current state, and  $d\varphi_{x_0}^\tau$  represents the differential of  $\varphi^\tau$  at  $x_0$ . Brackets have been included to emphasize that  $d\varphi_{x_0}^\tau$  is a matrix and represents the Jacobian of the flow. This transformation comes directly from the observation that advancing the state forward in time is a diffeomorphic reparameterization of the probability distribution (Lehmann and Casella, 1998).

We assert that in many stable systems,  $d\varphi_{x_0}^\tau$  becomes approximately low-rank as  $\tau$  grows large, and thus we can consider optimal data collection to be done by making measurements which minimize  $\Sigma_0$  along the non-zero modes of the Jacobians. This low-rank structure occurs due to systems converging to lower-dimensional limiting spaces. In Table 1, we provide a number of examples of dynamical systems with their ambient space and limiting space. In the most extreme case, the heat equation reduces from a Sobolev space  $W^{1,2}(\Omega)$ , which is countably infinite, to the real line which is one dimensional. Full system definitions are provided in Appendix A.

**Measurement Optimization** Representing the expected squared magnitude of the error of potential estimators, we consider our cost to be the trace of the CRLB. We consider both the time averaged cost which is applicable to problems such as limit cycles, as well as a discounted cost formulation more strongly applicable to chaotic dynamics due to the exponential growth of the CRLB based on the largest Lyapunov exponent (Bertsekas, 2017a; Strogatz, 2015).

To simplify the computational aspects while retaining the meaningful structure, we approximate our cost through sampling. Then, we have the following two infinite-horizon dynamic programming problems, where  $\gamma$  represents the discount factor and  $\mathcal{S}$  represents the set of indices for the measurement times.

**Problem 1** (Discounted Cost Optimal Design).

$$\arg \min_{\{u_i\}_{i \in \mathcal{S}}} \lim_{N \rightarrow \infty} \sum_{j=1}^N \gamma^j \text{Tr}(\Sigma_{\tau_j}(\{u_i\}_{i \in \mathcal{S}})) \quad (7)$$

**Problem 2** (Average Cost Optimal Design).

$$\arg \min_{\{u_i\}_{i \in \mathcal{S}}} \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N \text{Tr}(\Sigma_{\tau_j}(\{u_i\}_{i \in \mathcal{S}})) \quad (8)$$

Both discounted and average cost infinite-horizon problems are well-studied in stochastic control and in reinforcement learning, where numerous variants of value and policy iteration are applied (Bertsekas, 2017a).

An important note is that in unstable and chaotic systems, we still observe an approximate low dimensionality. Consider a system for which the two largest eigenvalues of  $d\varphi^\tau$  are  $\lambda_1, \lambda_2 > 1$ . Without additional sampling, the ratio of the two largest eigenvalues of the geometric sum  $\sum_j \gamma^j d\varphi^{\tau_j}$  in the discounted cost formulation will become

$$\frac{1 - \lambda_2 \gamma}{1 - \lambda_1 \gamma}. \quad (9)$$

There always exists a discount factor  $\gamma < 1$  for which this gap becomes arbitrarily large. Hence, there will always be some discounted cost formulation for which the cost matrix corresponding to the CRLB is approximately singular.

**Dynamic Programming Formulation** The standard formulation to infinite-horizon discounted and average cost problems is through the Bellman equation, enabling common algorithms such as value iteration and policy iteration. For a known state, the discounted cost Bellman equation becomes

$$\begin{aligned} x_+ &= \varphi^\tau(x) \\ \Sigma_+(u) &= [d\varphi_x^\tau] (\Sigma^{-1} + \|u\|_{\Lambda_\sigma}^{-2} u u^\top)^{-1} [d\varphi_x^\tau]^\top \quad (10) \\ J(x, \Sigma) &= \text{Tr} \Sigma + \gamma \max_{u \in \mathcal{A}} \{J(x_+, \Sigma_+(u))\}, \end{aligned}$$

where  $x_+$  and  $\Sigma_+$  represent the updated state and CRLB respectively,  $\|u\|_{\Lambda_\sigma}^2 = u^\top \Lambda_\sigma u$  represents the squared norm induced by the measurement noise covariance matrix  $\Lambda_\sigma$ ,  $J$  represents the value function, and  $\mathcal{A}$  denotes the set of possible normalized measurements functionals (Bertsekas, 2017a). The update to the CRLB comes from the total Fisher information of two independent observations being the sum of the Fisher information matrices. The CRLB update in Equation (10) comes directly from the form of the Fisher Information in Theorem 3.5.

To apply most common iterative algorithms for solving the Bellman equation, the state space must be discretized. Without strengthening the assumptions beyond the typical, already asserted, dynamical systems assumptions, the number of required samples can be shown to grow exponentially with the dimensionality of the system. As a baseline in this work, we include a local-averaging approximation, the details of which including sample complexity are available in Appendix C of the supplemental information.

### 3 DIMENSIONALITY REDUCTION TECHNIQUE

The computational cost of the dynamic programming approach with sampled state space approximation becomes prohibitively expensive as the ambient dimensionality of the system grows (See Appendix C). However, the future of a dynamical system will often be low dimensional, and in many cases approximately one dimensional. In this section, we analyze the problem when the governing dynamical system converges to a low dimensional limit set, enabling a computationally efficient alternative to the dynamic programming formulation for the average cost formulation. We then briefly discuss the extension to arbitrary low-dimensional embedded limit sets. While proofs for all results appear in Appendix D, we include numerous lemmas to aid in the exposition.

The core ideas of this section can be seen in Figure 1 applied to the Van Der Pol oscillator, a common example of a system exhibiting a limit cycle. In all panels, the plot axes represent the state of the 2D system. In panels A and B, we see the vector field and some example trajectories of system. Panel D illustrates  $\varphi^\infty := \lim_{\tau \rightarrow \infty} \varphi^\tau$  computed numerically, where the color indicates a representation of the phase along the limit cycle. As the mapping  $\varphi^\infty$  illustrated in panel D maps onto a one-dimensional space in a smooth manner outside of the origin, the Jacobian of the limiting flow in the ambient space is at most rank 1. The row space of  $d\varphi_x^\infty$  is plotted in panel C, representing the subspace that does not vanish as time goes to infinity. Our proposed method seeks to minimize the CRLB with respect to this subspace.

In Section 3.1, we introduce a series of statements char-




Figure 1: A: Vector field of Van Der Pol oscillator; B: Four example trajectories; C: Row space of limiting Jacobian; D: A smooth map  $\varphi^\infty : (\mathbb{R}^2 - \{0\}) \rightarrow \mathbb{S}^1$  defined by the dynamics.

acterizing the Fisher information matrix of the state vector for a single measurement of the form described in Equation (2). We show that the contribution of each measurement is, under a mild condition, the outer product of the measurement vector  $u$ , normalized by a norm defined by the noise distribution covariance matrix. Then, in Section 3.2 we show that this matrix structure enables a simple reformulation of the infinite horizon objective in order to simplify the problem of sequential optimal experimental design for forecasting.

### 3.1 Linear Measurements And Fisher Information

In this section, we show that independent NEFs preserve essential structures under a weighted average, resulting in a simple expression for the Fisher information of a linear functional measurement.

First, we must observe that the weighted average of independent random variables in NEFs is itself in another NEF.

**Lemma 3.1** (Proof: See Lemma D.1). *Given a finite set of independent random variables  $\{\xi_{(i)}\}$  in NEFs with log-partition functions  $\{A_i\}$ , parameters  $\{\theta_i\}$ , and a set of real-valued weights  $\{\alpha_i\}$ , then  $\bar{\xi} = \sum_i \alpha_i \xi_{(i)}$  is in a NEF with a log partition function*

$$A(\theta) := \sum_i A_i(\theta + \alpha_i \theta_i). \quad (11)$$

While we still lack a significant amount of knowledge regarding the specific distribution, fundamental properties of the log-partition function allow us to directly compute the




Figure 2: Commutative diagram of key maps and variables.

Fisher information of the observation about the parameter in the new family of probability distribution.

**Lemma 3.2** (Proof: See Lemma D.2). *The Fisher information of a random variable in a NEF is the variance of the random variable.*

The remaining task is to relate the Fisher information about the new parameter to the Fisher information regarding the mean values of the original distributions. The connections between the key parameters of interest are illustrated in the commutative diagram shown in Figure 2. In the illustration,  $\mu$  represents the mean of the noise vector  $\xi$ ,  $\bar{\mu}$  represents the mean of  $\bar{\xi}$ ,  $\theta$  represents the parameter vector of  $\xi$ , and  $\bar{\theta}$  represents the parameter of  $\bar{\xi}$  in the new model family. We can immediately observe  $h : \mu \mapsto \langle u, \mu \rangle$  by construction, and  $g : \theta \mapsto \frac{\partial}{\partial \Theta} A(\Theta)|_{\Theta=\theta}$  with the inverse denoted  $g^{-1}$ . Next,  $g_i$  and  $g$  are defined identically, but for  $\xi_{(i)}$  rather than  $\xi$ . The diagram commutes, and so we denote  $T$  to be the implicitly defined mapping from  $\mu$  to  $\bar{\theta}$ .

**Lemma 3.3** (Proof: See Lemma D.3). *The diagram in Figure 2 commutes for any vector  $\xi$  of independent random variables from NEFs.*

Before using the above quantities, we need to establish when  $g^{-1}$  exists. While full proofs are available in the appendix, invertibility comes from proving  $g$  is strictly monotonic.

**Lemma 3.4** (Proof: See Lemma D.4). *Assume the log-partition function  $A$  of a NEF is twice differentiable and that the second moment of the associated random variable exists. Then the inverse of  $g(\theta) := \frac{\partial}{\partial \Theta} A(\Theta)|_{\Theta=\theta}$  exists if and only if the random variable is non-degenerate for almost all  $\theta$ .*

Finally, observe that  $T$  is an injective function, and is thus invertible on the image. Furthermore, while  $g_i^{-1}$  is model specific, the Jacobian is directly related to the Fisher information. By recalling that Jacobian of  $g_i$  is the Fisher information of the probability distribution in the family, we can conclude the following theorem as the result of a reparameterization by  $T^{-1}$ .

**Theorem 3.5** (Proof: See Theorem D.5). *Suppose  $\xi$  is a vector of independent random variables from NEFs with twice differentiable log-partition functions, mean vector  $\mu$ , and a diagonal covariance matrix  $\Lambda_\sigma$ . Then the Fisher in-*

formation of the observation  $\langle u, \xi \rangle$  is

$$I(\mu) = \frac{uu^\top}{u^\top \Lambda_\sigma u} \quad (12)$$

As  $\|u\|_{\Lambda_\sigma} = \sqrt{u^\top \Lambda_\sigma u}$  is a valid norm, we additionally immediately observe that the Fisher information is invariant to the scaling of the measurement functional. This corollary allows us to restrict our search for informative measurements to unit spheres.

**Corollary 3.6** (Proof: See Corollary D.6). *The Fisher information for state estimation under independent noise is invariant to the scaling of the linear functional.*

Finally, when the noise is identically distributed, the form of the Fisher information matrix simplifies further.

**Corollary 3.7** (Proof: See Corollary D.7). *Suppose the elements of the noise vector are i.i.d. with variance  $\sigma^2$ . Then*

$$I(\mu) = \sigma^{-2} \frac{uu^\top}{\|u\|_2^2} \quad (13)$$

### 3.2 Transporting the CRLB Through Time

We now use the form of the Fisher information matrix in Theorem 3.5 to construct an alternative optimization objective in Theorem 3.9 with the same solution as the original infinite-horizon optimal experimental design problem. We observe that this reformulation implies the existence of local informative subspaces. Finally, we show that when the system converges to a one dimensional space, the reformulation can be solved exactly.

A key requirement for this analysis is a notion of stability of the nullspace of the Jacobian of the flow  $\varphi^\tau$ . For this lemma, sensitivities to normal perturbations in the neighborhood of the attracting manifold must decay monotonically. This allows the construction of an  $M - K$  dimensional vectorspace at each  $x$ , commonly called a vector bundle, for which the local sensitivity of the flow converges to zero for all vectors.

**Lemma 3.8** (Proof: See Lemma D.8). *If the system in Equation (1) converges to a  $K$ -dimensional smooth manifold  $\mathcal{M}$  such that there exists some  $\alpha > 0$  for which  $v^\top \left( [df_x] + [df_x]^\top \right) v + \alpha \|v\|^2 < 0$  for each  $v \in T_p^\perp \mathcal{M}$  normal to the manifold at each  $p \in \mathcal{M}$ , then there exists an  $M - K$  dimensional subspace  $S(x_0) \subset \mathbb{R}^M$  for which  $[d\varphi_{x_0}^\tau] u \rightarrow 0$  as  $\tau \rightarrow \infty$  for each  $u \in S(x_0)$ .*

We now introduce the main theorem in this work. By recalling that when a probability distribution is reparameterized the CRLB undergoes a linear transformation based on the Jacobian of the reparameterization, previously shown in Equation (6), the information in measurements is restricted to that contained in a  $K$ -dimensional subspace.

**Theorem 3.9** (Proof: See Theorem D.9). *Suppose the state space model in Equations (1) and (2) with noise satisfying the assumptions in Theorem 3.5 converges to a  $K$ -dimensional smooth manifold such that the Jacobian of the limiting flow is rank  $K$ . Then the linear functional which minimizes the CRLB for prediction in the infinite-horizon average cost formulation is the solution of*

$$\arg \max_{u: \|u\|_{(\Lambda_\sigma + \Sigma)} = 1} \sum_{i=1}^K \alpha_i \langle v_i, u \rangle_{\Sigma}^2 \quad (14)$$

for some  $\{\alpha_i\}$ , where  $v_i$  is the set of right singular vectors of the Jacobian of the limiting flow around the current state with nonzero singular vectors,  $\Sigma$  is the CRLB for estimating the current state based on the past measurements,  $\langle v, u \rangle_{\Sigma} = v^\top \Sigma u$ , and  $\|u\|_{(\Lambda_\sigma + \Sigma)} = \sqrt{u^\top (\Lambda_\sigma + \Sigma) u}$ .

It is worth noting that this problem reformulation additionally holds for minimizing the trace of the covariance of biased estimators, as the bias introduces an additional linear transformation on the low-dimensional subspace (Lehmann and Casella, 1998). The  $\{\alpha_i\}$  coefficients become dependent on the bias of the estimator, but dimensionality reduction insights still hold. Through Lagrange multipliers, we can compute the informative subspace exactly.

**Corollary 3.10** (Proof: See Corollary D.10). *Under the assumptions of Theorem 3.9, the optimal measurement vector exists in the subspace*

$$\text{Span} \left\{ (\Sigma + \Lambda_\sigma)^{-1} \Sigma v_i \right\}_{i=1}^K. \quad (15)$$

We can consider this to be a local informative subspace, or an informative vector bundle. By further observing that the information content is invariant to scaling of the measurement vector, we can construct an exact solution to the time-averaged infinite-horizon optimal experimental design problem.

**Corollary 3.11** (Proof: See Corollary D.11). *If, in addition to the assumptions of Theorem 3.9, the dynamical system converges to an isolated limit cycle, then*

$$u = (\Sigma + \Lambda_\sigma)^{-1} \Sigma v \quad (16)$$

is an optimal measurement vector, where  $v$  is the right singular vector.

At this point, it is worth considering practical issues which largely represent opportunities for future work. For any given state, computing  $\{v_i\}$  requires first-order local sensitivity analysis of the dynamical system. As this is a well-studied problem with numerous existing techniques, we do not address this problem further in this manuscript (Chen and Revels, 2016; Ma et al., 2021).

The informative subspace is dependent on the state of the system. In Section 4, we propose and discuss the usage of

the state estimate produced by an extended Kalman filter due to the ability to reuse computation.

Furthermore, we have not addressed how to optimize Equation (14) in the general case. We believe this local dimensionality reduction may be useful to reduce scaling issues in general approaches, reducing the action space to the dimensionality of the limiting manifold.

## 4 STATE ESTIMATION

In order to practically apply either the Bellman equation based dynamic programming algorithm or the proposed algorithm, we must include some form of state estimation. Due to its close connection with the propagation of the CRLB, we illustrate our algorithms using the extended Kalman filter (EKF).

The EKF is a two-step state estimation procedure which involves a prediction step followed by an estimation step. In its general form, the state vector is updated according to the system dynamics, then corrected based on the observation. The covariance matrices are propagated through a linearization of state transition operator and measurement operator. While the technique lacks the optimality guarantees of the original Kalman filter, it has remained a common method of state estimation in dynamical systems since its discovery in the 1960s (Smith et al., 1962).

In our application, the EKF is particularly efficient because the computationally expensive covariance update coincides with the CRLB update. Furthermore, as our measurements are scalars, the typically required matrix inversion in the estimation step becomes scalar division. When we incorporate our system model in Equations (1) and (2) and our noise model from Theorem 3.5, the EKF reduces to

$$\hat{x}_{i|i-1} = \varphi^\tau(\hat{x}_{i-1|i-1}), \quad (17)$$

$$\hat{x}_{i|i} = \hat{x}_{i|i-1} + \frac{\sum_i u_i (y_i - u_i^\top \hat{x}_{i|i-1})}{u_i^\top (\Lambda_\sigma + \sum_i) u}, \quad (18)$$

$$\Sigma_{i+1} = [d\varphi_x^\tau] \left( \Sigma_i - \frac{\sum_i u_i u_i^\top \Sigma_i}{u_i^\top (\Lambda_\sigma + \sum_i) u} \right) [d\varphi_x^\tau]^\top, \quad (19)$$

where the covariance update is identical to the CRLB update, and can thus be reused. The only additional computational cost to apply the EKF beyond the CRLB update is in Equation (18), which is negligible compared to the already required costs. A detailed description of operation counts is available in Appendix E.

## 5 EXPERIMENTAL RESULTS

The numerical simulations in this section were chosen to illustrate three key features of our proposed experimental design procedure. We begin by illustrating that, when the

state of the system is known exactly, both the proposed experimental design procedure and a computationally expensive dynamic programming baseline (Appendix C) achieve the expected CRLB reduction suggested by the dimensionality collapse. We then demonstrate that the application of the EKF is sufficiently accurate to maintain this improvement after a short delay. Finally, we demonstrate that the benefit improves linearly with the ambient dimensionality of the system. Simulations regarding computation time appear in Appendix E as further justification of the computational efficiency of the proposed algorithm.

Throughout this section, we simulate representative examples of three important classes of systems: linear systems, limit cycles, and chaotic systems. As many real-world systems exhibit these fundamental behaviors in different operating regimes, our chosen examples help to elucidate how the proposed algorithm will perform in each behavior. The three characteristic systems on which the observation policies were evaluated were a stable linear system defined with eigenvalues  $\lambda_1 = -10$  and  $\lambda_2 = -0.1$ , a Hopf bifurcation with dynamics  $\{\dot{x}_{(1)} = x_{(1)}(1 - \|x\|^2) - x_{(2)}; \dot{x}_{(2)} = x_{(2)}(1 - \|x\|^2) + x_{(1)}\}$ , and a Lorenz system with dynamics  $\{\dot{x}_{(1)} = 10(x_{(2)} - x_{(1)}); \dot{x}_{(2)} = x_{(1)}(28 - x_{(3)}) - x_{(2)}; \dot{x}_{(3)} = x_{(1)}x_{(2)} - \frac{8}{3}x_{(3)}\}$ , where  $x_{(i)}$  indicates the  $i$ 'th element of the state vector. While chaotic systems do not meet our theoretical convergence assumption, we include the Lorenz system to characterize performance on this related behavior. The specific parameters are available in Appendix H.

In Figure 3, we illustrate that the oracle solution of the optimal sampling procedure works across representative examples of each class of systems in which the true state is known to the decision process. A line corresponding to the optimal CRLB reduction enabled by the loss of dimensionality is included for reference. The key observation in this figure is that all techniques approximately capture the full improvement potential. Furthermore, as the plot represents the CRLB for state estimation rather than forecasting, we observe the fundamental trade-off made in this work: while we attain an improvement for future time-horizons, our ability to estimate the current state is negatively impacted in the short term. This is particularly visible in the dynamic programming plot in Panel C where the CRLB reduction does not materialize for about 250 samples, or 2.5 seconds of the system trajectory.

Next, we introduce the partially observed setting with the EKF in Panels A and B of Figure 4, normalizing the trace of the CRLB to that of the oracle solution in the Dynamic Programming approximation. Because the plots again illustrate the CRLB for the current state, we observe a short-term performance benefit for the EKF approach, as it is less aggressive in prioritizing the future. As the Lorenz system is not truly 1D, we see a similar benefit in Panel B. In these simulations, we observe empirically that the EKF based




Figure 3: Using uniform random sampling as a baseline, the oracle solutions for both the computationally expensive dynamic programming method and our proposed method capture the dimensionality collapse on each of the three systems. A: Linear System; B: Hopf Bifurcation; C: Lorenz System




Figure 4: (A,B): Trace of the CRLB normalized to that of the dynamic programming oracle solution for the (A) Hopf Bifurcation system and (B) Chaotic Lorenz system; (C): Trace of the CRLB of estimating the state of a limit cycle system normalized to that of the proposed method as a function of ambient dimensionality. Note the linear growth of the baseline random sampling approach.

policy often converges to similar levels of performance to the oracle policy, particularly in the proposed method.

Finally, in Panel C of Figure 4, we introduce a system based on expanding a Van der Pol oscillator, a common example of a system with a limit cycle, with additional dimensions governed by stable linear dynamics to illustrate the linear improvement as a function of the number of collapsing dimensions. In this system, the dynamics of the first two state variables are defined as  $\{\dot{x}_{(1)} = 3.5(x_{(1)} - \frac{1}{3}x_{(1)}^4 - x_{(2)}); \dot{x}_{(2)} = \frac{2}{7}x_{(1)}\}$ , while the remaining dimensions all follow  $\dot{x}_{(i)} = -x_{(i)}$ . Through these simulations, we remove any confounding issues and focus purely on the scaling with the dimensionality. The blue curve represents a baseline approach in which, at each decision point, the measurement vector is chosen uniformly at random over all unit vectors. This baseline represents an efficient method to collect measurements about the current state, particularly in the case when observations are made quickly relative to the rate of change of state in the system. We see that, in this greedy baseline approach, the trace of

the future CRLB increases linearly with dimensionality of the system while our proposed technique preserves a nearly constant value.

## 6 CONCLUSION

In this work, we introduced a dimensionality reduction method for the sequential selection of linear functional measurements of a vector-valued time series when the system converges to a low-dimensional manifold. Based on theoretical properties of natural exponential families of probability distributions, we reformulated the average cost infinite horizon problem for systems which converge to low-dimensional sets, resulting in a computationally efficient alternative to a Bellman equation approach.

Beyond theoretical results, we demonstrated the performance on three different common classes of dynamical systems: linear systems, limit cycles, and chaotic systems. Simulations were completed illustrating the performance



of the dynamic programming formulation and the 1D reformulation using an extended Kalman filter for state estimation. Simulations showed that the measurement selection performance remained constant irrespective of the ambient dimensionality of the space — in contrast with i.i.d. uniform random measurements which grow linearly.

The contributions in this work enable the collection of significantly more informative data than would be acquired through uniform observations. This improvement is attained irrespective of the eventual forecasting algorithm. We believe this work opens new directions in forecasting time series, and anticipate follow-on work regarding the co-design of observation and forecasting algorithms.

### Acknowledgements

This work was supported in part by the National Science Foundation under Grant 1936663.

### References

- Yaser S. Abu-Mostafa and Amir F. Atiya. Introduction to financial forecasting. *Applied Intelligence*, 6(3):205–213, Jul 1996. ISSN 1573-7497. doi: 10.1007/BF00126626.
- S.P. Asprey and S. Macchietto. Designing robust optimal dynamic experiments. *Journal of Process Control*, 12(4):545–556, 2002. ISSN 0959-1524. doi: 10.1016/S0959-1524(01)00020-8.
- Benny Avelin and Kaj Nyström. Neural ODEs as the deep limit of ResNets with constant weights. *Analysis and Applications*, 19(03):397–437, 2021. doi: 10.1142/S0219530520400023.
- Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Multiscale deep equilibrium models. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 5238–5250. Curran Associates, Inc., 2020.
- D. P. Bertsekas. *Dynamic Programming and Optimal Control: Approximate Dynamic Programming*, volume 2. Athena Scientific, Belmont, MA, 2017a.
- D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1, chapter 6, pages 296–424. Athena Scientific, Belmont, MA, 2017b.
- Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. ISSN 0027-8424. doi: 10.1073/pnas.1517384113.
- Lijuan Cao and Francis E.H Tay. Financial forecasting using support vector machines. *Neural Computing & Applications*, 10(2):184–192, May 2001. ISSN 1433-3058. doi: 10.1007/s005210170010.
- F.P. Casey, D. Baird, Q. Feng, R.N. Gutenkunst, J.J. Waterfall, C.R. Myers, K.S. Brown, R.A. Cerione, and J.P. Sethna. Optimal experimental design in an epidermal growth factor receptor signalling and down-regulation model. *IET Systems Biology*, 1:190–202(12), May 2007. ISSN 1751-8849.
- Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995. ISSN 08834237.
- Jiahao Chen and Jarrett Revels. Robust benchmarking in noisy environments. *arXiv preprint*, 1608.04295, 2016.
- Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.
- Raj Dandekar, Vaibhav Dixit, Mohamed Tarek, Aslan Garcia-Valadez, and Christopher Rackauckas. Bayesian neural ordinary differential equations. *arXiv preprint*, abs/2012.07244, 2021.
- Simon Danisch and Julius Krumbiegel. Makie.jl: Flexible high-performance data visualization for Julia. *Journal of Open Source Software*, 6(65):3349, 2021. doi: 10.21105/joss.03349.
- Holger Dette. Designing experiments with respect to ‘standardized’ optimality criteria. *Journal of the Royal Statistical Society. Series B (Methodological)*, 59(1):97–110, 1997. ISSN 00359246.
- Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural ODEs. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- M. Ehrendorfer. The Liouville equation and its potential usefulness for the prediction of forecast skill. part i: Theory. *Monthly Weather Review*, 122(4):703 – 713, 1994. doi: 10.1175/1520-0493(1994)122<0703:TLEAIP>2.0.CO;2.
- M. Ehrendorfer. The Liouville equation and atmospheric predictability. In T. Palmer and R. Hagedorn, editors, *Predictability of Weather and Climate*, page 59–98. Cambridge University Press, 2006. doi: 10.1017/CBO9780511617652.005.

- N. Benjamin Erichson, Lionel Mathelin, J. Nathan Kutz, and Steven L. Brunton. Randomized dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 18(4):1867–1891, 2019. doi: 10.1137/18M1215013.
- J. Fohring and E. Haber. Adaptive a-optimal experimental design for linear dynamical systems. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1138–1159, 2016. doi: 10.1137/15M1034738.
- William Gilpin. Deep reconstruction of strange attractors from time series. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 204–216. Curran Associates, Inc., 2020.
- Davis Gilton, Gregory Ongie, and Rebecca Willett. Deep equilibrium architectures for inverse problems in imaging. *arXiv preprint*, 2012.07944, 2021.
- Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Petr Hájek and Michal Johanis. Smooth approximations. *Journal of Functional Analysis*, 259(3):561–582, 2010. ISSN 0022-1236. doi: 10.1016/j.jfa.2010.04.020.
- Xun Huan and Youssef M. Marzouk. Sequential Bayesian optimal experimental design via approximate dynamic programming. *arXiv preprint*, 1604.08320, 2016.
- E. Kaiser, J. N. Kutz, and S. L. Brunton. Sparse identification of nonlinear dynamics for model predictive control in the low-data limit. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2219):20180335, 2018. doi: 10.1098/rspa.2018.0335.
- Hassan K. Khalil. *Nonlinear Systems*. Prentice-Hall, Upper Saddle River, NJ, 2002.
- John M. Lee. *Introduction to Riemannian Manifolds*. Springer Nature, Switzerland, 2018.
- E. L. Lehmann and George Casella. *Theory of Point Estimation*. Springer, New York, NY, 1998.
- Benjamin Letham, Portia A. Letham, Cynthia Rudin, and Edward P. Browne. Prediction uncertainty and optimal experimental design for learning dynamical systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 26(6):063110, 2016. doi: 10.1063/1.4953795.
- Yingbo Ma, Vaibhav Dixit, Michael J Innes, Xingjian Guo, and Chris Rackauckas. A comparison of automatic differentiation and continuous sensitivity analysis for derivatives of differential equation solutions. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–9, 2021. doi: 10.1109/HPEC49654.2021.9622796.
- Sonia Martínez and Francesco Bullo. Optimal sensor placement and motion coordination for target tracking. *Automatica*, 42(4):661–668, 2006. ISSN 0005-1098. doi: 10.1016/j.automatica.2005.12.018.
- Stefano Massaroli, Michael Poli, Jinkyoo Park, Atsushi Yamashita, and Hajime Asama. Dissecting neural ODEs. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 3952–3963. Curran Associates, Inc., 2020.
- Daniel A. Messenger and David M. Bortz. Weak SINDy for partial differential equations. *Journal of Computational Physics*, 443:110525, 2021. ISSN 0021-9991. doi: 10.1016/j.jcp.2021.110525.
- Francesco Mezzadri. How to generate random matrices from the classical compact groups. *arXiv preprint*, math-ph/0609050, 2006.
- Carl N. Morris. *Natural Exponential Families*. John Wiley & Sons, Ltd, 2006. ISBN 9780471667193. doi: 10.1002/0471667196.ess1759.pub2.
- Chirag Pabbaraju, Ezra Winston, and J Zico Kolter. Estimating Lipschitz constants of monotone deep equilibrium models. In *International Conference on Learning Representations*, 2021.
- C. Penland and T. Magorian. Prediction of Niño 3 sea surface temperatures using linear inverse modeling. *Journal of Climate*, 6(6):1067–1076, June 1993. ISSN 0894-8755. doi: 10.1175/1520-0442.
- Joshua L. Proctor, Steven L. Brunton, and J. Nathan Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016. doi: 10.1137/15M1013857.
- Christopher Rackauckas and Qing Nie. Differentialequations.jl—a performant and feature-rich ecosystem for solving differential equations in Julia. *Journal of Open Research Software*, 5(1):15, 2017.
- M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 0021-9991. doi: 10.1016/j.jcp.2018.10.045.
- Maziar Raissi and George Em Karniadakis. Hidden physics models: Machine learning of nonlinear partial differential equations. *Journal of Computational Physics*, 357:125–141, 2018. ISSN 0021-9991. doi: 10.1016/j.jcp.2017.11.039.
- Juri Ranieri, Amina Chebira, and Martin Vetterli. Near-optimal sensor placement for linear inverse problems. *IEEE Transactions on Signal Processing*, 62(5):1135–1146, 2014. doi: 10.1109/TSP.2014.2299518.

- Suman Ravuri, Karel Lenc, Matthew Willson, Dmitry Kangin, Remi Lam, Piotr Mirowski, Megan Fitzsimons, Maria Athanassiadou, Sheleem Kashem, Sam Madge, Rachel Prudden, Amol Mandhane, Aidan Clark, Andrew Brock, Karen Simonyan, Raia Hadsell, Niall Robinson, Ellen Clancy, Alberto Arribas, and Shakir Mohamed. Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597(7878):672–677, Sep 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03854-z.
- J. Revels, M. Lubin, and T. Papamarkou. Forward-mode automatic differentiation in Julia. *arXiv preprint*, 1607.07892, 2016.
- Peter J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010. doi: 10.1017/S0022112010001217.
- Jack Sherman and Winifred J. Morrison. Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix. *The Annals of Mathematical Statistics*, 21(1):124 – 127, 1950. doi: 10.1214/aoms/1177729893.
- Xingjian Shi, Zhihan Gao, Leonard Lausen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Deep learning for precipitation nowcasting: A benchmark and a new model. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Gerald L. Smith, Stanley F. Schmidt, and Leonard A. McGee. Application of statistical filtering theory to the optimal estimation of position and velocity on board a circumlunar vehicle. Technical report, NASA Ames Research Center, 1962.
- Steven H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering*. CRC Press, Boca Raton, FL, 2015.
- D. M. Titterton. Aspects of optimal design in dynamic systems. *Technometrics*, 22(3):287–299, 1980. doi: 10.1080/00401706.1980.10486160.
- Mark K. Transtrum and Peng Qiu. Optimal experiment selection for parameter estimation in biological differential equation models. *BMC Bioinformatics*, 13(1):181, Jul 2012. ISSN 1471-2105. doi: 10.1186/1471-2105-13-181.
- L.N. Trefethen and D. Bau. *Numerical Linear Algebra*, page 59. SIAM, Philadelphia, PA, 1997.
- Matthew O. Williams, Ioannis G. Kevrekidis, and Clarence W. Rowley. A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, Dec 2015. ISSN 1432-1467. doi: 10.1007/s00332-015-9258-5.
- Themistoklis C. Xygkis, George N. Korres, and Nikolaos M. Manousakis. Fisher information-based meter placement in distribution grids via the d-optimal experimental design. *IEEE Transactions on Smart Grid*, 9(2): 1452–1461, 2018. doi: 10.1109/TSG.2016.2592102.
- Linan Zhang and Hayden Schaeffer. On the convergence of the SINDY algorithm. *Multiscale Modeling & Simulation*, 17(3):948–972, 2019. doi: 10.1137/18M1189828.

## APPENDIX

The appendix is organized as follows. Appendix A describes the systems shown in Table 1. Appendix C provides derivations of the performance of the dynamic programming base case. Appendix D provides the proofs for Section 3. Appendix E contains counts of multiplications required in the algorithms, as well as timing simulations and details regarding the gradient ascent procedure. Appendix F discusses potential numerical issues in the work relating to the requirement to invert approximately singular matrices, as well as the chosen solution to the problem. Appendix G describes the interface for incorporating new dynamical systems into the library produced in the course of this work. Finally, Appendix H includes all parameter used in generating the plots in the main body of the manuscript.

### A TABLE 1 DETAILS

Representative trajectories for the systems in Table 1 are shown in Figure 5. Initializations were chosen in an arbitrary, often randomized manner as the chosen systems are known to contain only a small number of stable global behaviors. In particular, all systems other than the gradient flow globally converge to a fixed behavior, while the chosen gradient flow contains only two stable equilibria.




Figure 5: From left to right, Morse Flow, Van Der Pol, Hopf, Lorenz, Heat equation

**Gradient Flow** Let  $h : \mathbb{R}^M \rightarrow \mathbb{R}$  on some bounded interval be a Morse function. The gradient flow is defined as

$$\dot{x} = -\nabla h(x). \quad (20)$$

The key property of a Morse function in this example is that all critical points occur in the interior of the manifold  $\mathcal{M}/\partial\mathcal{M}$  and that no critical points are degenerate. The lack of degeneracy is well-known to imply that critical points are isolated, thus resulting in a discrete-set of points.

In the particular simulation for Figure 5,  $h(x) = x^4 - 5x^2 + 4$ .

#### Van Der Pol Oscillator

$$\dot{x} = y \quad (21)$$

$$\dot{y} = \mu(1 - x^2)y - x \quad (22)$$

The Van Der Pol Oscillator is well-known to have a single globally stable limit cycle, though there is no closed form.

In the Figure 5,  $\mu = 1$ .

#### Hopf Bifurcation

$$\dot{x} = x(\lambda + b(x^2 + y^2)) - y \quad (23)$$

$$\dot{y} = y(\lambda + b(x^2 + y^2)) + x \quad (24)$$

The parameters are chosen such that the system exhibits a stable limit cycle, for example  $b = -1$  and  $\lambda = 1$  in Figure 5. Under the above choice of parameters, there is a stable limit cycle at  $x^2 + y^2 = 1$ . By substituting in the constraint into the

above expression, the system becomes

$$\dot{x} = -y \quad (25)$$

$$\dot{y} = x \quad (26)$$

which is linear with purely imaginary eigenvalues.

### Lorenz System

$$\dot{x} = \sigma(y - x) \quad (27)$$

$$\dot{y} = x(\rho - z) - y \quad (28)$$

$$\dot{z} = xy - \beta z \quad (29)$$

The almost canonical set of parameters for the system is  $\sigma = 10$ ,  $\rho = 28$ ,  $\beta = 8/3$ , which generates the butterfly-shaped attractor pictured in Figure 5. In the case of this well-known chaotic dynamical system, the low-dimensional structure is less obvious. In fact, the attractor is not embedded in a 2D manifold at all. Instead, due to deep connections between chaos and fractals, chaotic attractors are often described by their fractal dimension. While there are many definitions of fractal dimensions, a common definition is the Hausdorff dimension, defined based on covering the space with arbitrarily small balls. In the case of the Lorenz system with the chosen parameters, the dimensionality is well-known to be approximately 2.06.

**Heat Equation** The heat equation is one of the most fundamental linear partial differential equations, and represents spatial diffusion over time.

$$\frac{\partial u}{\partial t} = \alpha \Delta u \quad (30)$$

In Figure 5, the heat equation is simulated according to a finite-difference discretization with  $\alpha = 0.2$ . The only eigenfunction of the system which does not exponentially decay is the average value of  $u$ , and so the system converges only to a 1D space.

## B CRAMÉR–RAO LOWER BOUND

In this appendix, we provide a brief overview of the well-known Cramér–Rao Lower Bound (CRLB), with regularity conditions adapted from Lehmann and Casella (1998).

The multivariate CRLB, sometimes known as an information inequality, is a matrix inequality on the covariance matrix of all estimators of an unknown, deterministic parameter of a probability distribution. Specifically, let  $p(\xi; \theta)$  be a family of density functions of a random variable  $\xi$  parameterized by a vector  $\theta$  such that  $p$  satisfies the regularity conditions:

1.  $\Theta$  is an open set of possible parameters  $\theta$ .
2. The densities  $p(\cdot; \theta)$  have common support.
3. For each  $\xi \in \Xi$ ,  $\theta \in \Theta$ , and  $i \in 1, \dots, M$ ,  $\frac{\partial p(\xi; \theta)}{\partial \theta_i} < \infty$ .

Then, for any estimator of the parameter  $\hat{\theta}$  such that  $i \in 1, \dots, M$ ,  $\frac{\partial}{\partial \theta_i} \mathbb{E} [\hat{\theta}] = \mathbb{E} \left[ \frac{\partial}{\partial \theta_i} \hat{\theta} \right]$ ,

$$\Sigma \geq \left[ \frac{d\mathbb{E} [\hat{\theta}]}{d\theta} \right]^\top J^{-1}(\theta) \left[ \frac{d\mathbb{E} [\hat{\theta}]}{d\theta} \right], \quad (31)$$

where  $\Sigma$  is the covariance matrix of the estimator and  $J^{-1}(\theta)$  is the inverse Fisher information matrix. The matrix inequality is interpreted such that a matrix  $A \geq 0$  is positive semi-definite. When applied to unbiased estimators where  $\mathbb{E} [\hat{\theta}] = \theta$ , the CRLB reduces to the exceedingly common unbiased form

$$\Sigma \geq J^{-1}(\theta). \quad (32)$$

**Algorithm 1** Interpolation Point Sampling

---

|   |   |
|---|---|
| <b>Input:</b> Eigenvalue Distribution $\mu$ |   |
| Sample diagonal matrix $\Lambda \sim \mu$   | $(\Lambda \in \mathbb{R}^{M \times M})$ |
| Sample $G \sim \mathcal{N}(0, I)$           | $(G \in \mathbb{R}^{M \times M})$       |
| $Q, R \leftarrow QR(G)$                     | $(Q \sim \text{Haar Measure})$          |
| $x_i \leftarrow Q^\top \Lambda Q$           | $(x_i \in \mathbb{R}^{M \times M})$     |

---

The Fisher information matrix is the covariance of the derivative of the log-likelihood function, i.e.

$$J(\theta) = \mathbb{E} \left[ \left( \frac{d}{d\theta} \log p(\xi; \theta) \right) \left( \frac{d}{d\theta} \log p(\xi; \theta) \right)^\top \right]. \quad (33)$$

Of note in this work is the impact of diffeomorphic reparameterization on the Fisher information matrix, which can be seen immediately by applying the chain rule to Equation (33).

## C DYNAMIC PROGRAMMING BASELINE

While the dynamic programming problem maps directly into the stochastic control and reinforcement learning frameworks, there are still a number of application-specific details in all such algorithms. In this section, we address the choice of an appropriate discretization scheme for the state space.

First, we need to consider the structure of the state space and the action space. The policy is dependent both on the previous CRLB and on the previous system state. But note that the CRLB is a positive semi-definite (PSD) matrix, and so if the system state is in  $\mathbb{R}^M$ , then the space of CRLB matrices is embedded in  $\mathbb{R}^{(M^2+M)/2}$ . Furthermore, the space of PSD matrices trivially forms a cone, which further reduces the space. An equivalent representation of the space is through the diagonalization  $A = U\Lambda U^\top$ , where  $\Lambda$  is a diagonal matrix and  $U \in O(M)$  is an orthonormal matrix.

The space of PSD matrices has a non-trivial structure, yet we have strong heuristics for the problem. Thus, we opt to approximate through the interpolation of random sample points. We denote the set of sample points by  $\mathcal{X} = \{x_1, \dots, x_N\}$  and our approximation of the value function to be  $\hat{J}_{\mathcal{X}}(x, \Sigma)$ . We first describe the sampling distribution of  $\mathcal{X}$ , before discussing the interpolation methods.

**Random Sampling** A practical approach to generating random PSD matrices is to decompose the problem into sampling from the space of orthonormal matrices and sampling non-negative eigenvalues. Not only is this approach straightforward, but as we have more intuition regarding the distribution of eigenvalues, it allows easy tuning of the approximation of the distribution. The full sampling routine is shown in Algorithm 1.

Using a naive prior to model the problem, we begin with sampling the orthonormal matrix components from a Haar measure. The Haar measure can be thought of as a uniform distribution — it is the unique measure that is invariant to the action of elements in the group (Mezzadri, 2006). It is exceedingly easy to generate random orthonormal matrices from this measure. It can be shown that due to the isotropic behavior, diagonalizing vectors of a matrix populated with i.i.d. Gaussian random variables results in such a distribution. Practically, sampling is done using QR factorization, typically with an eigenvalue correction term that cancels in similarity transforms (Mezzadri, 2006).

Understanding that each measurement is optimized to minimize the eigenvalues of the CRLB, we follow the heuristic that smaller eigenvalues will be encountered more frequently and sample from an exponential distribution. The parameter of the distribution is then a hyperparameter of the problem.

**Interpolation Methods** While one could choose any one of a large number of interpolation methods, in this work, we only analyze nearest neighbor and a locally averaged perturbation of nearest neighbor. While alternative approaches may yield higher accuracy, the approximation quality for classification tasks is extremely well studied with strong guarantees, typically providing both upper and lower bounds due to the connection to density estimation techniques.

It is clear that the CRLB of our measurement system will be unevenly distributed in any practical setting: as we gather more data, the CRLB decreases. Thus, we would like to weight our approximation quality to account for this phenomenon.




Figure 6: Numerically estimated expected minimum distance as a function of the size of the dataset.

Here, we provide a bound for function approximation accuracy under a non-uniform measure defined by a probability distribution.

**Lemma C.1.** *Assume for every  $x \in \mathcal{X}$ , the approximation of the value function is exact, i.e.  $\hat{J}_{\mathcal{X}}(x) = J(x)$ , where  $\hat{J}_{\mathcal{X}}$  is our approximation. Furthermore, assume  $J$  is Lipschitz with coefficient  $L$ . Then the mean absolute error of the function approximation is upper bounded by*

$$\mathbb{E}_{\mathcal{X},x} [|\hat{J}_{\mathcal{X}}(x) - J(x)|] \leq L \mathbb{E}_{\mathcal{X},x} \left[ \min_{x' \in \mathcal{X}} \|x - x'\| \right] \quad (34)$$

*Proof.* By the Lipschitz assumption,  $|J(x) - J(\tilde{x})| \leq L\|x - \tilde{x}\|$ , and so

$$\mathbb{E}_{\mathcal{X},x} [|J(x) - J(\tilde{x})|] \leq L \mathbb{E}_{\mathcal{X},x} [\|x - \tilde{x}\|]. \quad (35)$$

By construction of nearest neighbor regression,  $\hat{J}_{\mathcal{X}}(x) = J(x')$  where  $x' = \arg \min_{\tilde{x} \in \mathcal{X}} \|x - \tilde{x}\|$ . Thus

$$\mathbb{E}_{\mathcal{X},x} [|\hat{J}_{\mathcal{X}}(x) - J(x)|] = \mathbb{E}_{\mathcal{X},x} [|J(x) - J(x')|] \leq L \mathbb{E}_{\mathcal{X},x} \left[ \min_{x' \in \mathcal{X}} \|x - x'\| \right] \quad (36)$$

□

The expected minimum distance to the dataset is dependent on the distribution and is non-trivial to compute analytically. We include a simulation in Figure 6 illustrating the expected minimum distance as a function of the number of samples. One important observation is the well-known curse of dimensionality, and so this particular function approximation would scale poorly to higher dimensions.

Finally, the nearest neighbor model fails to produce reasonable decisions when the CRLB changes by less than twice the distance between sample points — resulting in the state staying constant in value iteration steps. To prevent this issue, we apply local averaging motivated by linear interpolation. In particular, in our local averaging model, an approximate point is evaluated as

$$\hat{J}_{\mathcal{X}}^{(loc)}(\hat{x}) = \frac{\sum_{x \in \mathcal{X}} J(x) \min(0, d_{max} - \|x - \hat{x}\|)}{\sum_{x \in \mathcal{X}} \min(0, d_{max} - \|x - \hat{x}\|)}, \quad (37)$$

where  $d_{max}$  is a chosen maximum distance for averaging and  $\hat{J}_{\mathcal{X}}(\hat{x})$  is replaced by a nearest neighbor approximation if the denominator becomes 0. We can readily bound the perturbation of our locally averaged interpolation from our nearest neighbor interpolation based on the Lipschitz assumption.

**Theorem C.2.** *The mean absolute error of the locally averaged approximation is at most*

$$\mathbb{E}_{\mathcal{X},x} [|\hat{J}_{\mathcal{X}}^{(loc)}(x) - J(x)|] \leq L \left( \mathbb{E}_{\mathcal{X},x} \left[ \min_{x' \in \mathcal{X}} \|x - x'\| \right] + 2d_{max} \right). \quad (38)$$

*Proof.* The inequality comes from observing the maximum perturbation from the nearest neighbor solution.

First, observe that if the  $d_{max}$  ball around  $x$  contains 1 or fewer points, then there is no change to the approximation. Next, consider the case when there are exactly two representation points in the ball. Then  $\hat{J}_{\mathcal{X}}^{(loc)}(x) = \frac{(d_{max}-d_1)J(x_1)+(d_{max}-d_2)J(x_2)}{2d_{max}-d_1-d_2}$ , where  $d_i$  is the distance from  $x$  to  $x_i$ . Now suppose  $d_1 < d_2$ , then  $\hat{J}_{\mathcal{X}}(x) = J(x_1)$ .

$$\hat{J}_{\mathcal{X}}(x) - \hat{J}_{\mathcal{X}}^{(loc)}(x) = (J(x_1) - J(x_2)) \frac{d_{max} - d_2}{(d_{max} - d_1) + (d_{max} - d_2)} \quad (39)$$

Now,  $d_1 < d_2 \Rightarrow (d_{max} - d_1) > (d_{max} - d_2) \Rightarrow \frac{d_{max}-d_2}{(d_{max}-d_1)+(d_{max}-d_2)} \leq \frac{1}{2}$ . It then follows that

$$\left| \hat{J}_{\mathcal{X}}(x) - \hat{J}_{\mathcal{X}}^{(loc)}(x) \right| \leq \frac{1}{2} |J(x_1) - J(x_2)| \quad (40)$$

The distance between  $x_1$  and  $x_2$  is upper bounded by  $2d_{max}$  which represents opposite ends of the ball. Thus, by the Lipschitz assumption

$$\left| \hat{J}_{\mathcal{X}}(x) - \hat{J}_{\mathcal{X}}^{(loc)}(x) \right| \leq Ld_{max}. \quad (41)$$

Consider the implication of adding additional points to the ball. Note that again, the worst-case scenario is when all points are nearly equidistant from the center, resulting in

$$\left| \hat{J}_{\mathcal{X}}(x) - \hat{J}_{\mathcal{X}}^{(loc)}(x) \right| \leq \frac{1}{N} \left| J(x_1) - \sum_{i=2}^N J(x_i) \right|. \quad (42)$$

Once again,  $J(x_1) - 2Ld_{max} \leq J(x_i) \leq J(x_1) + 2Ld_{max}$ , and thus

$$\left| \hat{J}_{\mathcal{X}}(x) - \hat{J}_{\mathcal{X}}^{(loc)}(x) \right| \leq \left( \frac{N-1}{N} \right) 2Ld_{max} \quad (43)$$

By observing  $(N-1)/N < 1$  for all  $N > 1$ , we find one irrespective of the number of points in the ball.

$$\left| \hat{J}_{\mathcal{X}}(x) - \hat{J}_{\mathcal{X}}^{(loc)}(x) \right| \leq 2Ld_{max} \quad (44)$$

Finally, apply the triangle inequality to conclude the original statement, i.e.

$$\left| \hat{J}_{\mathcal{X}}^{(loc)} - J(x) \right| = \left| \hat{J}_{\mathcal{X}}^{(loc)} - \hat{J}_{\mathcal{X}}(x) + \hat{J}_{\mathcal{X}}(x) - J(x) \right| \quad (45)$$

$$\leq \left| \hat{J}_{\mathcal{X}}^{(loc)} - \hat{J}_{\mathcal{X}}(x) \right| + \left| \hat{J}_{\mathcal{X}}(x) - J(x) \right| \quad (46)$$

$$\leq L \left( 2d_{max} + \min_{x' \in \mathcal{X}} \|x - x'\| \right) \quad (47)$$

$$\Rightarrow \mathbb{E}_{\mathcal{X},x} \left[ \left| \hat{J}_{\mathcal{X}}^{(loc)}(x) - J(x) \right| \right] \leq L \left( \mathbb{E}_{\mathcal{X},x} \left[ \min_{x' \in \mathcal{X}} \|x - x'\| \right] + 2d_{max} \right). \quad (48)$$

□

Conveniently, as  $N$  increases, we can decrease  $d_{max}$  while retaining the averaging behavior. Thus, we choose  $d_{max}$  proportional to expected minimum distance to preserve the behavior of the bound.

## D PROOFS FOR SECTION 3

**Lemma D.1** (Context: See Lemma 3.1). *Given a finite set of independent random variables  $\{\xi_{(i)}\}$  in NEFs with log-partition functions  $\{A_i\}$ , parameters  $\{\theta_i\}$ , and a set of real-valued weights  $\{\alpha_i\}$ , then  $\tilde{\xi} = \sum_i \alpha_i \xi_{(i)}$  is in a NEF with a log partition function*

$$A(\theta) := \sum_i A_i(\theta + \alpha_i \theta_i). \quad (49)$$



*Proof.* Without loss of generality, we prove the case of two random variables. The result can then be trivially extended to any finite set, either by induction or through a modification of the below proof.

The moment generating function of random variable  $X$  in a NEF is well-known to be

$$M_X(s) = \exp\{A(\theta + s) - A(\theta)\}. \quad (50)$$

If we denote the log-partition function of  $\xi_i$  to be  $A_i$ , then

$$M_{\xi_1 + \xi_2}(s) = \exp\{A_1(\theta_1 + s) + A_2(\theta_2 + s) - A_1(\theta_1) - A_2(\theta_2)\}. \quad (51)$$

We then define the new log-partition function  $A : s \mapsto A_1(\theta_1 + s) + A_2(\theta_2 + s)$ . Thus, by the structure of the MGF, the distribution itself is in a NEF.  $\square$

**Lemma D.2** (Context: See Lemma 3.2). *The Fisher information of a random variable in a NEF is the variance of the random variable.*

*Proof.* By definition of the NEF, the log-likelihood takes the form

$$\theta\xi - A(\theta) + B(\xi). \quad (52)$$

Recall that  $A'(\theta) = \mathbb{E}[\xi]$ . Thus, by computing the derivative, the Fisher information is

$$\mathbb{E}[(\xi - A'(\theta))^2] = \mathbb{E}[(\xi - \mathbb{E}[\xi])^2] \quad (53)$$

$\square$

**Lemma D.3** (Context: See Lemma 3.3). *The diagram in Figure 2 commutes for any vector  $\xi$  of independent random variables from NEFs.*

*Proof.* We must show that  $\{g_i\} \circ h = h \circ g$ . It suffices to show that each connection as drawn is accurate. The application of  $g$  and  $\{g_i\}$  are well-known from the NEF. For completeness, recall that the expected value of the score function of a random variable is 0, i.e.

$$\mathbb{E}\left[\frac{\partial}{\partial\theta} \log p(x; \theta) | \theta\right] = \frac{\partial}{\partial\theta} \mathbb{E}[\log p(x; \theta) | \theta] = 0. \quad (54)$$

Then

$$\mathbb{E}\left[\frac{\partial}{\partial\theta} (x\theta - A(\theta) + B(x))\right] = \mathbb{E}\left[x - \frac{\partial}{\partial\theta} A(\theta)\right] \quad (55)$$

$$= \mu - g(\theta) = 0 \quad (56)$$

$$\Rightarrow \mu = g(\theta) \quad (57)$$

Furthermore,  $h : \{\mu_i\} \mapsto \mu$  follows by linearity of expectation.

Thus, the only tricky connection is to show  $\theta = h(\{\theta_i\})$ . Fortunately, this follows by observing we have flexibility in the definition of the NEF of  $\xi$ . In particular, adding a constant offset to our definition of  $\theta$  does not impact the definition of the family. Thus, we are free to center  $\theta$  and our NEF definition around  $\theta = h(\{\theta_i\})$ .  $\square$

**Lemma D.4** (Context: See Lemma 3.4). *Assume the log-partition function  $A$  of a NEF is twice differentiable and that the second moment of the associated random variable exists. Then the inverse of  $g(\theta) := \frac{\partial}{\partial\Theta} A(\Theta) \Big|_{\Theta=\theta}$  exists if and only if the random variable is non-degenerate for almost all  $\theta$ .*

*Proof.* First, assume that  $\xi$  is non-degenerate. Then, we will show  $g$  is strictly monotonic and thus invertible. Recall that  $g := A'$ , where  $A'$  is the derivative of  $A$  with respect to  $\theta$ . Furthermore, recall that under the above regularity conditions,  $A''$  is the Fisher information. Finally, by Lemma 3.2.

$$\frac{\partial}{\partial\theta} g(\theta) = A''(\theta) = \mathbb{E}[(\xi - \mathbb{E}[\xi])^2] = \text{Var}(\xi) > 0. \quad (58)$$

Thus,  $g$  is strictly increasing and is therefore invertible.

Now, by contrapositive, assume there exists an interval for which  $\xi$  deterministically takes a constant value. Then by the above argument,  $\frac{\partial}{\partial\theta} g(\theta) = 0$  for all  $\theta$  on the interval, and thus  $g$  is not invertible on the interval.  $\square$

**Theorem D.5** (Context: See Theorem 3.5). *Suppose  $\xi$  is a vector of independent random variables from NEFs with twice differentiable log-partition functions, mean vector  $\mu$ , and a diagonal covariance matrix  $\Lambda_\sigma$ . Then the Fisher information of the observation  $\langle u, \xi \rangle$  is*

$$I(\mu) = \frac{uu^\top}{u^\top \Lambda_\sigma u} \quad (59)$$

*Proof.* The proof follows from the transformation to reparameterized Fisher information. That is, we can instead write  $\bar{\theta} = T(\mu)$ .

$$I(\mu) = [dT_\mu] I(\bar{\theta}) [dT_\mu]^\top \quad (60)$$

$$= [dh_\mu] dg_{\bar{\mu}}^{-1} I(\bar{\theta}) dg_{\bar{\mu}}^{-1} [dh_\mu]^\top \quad (61)$$

Now, recall that  $I(\bar{\theta}) = dg_{\bar{\mu}} = \text{Var}(\bar{\xi}) = \sum_i u_{(i)}^2 \sigma_i^2 = u^\top \Lambda_\sigma u$ , where  $u_{(i)}$  represents the  $i$ 'th element of  $u$ . Then by noting the inverse commutes with the derivative and that scalar multiplication commutes,

$$I(\mu) = [dh_\mu] [dh_\mu]^\top \frac{1}{u^\top \Lambda_\sigma u}. \quad (62)$$

Finally, observe  $[dh_\mu] = u$  to conclude the proof.  $\square$

**Corollary D.6** (Context: See Corollary 3.6). *The Fisher information for state estimation under i.i.d. noise is invariant to the scaling of the linear functional.*

*Proof.* Follows immediately from the previous corollary by writing it as

$$I(\mu) = \sigma^{-2} \frac{uu^\top}{\|u_i\|^2} = \sigma^{-2} \left( \frac{u}{\|u\|} \right) \left( \frac{u}{\|u\|} \right)^\top, \quad (63)$$

and observing the absolute homogeneity of norms.  $\square$

**Corollary D.7** (Context: See Corollary 3.7). *Suppose the elements of the noise vector are i.i.d. with variance  $\sigma^2$ . Then*

$$I(\mu) = \sigma^{-2} \frac{uu^\top}{\|u\|^2} \quad (64)$$

*Proof.* As the noise is identically distributed, it has the same variance and the covariance matrix becomes a scaled version of the identity. From Theorem 3.5

$$I(\mu) = \frac{uu^\top}{u^\top \Lambda_\sigma u} = \sigma^{-2} \frac{uu^\top}{u^\top u} = \sigma^{-2} \frac{uu^\top}{\|u\|^2} \quad (65)$$

$\square$

**Lemma D.8** (Context: See Lemma 3.8). *If the system in Equation (1) converges to a  $K$ -dimensional smooth manifold  $\mathcal{M}$  such that there exists some  $\alpha > 0$  for which  $v^\top \left( [df_x] + [df_x]^\top \right) v + \alpha \|v\|^2 < 0$  for each  $v \in T_p^\perp \mathcal{M}$  normal to the manifold at each  $p \in \mathcal{M}$ , then there exists an  $M - K$  dimensional subspace  $S(x_0) \subset \mathbb{R}^M$  for which  $[d\varphi_{x_0}^\tau] u \rightarrow 0$  as  $\tau \rightarrow \infty$  for each  $u \in S(x_0)$ .*

Note that this proof is long. As such, we have split some of the significant intermediate steps into italicised statements. Proofs of these intermediate statements are terminated with  $\blacksquare$ , while the full proof is terminated with  $\square$ .

*Proof.* By considering the sensitivity equation (Khalil, 2002)

$$\dot{x} = f(x) \quad (66)$$

$$\left[ \frac{d\dot{\varphi}^\tau}{dx} \right] = \left[ \frac{df}{dx} \right] \left[ \frac{d\varphi^\tau}{dx} \right], \quad (67)$$

we will see that the normal sensitivity is asymptotically stable once sufficiently close to the manifold. By our assumption that the trajectory converges, we can represent the state using Fermi coordinates on a tubular neighborhood of the manifold.

That is, we can represent the space as the product of the manifold  $\mathcal{M}$  and an  $M - K$  dimensional vectorspace representing the normals of the manifold. If  $\mathcal{M}$  is smooth, such a neighborhood always exists by the tubular neighborhood theorem (Lee, 2018).

For any  $\tau$ , there exists a vector bundle  $S_\tau(x) \subset \mathbb{R}^M$  such that  $\frac{d\varphi}{dx} S_\tau(x) = T_{\varphi^\tau(x)}^\perp \mathcal{M}$  is normal to the manifold.

Because  $\varphi^\tau$  is a diffeomorphism,  $\frac{d\varphi^\tau}{dx}$  is full rank. We can construct any vector  $v = \left[\frac{df}{dx}\right]^{-1} u$  for some  $u$ . Therefore, there exists some subspace  $S(x)$  such that  $T_p^\perp \mathcal{M} = \left[\frac{df}{dx}\right]^{-1} S(x)$  where  $p = \pi_{\mathcal{M}}(\varphi^\tau(x))$  is the projection of  $\varphi^\tau(x)$  onto the manifold.  $\blacksquare$

Let  $\Pi_{\mathcal{M}}^\perp$  be the orthogonal projection onto the normal space. Then  $\lim_{\tau \rightarrow \infty} v_0^\top \left[\frac{d\varphi^\tau}{dx}\right]^\top \Pi_{\mathcal{M}}^\perp \left[\frac{d\varphi^\tau}{dx}\right] v_0 = 0$ .

Observe that  $\left[\frac{d\varphi^\tau}{dx}\right] v_0$  is the solution of the sensitivity equation  $\dot{v} = \left[\frac{df}{dx}\right] v$ . Let  $V(v, \tau) = v^\top \Pi_{\mathcal{M}}^\perp v$ . We can compute the time derivative of  $V(v, \tau) = v^\top \Pi_{\mathcal{M}}^\perp v$  and show that the quantity decreases monotonically. Note that through the use of Fermi coordinates, we have removed the dependency on  $\tau$  in  $\Pi_{\mathcal{M}}^\perp$ . The derivative then becomes

$$\dot{V}(v, \tau) = v^\top \left( \Pi_{\mathcal{M}}^\perp \left[\frac{df}{dx}\right] \Big|_{x=x_0} + \left[\frac{df}{dx}\right] \Big|_{x=x_0}^\top \Pi_{\mathcal{M}}^\perp \right) v \quad (68)$$

$$= v^\top \left( \Pi_{\mathcal{M}}^\perp \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x_0)} + \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x_0)}^\top \Pi_{\mathcal{M}}^\perp \right) v + v^\top (\Pi_{\mathcal{M}}^\perp \Gamma(x) + \Gamma(x)^\top \Pi_{\mathcal{M}}^\perp) v, \quad (69)$$

where  $\|\Gamma(x)\| < L\epsilon$  by the Lipschitz assumption. Because  $\Pi_{\mathcal{M}}^\perp$  is an orthogonal projection,  $\|\Pi_{\mathcal{M}}^\perp\| = 1$ . The operator norms can then be used to upper bound the derivative as

$$\dot{V}(v, \tau) \leq v^\top \left( \Pi_{\mathcal{M}}^\perp \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x)} + \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x)}^\top \Pi_{\mathcal{M}}^\perp \right) v + 2\epsilon L \|v\|^2 \quad (70)$$

$$= 2v^\top \Pi_{\mathcal{M}}^\perp \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x)} v + 2\epsilon L \|v\|^2. \quad (71)$$

Decompose  $v$  into  $v_\perp \in T_p^\perp \mathcal{M}$  and  $v_\parallel \in T_p \mathcal{M}$  such that  $v = v_\perp + v_\parallel$  and reorder terms as

$$\dot{V}(v, \tau) \leq 2v^\top \Pi_{\mathcal{M}}^\perp \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x)} (v_\perp + v_\parallel) + 2\epsilon L \|v\|^2 \quad (72)$$

$$= 2v_\perp^\top \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x)} (v_\perp + v_\parallel) + 2\epsilon L \|v\|^2 \quad (73)$$

$$= 2v_\perp^\top \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x)} v_\perp + 2v_\perp^\top \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x)} v_\parallel + 2\epsilon L \|v\|^2 \quad (74)$$

$$\leq 2(\epsilon L - \alpha) \|v_\perp\|^2 + 2\epsilon L \|v_\perp\| \|v_\parallel\| + 2v_\perp^\top \left[\frac{df}{dx}\right] \Big|_{x=\pi_{\mathcal{M}}(x)} v_\parallel, \quad (75)$$

where the final line comes from our eigenvalue assumption and the triangle inequality.

By recalling that  $f$  determines the geometry of the manifold through the trajectories, the rightmost term can be seen to represent the derivative of the normal vector along a curve in the manifold. By further considering the osculating circle at the point on the associated curve, or the unique circle which matches at least the first two derivatives (Lee, 2018), we can see that the directional derivatives of the normal vectors along the tangent space are zero. This is because the derivative of the normal component of the circle is zero along the tangent, i.e.  $\frac{d}{dx} \cos(x)|_{x=0} = \sin(0) = 0$ . Thus

$$\dot{V}(v, \tau) \leq 2(\epsilon L - \alpha) \|v_\perp\|^2 + 2\epsilon L \|v_\perp\| \|v_\parallel\|, \quad (76)$$

which is strictly less than zero if

$$\frac{\|v_\perp\|}{\|v_\parallel\|} \geq \frac{\epsilon L}{\alpha - \epsilon L}. \quad (77)$$

Thus, the normal components of all vectors must decay to zero as  $\epsilon \rightarrow 0$ , and so  $\Pi_{\mathcal{M}}^\perp \left[ \frac{d\varphi^\tau}{dx} \right] v_0 \rightarrow 0$  as  $\tau \rightarrow \infty$ .  $\blacksquare$

Finally, by decomposing the transformation into two distinct functions through the semigroup property, we conclude the overall proof. Consider  $\varphi^{\tau+T}$ , where  $T$  is sufficiently large that the distance from  $\mathcal{M}$  is at most  $\epsilon$ . Decompose the function as  $\varphi^{\tau+T} = \varphi^\tau \circ \varphi^T$  and apply the chain rule as

$$\frac{d\varphi^{\tau+T}}{dx} = \frac{d\varphi^\tau}{d\varphi^T(x)} \frac{d\varphi^T}{dx}. \quad (78)$$

From our intermediate results, for any  $\delta$  we can always choose an appropriate  $\tau$  and  $v$  such that

$$\frac{d\varphi^T}{dx} v \in T_p \mathcal{M} \quad \text{and} \quad \frac{d\varphi^\tau}{d\varphi^T(x)} \frac{d\varphi^T}{dx} v < \delta. \quad (79)$$

We now argue that such  $v$  converges to a subspace based on the adjoint sensitivity equation

$$\dot{u} = - \left[ \frac{df}{dx} \right]^\top u. \quad (80)$$

The adjoint sensitivity equation runs backwards in time, identifying the directional derivative of the inverse function. We will characterize the limiting nullspace of  $\frac{d\varphi^\tau}{dx}$  based on the directional derivatives of  $\varphi^{-\tau}$  along the normals, which we have shown becomes the left nullspace of the matrix.

Recall that we can expand the derivative as the derivative on the manifold and an  $\epsilon$  perturbation, i.e.

$$\dot{u} = \left( - \left[ \frac{df}{dx} \right] \Big|_{x=\pi_{\mathcal{M}}(x)}^\top + \Gamma(x) \right) u \quad (81)$$

where  $\|\Gamma(x)\| < \epsilon$ .

By observing that  $\left[ \frac{df}{dx} \right] \Big|_{x=\pi_{\mathcal{M}}(x)}^\top u \in T_{\pi_{\mathcal{M}}(x)}^\perp \mathcal{M}$ , we see that the dynamics converge such that  $u$  cannot rotate out of the normal space as  $\epsilon \rightarrow 0$ .

Thus, as  $T \rightarrow \infty$ , the normal space becomes approximately closed in the adjoint analysis, and the vectors which map onto the normal space become  $S(x_0)$ .  $\square$

**Theorem D.9** (Context: See Theorem 3.9). *Suppose the state space model in Equations (1) and (2) with noise satisfying the assumptions in Theorem 3.5 converges to a  $K$ -dimensional smooth manifold such that the Jacobian of the limiting flow is rank  $K$ . Then the linear functional which minimizes the CRLB for prediction in the infinite-horizon average cost formulation is the solution of*

$$\arg \max_{u: \|u\|_{(\Lambda_\sigma + \Sigma)} = 1} \sum_{i=1}^K \alpha_i \langle v_i, u \rangle_\Sigma^2 \quad (82)$$

for some  $\{\alpha_i\}$ , where  $v_i$  is the set of right singular vectors of the Jacobian of the limiting flow around the current state with nonzero singular vectors,  $\Sigma$  is the CRLB for estimating the current state based on the past measurements,  $\bar{u} = \frac{u}{\|u\|_\Sigma}$  is a unit-length version of  $u$ ,  $\langle v, u \rangle_\Sigma = v^\top \Sigma u$ , and  $\|\cdot\|_\Sigma$  is the associated induced seminorm.

*Proof.* Apply the singular value decomposition to the Jacobian of the flow as  $[d\varphi_{x_0}^\tau] = U_\tau D_\tau V_\tau^*$ . For sufficiently large  $\tau$ , by our low-rank assumption, there are at most  $K$  non-zero singular values. By Lemma 3.8, the row space  $V_\tau^*$  is invariant to  $\tau$ . Thus, we let  $V_\tau^* = R_\tau^* V^*$ , where  $R_\tau$  is some  $K \times K$  unitary transformation.

It follows that

$$\Sigma_\tau = U_\tau D_\tau R_\tau^* \underbrace{[V^* \Sigma_0 V]}_{K \times K} R_\tau D_\tau U_\tau^*. \quad (83)$$

Thus, by plugging the CRLB into our average cost objective

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N \text{Tr}(\Sigma_{\tau_j}) = \frac{1}{N} \sum_{j=1}^N \text{Tr}(U_\tau D_\tau R_\tau^* [V^* \Sigma_0 V] R_\tau D_\tau U_\tau^*) \quad (84)$$

$$= \frac{1}{N} \sum_{j=1}^N \text{Tr}(D_\tau R_\tau^* [V^* \Sigma_0 V] R_\tau D_\tau). \quad (85)$$

Using the Sherman–Morrison formula (Sherman and Morrison, 1950) for our Fisher Information update, i.e.

$$\left[ \Sigma_{-1} + \frac{1}{\|u_i\|_{\Lambda_\sigma}^2} u_i u_i^\top \right]^{-1} = \Sigma_{-1} - \frac{\Sigma_{-1} u_i u_i^\top \Sigma_{-1}}{\|u_i\|_{\Lambda_\sigma}^2 + u_i^\top \Sigma_{-1} u_i}, \quad (86)$$

the update to the future cost becomes

$$\frac{-1}{N \left( \|u_i\|_{\Lambda_\sigma}^2 + \|u_i\|_{\Sigma_{-1}}^2 \right)} \sum_{j=1}^N \text{Tr}(D_\tau R_\tau^* [V^* \Sigma_{-1} u_i u_i^\top \Sigma_{-1} V] R_\tau D_\tau). \quad (87)$$

Next, note that by expressing terms as

$$V^* \Sigma_{-1} u_i = \begin{bmatrix} \langle v_1, u_i \rangle_{\Sigma_{-1}} \\ \vdots \\ \langle v_K, u_i \rangle_{\Sigma_{-1}} \end{bmatrix} \quad R_\tau D_\tau D_\tau R_\tau^* = \begin{bmatrix} \alpha_{1,1}^{(\tau)} & \alpha_{1,2}^{(\tau)} & \cdots & \alpha_{1,K}^{(\tau)} \\ \alpha_{2,1}^{(\tau)} & \alpha_{2,2}^{(\tau)} & \cdots & \alpha_{2,K}^{(\tau)} \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_{K,1}^{(\tau)} & \alpha_{K,2}^{(\tau)} & \cdots & \alpha_{K,K}^{(\tau)} \end{bmatrix}, \quad (88)$$

we can rewrite the expression in the trace as

$$\text{Tr}(D_\tau R_\tau^* [V^* \Sigma_{-1} u_i u_i^\top \Sigma_{-1} V] R_\tau D_\tau) = \sum_{k=1}^K \alpha_{k,k}^{(\tau)} \langle v_k, u_i \rangle_{\Sigma_{-1}}^2. \quad (89)$$

Rewrite the summation in terms of our new expression for the trace and commute the summations as

$$\arg \max_{u_i} \sum_{k=1}^K \left( \frac{\langle v_k, u_i \rangle_{\Sigma_{-1}}^2}{\|u_i\|_{\Lambda_\sigma}^2 + \|u_i\|_{\Sigma_{-1}}^2} \left( \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N \alpha_{k,k}^{(\tau)} \right) \right) \quad (90)$$

Finally, let  $\alpha_k := \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N \alpha_{k,k}^{(\tau)}$  to express the objective as

$$\arg \max_{u_i} \frac{\sum_{k=1}^K \alpha_k \langle v_k, u_i \rangle_{\Sigma_{-1}}^2}{\|u_i\|_{\Lambda_\sigma}^2 + \|u_i\|_{\Sigma_{-1}}^2}. \quad (91)$$

Observe that  $\|u_i\|_{\Lambda_\sigma}^2 + \|u_i\|_{\Sigma_{-1}}^2 = \|u_i\|_{\Lambda_\sigma + \Sigma_{-1}}^2$  and apply bilinearity to move it inside the inner products as

$$\arg \max_{u_i} \sum_{k=1}^K \alpha_k \left\langle v_k, \frac{u_i}{\|u_i\|_{(\Lambda_\sigma + \Sigma_{-1})}} \right\rangle_{\Sigma_{-1}}^2. \quad (92)$$

Finally, replace the normalization with a constraint that  $\|u_i\|_{(\Lambda_\sigma + \Sigma_{-1})} = 1$ , noting that by the absolute homogeneity of norms, the result is invariant to scaling:

$$\arg \max_{u_i: \|u_i\|_{(\Lambda_\sigma + \Sigma_{-1})} = 1} \sum_{k=1}^K \alpha_k \langle v_k, u_i \rangle_{\Sigma_{-1}}^2. \quad (93)$$

□

**Corollary D.10** (Proof: See Corollary 3.10). *Under the assumptions of Theorem 3.9, the optimal measurement vector exists in the subspace*

$$\text{Span} \left\{ (\Sigma + \Lambda_\sigma)^{-1} \Sigma v_i \right\}_{i=1}^K. \quad (94)$$

*Proof.* This statement follows quickly from the application of Lagrange multipliers. Begin with the statement from Theorem 3.9,

$$\arg \max_{u: \|u\|_{(\Lambda_\sigma + \Sigma)} = 1} \sum_{i=1}^K \alpha_i \langle v_i, u \rangle_{\Sigma}^2. \quad (95)$$

Introduce the construct the Lagrangian

$$\mathcal{L}(u, \lambda) = \sum_{i=1}^K \alpha_i \langle v_i, u \rangle_{\Sigma}^2 + \lambda (1 - \|u\|_{(\Lambda_\sigma + \Sigma)}^2), \quad (96)$$

where the norm has been squared, leaving the constraint unchanged.

Expand the expressions

$$\mathcal{L}(u, \lambda) = \sum_{i=1}^K \alpha_i (v_i^\top \Sigma u)^2 + \lambda (1 - u^\top (\Lambda_\sigma + \Sigma) u). \quad (97)$$

Compute the derivative with respect to  $u$ :

$$\frac{d\mathcal{L}}{du} = \sum_{i=1}^K \alpha_i (v_i^\top \Sigma u) \Sigma v_i - 2\lambda (\Lambda_\sigma + \Sigma) u. \quad (98)$$

Set equal to zero and rearrange terms to find

$$u = \frac{1}{2\lambda} \sum_{i=1}^K \underbrace{\alpha_i (v_i^\top \Sigma u)}_{\in \mathbb{R}} (\Lambda_\sigma + \Sigma)^{-1} \Sigma v_i. \quad (99)$$

Finally, observe that  $\alpha_i (v_i^\top \Sigma u) \in \mathbb{R}$ , and thus

$$u \in \text{Span} \left\{ (\Sigma + \Lambda_\sigma)^{-1} \Sigma v_i \right\}_{i=1}^K. \quad (100)$$

□

**Corollary D.11** (Proof: See Corollary 3.11). *If, in addition to the assumptions of Theorem 3.9, the dynamical system converges to an isolated limit cycle, then*

$$u = (\Sigma + \Lambda_\sigma)^{-1} \Sigma v \quad (101)$$

*is an optimal measurement vector, where  $v$  is the right singular vector.*

*Proof.* By Corollary 3.6, the information content is invariant to the scaling of the linear functional. By Corollary 3.10, the optimal measurement functional is in the span of  $(\Sigma + \Lambda_\sigma)^{-1} \Sigma v$ . Thus, the vector generating the subspace is one possible optimal measurement. □

## E COMPUTATIONAL COSTS

In this section, we detail some of the specific aspects of the computation in this work. In particular, we detail the computational costs of the various operations, as well as provide more detail on the gradient ascent procedure for optimizing the 1D approximation. In Table 2 we include a summary of the computational cost scaling for the dynamic programming approach and the proposed gradient ascent based dimensionality reduction approach. Note that, as described in Section E.2, the evaluation cost of the proposed method can be reduced to  $O(CM^2)$ , where  $C$  is the number of gradient ascent steps, when the governing dynamics are particularly easy to solve.

|                     | INITIALIZATION                 | EVALUATION                  |
|---------------------|--------------------------------|-----------------------------|
| DYNAMIC PROGRAMMING | $O((MN^2 + NKP) \mathcal{A} )$ | $O((M^3 + K) \mathcal{A} )$ |
| PROPOSED METHOD     | $O(1)$                         | $O(M^2C + M^3)$             |

Table 2: Multiplication scaling for the two approaches outlined in this work, where  $M$  is the dimensionality of the system,  $N$  is the number of samples in the state space approximation,  $K$  is the average number of state space representation points in a  $d_{max}$  ball,  $P$  is the number of steps in value iteration,  $|\mathcal{A}|$  is the number of possible actions, and  $C$  is the number of gradient ascent steps.

From this small table, the main observation is that the dynamic programming approach requires a prohibitively expensive initialization, as  $N$  is typically significantly larger than any other variables in the expressions. Furthermore, we do not necessarily incur a significant cost in the proposed method for evaluating the policy. Differences in iteration cost essentially reduce to a comparison of the number of gradient ascent steps  $C$  to the product of the number of discrete actions and the dimensionality of the system  $M|\mathcal{A}|$ . As the surface area of the hypersphere grows exponentially with the dimensionality, the latter is often significantly larger.

### E.1 Sampling State Space

$O(NM^3)$  multiplications, where  $N$  is the total number of samples.

First, we must generate  $M^2 + M$  scalar random variables. Then, QR factorization requires  $O(M^3)$  multiplications (Trefethen and Bau, 1997). As  $N$  samples are required, the entire cost is  $O(NM^3)$  multiplications.

### E.2 Advancing State and CRLB in Time

Worst Case:  $O(M^3)$  multiplications; Best Case:  $O(M^2)$  multiplications

As evaluating  $\varphi^\tau$  and  $d\varphi^\tau$  are heavily dependent on the particular system, we treat these as some fixed constant cost, though it ranges from  $O(M)$  to  $O(M^3)$ .

The remainder of the update is then  $[d\varphi_x^\tau] (\Sigma^{-1} + \sigma^{-2}uu^\top)^{-1} [d\varphi_x^\tau]^\top$ , where the inverse is computed using the Sherman–Morrison Formula (Sherman and Morrison, 1950)

$$(\Sigma^{-1} + \sigma^{-2}uu^\top)^{-1} = \Sigma - \frac{\Sigma uu^\top \Sigma}{\sigma^2 + u^\top \Sigma u}. \quad (102)$$

The number of multiplications for equation (102) can be limited through the following order of computations.

1. Compute  $\Sigma u$  for  $M^2$  multiplications
2. Compute  $u^\top (\Sigma u)$  for  $M$  multiplications
3. Compute  $(\Sigma u)/(\sigma^2 + u^\top \Sigma u)$  for  $M$  multiplications
4. Compute  $[(\Sigma u)/(\sigma^2 + u^\top \Sigma u)] [\Sigma u]^\top$  for  $M^2$  multiplications.

Thus, the update equation requires  $2(M^2 + M)$  multiplications.

### E.3 Computation of Local Average Weights

$O(MN^2|\mathcal{A}|)$  multiplications, where  $|\mathcal{A}|$  is the cardinality of the finite action space.

This involves computing the distance between all potential states after the observation at a given sample to all samples. Thus, as  $\|x - x'\|$  requires  $M + 1$  multiplications, the computation of the entire set of distances requires  $MN^2|\mathcal{A}|$  multiplications, where  $|\mathcal{A}|$  is the cardinality of the finite action space.

It is worth noting that this represents the primary cost of the dynamic programming approach, as the required value of  $N^2$  grows rapidly with the dimensionality of the system.

#### E.4 Value Iteration

$O(|\mathcal{A}|NK^2P)$  multiplications, where  $K$  is the average number number of points in a  $d_{max}$  ball, and  $P$  is the number of steps in value iteration.

In this section, assume  $d_{max}$  is appropriately scaled such that states have an average number of  $K \ll N$  neighbors. Then, each action for each state sample requires on average  $K$  multiplications, resulting in a total of  $|\mathcal{A}|NK$  multiplications.

#### E.5 Dynamic Programming Optimal Action Evaluation

$O((M^3 + K)|\mathcal{A}|)$  multiplications

For each possible action, we must advance the state and CRLB forward in time by one timestep, requiring  $O(M^3)$  multiplications. We then must compute the local average of the sample value function for  $K$  multiplications.

#### E.6 Extended Kalman Filter

$O(M^2)$  multiplications

$$\hat{x}_{i|i} = \hat{x}_{i|i-1} + \frac{\sum_i u_i (y_i - u_i^\top \hat{x}_{i|i-1})}{u_i^\top \sum_i u_i + \sigma^2}, \quad (103)$$

1. Compute  $\sum u$  for  $M^2$  multiplications.
2. Compute  $u^\top (\sum u)$  for  $M$  multiplications.
3. Compute  $u^\top \hat{x}_{i|i-1}$  for  $M$  multiplications
4. Combine multiplicative constants for 1 multiplication.
5. Remaining update for  $M$  multiplications.

Adding the above costs shows that the EKF update step requires  $M^2 + 3M + 1$  multiplications.

#### E.7 Gradient Ascent Optimization

$O(M^2)$  multiplications per gradient step

We now address the details of the iterative optimization of

$$\arg \max_{u: \|u\|=1} \frac{(v^\top \sum u)^2}{u^\top \sum u + \sigma^2 u^\top u}. \quad (104)$$

To compute the gradient of the function on the sphere, we first need to compute the gradient in the ambient Euclidean space, then project onto the tangent space.

By the application of standard differentiation rules

$$d \left( \frac{(v^\top \sum u)^2}{u^\top \sum u + \sigma^2 u^\top u} \right)_u = 2 \frac{(u^\top \sum u + \sigma^2 u^\top u) v^\top \sum u v^\top \sum - (v^\top \sum u)^2 (u^\top \sum + u^\top)}{(u^\top \sum u + \sigma^2 u^\top u)^2} \quad (105)$$

$$= 2 \frac{(u^\top \sum u + \sigma^2) v^\top \sum u v^\top \sum - (v^\top \sum u)^2 (u^\top \sum + u^\top)}{(u^\top \sum u + \sigma^2)^2}, \quad (106)$$

where the second line comes from noting that  $\|u\| = 1$ . Finally, we can slightly reduce the computation further by noting that we eventually orthogonalize the gradient to  $u$ . Thus, the following is an equivalent computation

$$2 \frac{(u^\top \sum u + \sigma^2) v^\top \sum u v^\top \sum - (v^\top \sum u)^2 u^\top \sum}{(u^\top \sum u + \sigma^2)^2}, \quad (107)$$



By performing operations in the following order, the differential of the value in Euclidean space requires few operations. Below, we count the number of multiplications.

1. Compute  $\underline{\Sigma}u$  for  $M^2$  multiplications.
2. Compute  $u^\top (\underline{\Sigma}u)$  for  $M$  multiplications.
3. Compute  $v^\top (\underline{\Sigma}u)$  for  $M$  multiplications.
4. Compute  $v^\top \underline{\Sigma}$  for  $M^2$  multiplications.
5. Computing remaining terms in numerator for  $2M + 2$  multiplications.
6. Divide by numerator for  $M + 1$  multiplications.

Finally, we must project the result onto the tangent space of the sphere. Fortunately, the tangent space is straightforward for a sphere, in particular, it is the space orthogonal to  $u$ . Thus, we project the vector  $u$  out with  $2M$  additional multiplications.

By adding all of the above, we find that computation of the gradient requires  $2M^2 + 7M + 3$  multiplications.

Once the gradient has been computed, it defines the discrete-time approximation of the gradient flow by following the geodesic curve

$$u_{i+1} = \cos(\alpha_i \|s_i\|)u_i + \sin(\alpha_i \|s_i\|)s_i / \|s_i\|, \quad (108)$$

where  $s_i$  is the gradient in the ambient space. The above computation requires another  $3M$  multiplications.

Thus, each gradient ascent step requires  $2M^2 + 10M$  plus some constant number of multiplications from evaluating sine and cosine.

### E.7.1 Gradient Ascent Simulations

To give a sense of the non-linear optimization problem, we include visualizations of two potential optimization surfaces for 3D systems, shown as two different panels in Figure 7. Because the domain of the function is the unit sphere, we include a pair of heatmaps in each panel representing the front and back of the sphere by a z-coordinate threshold to show the occluded side.

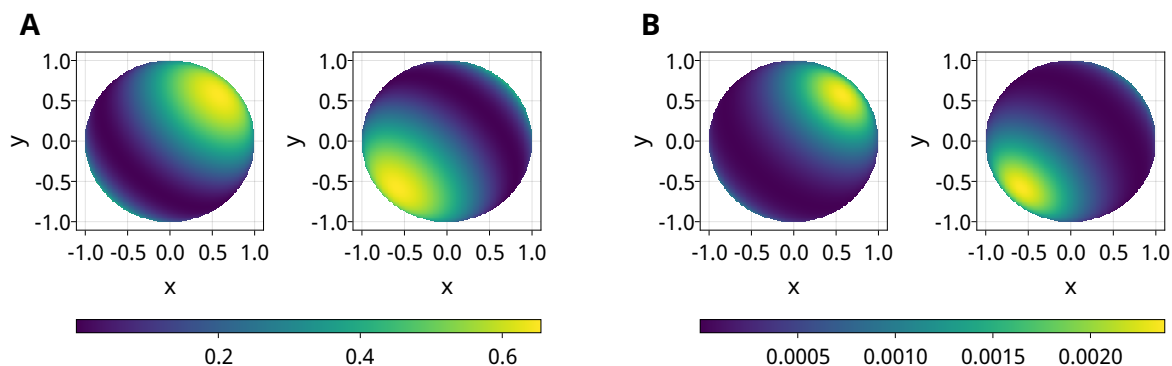


Figure 7: Optimization surface for gradient ascent for two different estimation CRLBs. In Panel B, the estimation CRLB is chosen to be poorly aligned with the limit vector.

There are two useful observations to be made. First, as was clear from both intuition and the form of the optimization problem, the function is symmetrical. Next, in these particular simulations, the function is relatively smooth with a “unique” maximum, only mirrored on the opposite side. This second observation allows the gradient ascent to be rather straightforward and gives us some certainty that, given appropriate step sizes, we will identify the true optimum. In these

plots,  $v = [3^{-1/2}, 3^{-1/2}, 3^{-1/2}]$  and  $\sigma^2 = 1$ . The estimation CRLB took the form:

$$\text{Panel A: } \Sigma = \begin{bmatrix} 1.0 & 0.1 & 0.1 \\ 0.1 & 1.0 & 0.1 \\ 0.1 & 0.1 & 1.0 \end{bmatrix} \quad \text{Panel B: } \Sigma = \begin{bmatrix} 1.0 & 0.0 & 0.0 \\ 0.0 & 1.0 & 0.0 \\ 0.0 & 0.0 & 1.0 \end{bmatrix} - 0.95vv^\top. \quad (109)$$

Figure 8 provides a representative trajectory of the gradient ascent operation over the previously illustrated objective.




Figure 8: Example gradient ascent trajectory on our objective function. Panel A: Objective function on sphere with black line corresponding to the ascent trajectory; Panel B: Objective function value as a function of iteration.

In Panel A of Figure 8, we plot the trajectory of the gradient ascent on top of the objective function. The initialization was chosen to emphasize that the trajectory approximately follows a geodesic: when the projection of the sphere onto a plane is centered on the initial position, geodesic curves appear straight, similar to the behavior in this plot. Panel B shows the evaluation of the objective function over time, illustrating the rapid convergence in this simulation. In this simulation, a constant step size was chosen of  $\alpha = 0.1$ .

### E.8 Timing Simulations

While the implementations in this work are not optimized, we include some timing plots below to give a sense of the rate at which the computational costs grow for the different approaches and provide a general order of magnitude for the performance. In our timing plots, a Van der Pol system is being observed. In order to best elucidate the scaling aspects, in the proposed method, the computation time is plotted as a function of dimensionality, while the dynamic programming computation time is plotted as a function of the number of state space representation points. The timing plots are available in Figure 9.




Figure 9: A: Timing of proposed method as a function of dimensions. B: Timing of the dynamic programming baseline as a function of the number of samples. Note that the required samples may grow exponentially with the dimensionality of the space, resulting in very fast growth of the computational costs.

To generate the plots, the time required to initialize the algorithms as well as compute 1000 decision steps was recorded. Timing was done using the default number of runs in the BenchmarkTools library for Julia, which returns the minimum time of multiple runs to eliminate issues relating to just-in-time compilation and unrelated computer processes interfering with a run.

Timesteps were spaced by 0.05 units of time in both simulations. In the proposed method, the 1D approximation was made at  $T = 10$  units of time in the future, and 100 gradient ascent steps were computed. In the dynamic programming approach, the cardinality of the action space was 31,  $d_{max} = 1$ , and 100 steps of value iteration were completed in the initialization.

In the simulation of the proposed method shown in Panel A of Figure 9, we observe nearly linear growth in the computation time as a function of the dimensionality of the system with a sudden change around 40 dimensional systems potentially resulting in a change in caching at that point. This suggests that for the chosen parameters, the cost is dominated by system simulation costs which grow linearly for the chosen system as a function of dimension. In Panel B, the growth is clearly faster than linear.

It is additionally worth noting that, for reasonable performance, Panel B represents only a small subset of the domain of Panel A. As previously shown in Figure 6, the required number of samples rapidly increases with the dimensionality of the system, resulting in the dynamic programming approach becoming intractable while the proposed method still requires less than 1 second of computation.

## E.9 System Specifications

The code in this project was run on an AMD Ryzen Threadripper 3960X with 64 GB of RAM and 8GB swap running Ubuntu 20.04.2 LTS. All code was written in Julia 1.7.1, using the libraries

1. DifferentialEquations v7.0.0, MIT "Expat" License, (Rackauckas and Nie, 2017)
2. DiffEqSensitivity v6.69.0, MIT "Expat" License, (Ma et al., 2021)
3. EllipsisNotation v1.1.3, MIT "Expat" License
4. ForwardDiff v0.10.25, MIT License, (Revels et al., 2016)
5. SafeTestsets v0.0.1, MIT "Expat" License
6. Documenter v0.27.12, MIT License
7. CSV v0.10.2, MIT "Expat" License
8. BenchmarkTools v1.2.2, MIT "Expat" License, (Chen and Revels, 2016)
9. CairoMakie v0.6.4, MIT "Expat" License, (Danisch and Krumbiegel, 2021)

## F NUMERICAL ISSUES

While the approximately singular behavior in this work benefit the statistical properties, they can introduce numerical issues in the implementation. While the algorithm introduced in the main paper functions in exact arithmetic, rounding errors in floating point arithmetic result in matrices becoming singular. In this section, we discuss the required modifications to account for these numerical issues relating to inverting such matrices.

The main issue is from the CRLB update formula

$$\Sigma_{i+1} = [d\varphi_x^\tau] (\Sigma_i^{-1} + \sigma^{-2}uu^\top)^{-1} [d\varphi_x^\tau]^\top. \quad (110)$$

Recalling that  $\Sigma$  and  $\Sigma^{-1}$  both begin as positive definite matrices, consider the operations that could make the inverse impossible. First,  $\Sigma_i^{-1} + \sigma^{-2}uu^\top$  may only become singular if  $\sigma^{-2}uu^\top$  is significantly large compared to an eigenvalue of the Fisher information. Fortunately, due to the dimensionality collapse and the continual collection of data, this issue does not appear.

The other, more likely, issue is that the application of the Jacobian creates a singular matrix. In practice, this happens quite often and thus must be addressed by replacement with an equivalent computation.

To resolve these issues, we complete the updates through the Sherman–Morrison formula. We include here a brief note on the behavior of the Sherman–Morrison formula when a matrix is singular.

$$(\Sigma^{-1} + \sigma^{-2}uu^T)^{-1} = \Sigma - \frac{\Sigma uu^T \Sigma}{\sigma^2 + u^T \Sigma u}. \quad (111)$$

We argue that, when  $\Sigma$  is singular, this update formula is equivalent to updating the subspace containing non-zero eigenvalues. To see this, we simply replace  $\Sigma^{-1}$  with the Moore–Penrose pseudoinverse before multiplying the terms to verify the property. Thus, replacement with the pseudo-inverse produces a new pseudo-inverse in the formula. By recalling that the pseudoinverse can be computed by inverting all non-zero singular values, we can conclude that the formula operates correctly in the subspace.

## G CODE USAGE

The library produced for this manuscript is publicly available under the MIT license at [github.com/Helmuthn/naumer\\_Dimensionality\\_2022.jl](https://github.com/Helmuthn/naumer_Dimensionality_2022.jl).

The implementation is a Julia package, and while not registered, can be installed through the package manager in the REPL through `] add https://github.com/Helmuthn/naumer_Dimensionality_2022.jl`. We include here a short summary of the interface which much be defined to apply the library to a new problem.

The functions in the library make use of the type `AbstractSystem`, which is an abstract type with the following three functions:

`dimension`

This function accepts an instance of `AbstractSystem` and returns the ambient dimensionality.

`flow`

This function accepts an instance of `abstractSystem`, a state vector  $x$ , and a duration of time  $\tau$  and returns the state advanced by  $\tau$ .

`flowJacobian`

This function accepts an instance of `abstractSystem`, a state vector  $x$ , and a duration of time  $\tau$  and returns the Jacobian of the flow advancing by  $\tau$  around initial condition  $x$ .

While trivial, we now include the complete implementation of `LinearSystem`, which represents a linear system of ordinary differential equations, below as an example of an implementation of the interface.

```
"""
    LinearSystem{T} <: AbstractSystem{T}
Defines a linear system dx/dt = Ax
### Fields
- `dynamics::Matrix{T}` - Matrix defining dynamics
"""
struct LinearSystem{T} <: AbstractSystem{T}
    dynamics::Matrix{T}
end

function flow(x::Vector, τ, system::LinearSystem)
    return exp(τ * system.dynamics) * x
end

function flowJacobian(x::Vector, τ, system::LinearSystem)
    return exp(τ * system.dynamics)
end

function dimension(system::LinearSystem)
```

```
    return size(system.dynamics)[1]
end
```

In typical applications, there will be no closed form solution for the dynamics. When this is the case, one can use standard libraries such as `DifferentialEquations.jl` and `ForwardDiff.jl` to define the `AbstractSystem`. An example is included below.

```
using DifferentialEquations, ForwardDiff

struct VanDerPolSystem{T} <: AbstractSystem{T}
    μ::T
end

function dimension(system::VanDerPolSystem)
    return 2
end

function vanderpolDynamics!(du, u, p, t)
    x, y = u
    μ = p
    du[1] = y
    du[2] = μ * (1 - x^2) * y - x
end

function flow(x::Vector, τ, system::VanDerPolSystem)
    problem = ODEProblem(vanderpolDynamics!, x, (0.0, τ), system.μ)
    sol = solve(problem, Tsit5(), reltol=1e-8)

    return sol[end]
end

function flowJacobian(x::Vector, τ, system::VanDerPolSystem)
    problem = ODEProblem(hopfDynamics!, x, (0.0, τ), system.μ)

    function solvesystem(init)
        prob = remake(problem, u0=init)
        sol = solve(prob, Tsit5(), reltol=1e-8)
        return sol[end]
    end

    return ForwardDiff.jacobian(solvesystem, x)
end
```

Discrete-time systems can be defined by restricting  $\tau$  in the function definitions, i.e.

```
flow(x::Vector, τ::Integer, system::AbstractSystem)
```

```
struct DiscreteSystem{T} <: AbstractSystem{T}
    dynamics::Matrix{T}
end

function flow(x::Vector, τ::Integer, system::DiscreteSystem)
    out = copy(x)
    for i in 1:τ
        out .= system.dynamics * out
    end
    return out
end

function flowJacobian(x::Vector, τ::Integer, system::DiscreteSystem)
    return system.dynamics^τ
end
```

```

end

function dimension(system::DiscreteSystem)
    return size(system.dynamics) [1]
end

```

## H SIMULATION DETAILS

In this section, we provide all parameters used in the generation of the figures in the main work. All code used to generate the figures is additionally available in the previously mentioned repository.

### H.1 Figure 1

Simulations regarding the Van der Pol oscillator with parameter  $\mu = 1$  were completed for this figure.

In Panel A, the direction of the flow was plotted on a grid with samples spaced by approximately  $1/3$ .

In Panel B, four trajectories initialized at

$$x_0 = \begin{bmatrix} 0.0 \\ -4.0 \end{bmatrix} \quad x_0 = \begin{bmatrix} 0.0 \\ -0.1 \end{bmatrix} \quad x_0 = \begin{bmatrix} 0.0 \\ 0.1 \end{bmatrix} \quad x_0 = \begin{bmatrix} 0.0 \\ 4.0 \end{bmatrix}, \quad (112)$$

and were simulated for 20 seconds of system time.

In Panel C, systems initialized at each of the grid points from Panel A were simulated for 20 seconds of system time with relative and absolute tolerance of  $10^{-6}$ . We then numerically differentiated the solution using the ForwardDiff library, a common forward mode automatic differentiation library in Julia with good support for differential equations sensitivity analysis. The singular value decomposition of the resulting Jacobian was computed, and the right singular vector associated with the largest singular value was plotted.

In Panel D, to compute a representation of the limit cycle, the system was initialized at

$$x_0 = [1.0 \quad 1.0]^\top \quad (113)$$

and was then simulated for 20 seconds of system time. We then computed the squared distances between the final point of the trajectory and the previous 100 positions spaced by 0.01 seconds of system time. As the period of the limit cycle is between 1 and 2 seconds of system time, we then computed the nearest position in the list. The trajectory from that index until the end was used to represent the phase of the limit cycle.

At each of the grid points from Panel A, the system trajectory was simulated for 40 seconds of system time. It was then assigned the index of the closest point in the phase representation of the limit cycle as its value and plotted as a heatmap with contours.

### H.2 Figure 3

#### H.2.1 Panel A

Random sampling was completed through MersenneTwister in the Random library of Julia with the seed 1234, chosen arbitrarily.

The system took a constant value of  $x = [1 \quad 1]^\top$ .

As the system trajectory value does not impact the CRLB updates in this system, 1000 PSD matrix samples were used with eigenvalues following exponential distributions with parameter  $\lambda = 1$ , while only one state space sample was chosen. The maximum distance for averaging was chosen as  $d_{max} = 1$ .

Observations of the system were made under i.i.d. Gaussian noise with variance  $\sigma^2 = 1$ . The action space for both the dynamic programming baseline approach and the proposed method was the set of vectors along the unit circle spaced by 0.1 radians between 0 and  $\pi$ .

The CRLB was initialized as

$$\Sigma = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}, \quad (114)$$

while the state was randomly initialized by sampling two i.i.d. Gaussian random variables with variance 4.

Value iteration was run for 1000 iterations with a discount factor of  $\gamma = 0.99$  to ensure a focus on distant time horizons.

The proposed method was not run for this system, as it relies inherently on uneven singular values.

### H.2.2 Panel B

Random sampling was completed through MersenneTwister in the Random library of Julia with the seed 1234, chosen arbitrarily.

The system followed the linear dynamics

$$\dot{x} = \begin{bmatrix} -10.0 & 0 \\ 0 & -0.1 \end{bmatrix} x, \quad (115)$$

discretized into timesteps of duration 0.01.

As the system trajectory value does not impact the CRLB updates in this system, 100 PSD matrix samples were used with eigenvalues following exponential distributions with parameter  $\lambda = 1$ , while only one state space sample was chosen. The maximum distance for averaging was chosen as  $d_{max} = 2$ .

Observations of the system were made under i.i.d. Gaussian noise with variance  $\sigma^2 = 1$ . The action space for both the dynamic programming baseline approach and the proposed method was the set of vectors along the unit circle spaced by 0.05 radians between 0 and  $\pi$ .

The CRLB was initialized as

$$\Sigma = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}, \quad (116)$$

while the state was randomly initialized by sampling two i.i.d. Gaussian random variables with variance 4.

Value iteration was run for 10000 iterations with a discount factor of  $\gamma = 0.95$  to ensure a focus on distant time horizons.

For the proposed technique, the limiting singular vector was approximated at 10 seconds of system time in the future.

### H.2.3 Panel C

Random sampling was completed through MersenneTwister in the Random library of Julia with the seed 1234, chosen arbitrarily.

The system followed a Hopf bifurcation dynamic

$$\begin{aligned} \dot{x}_{(1)} &= x_{(1)}(1 - \|x\|^2) - x_{(2)} \\ \dot{x}_{(2)} &= x_{(2)}(1 - \|x\|^2) + x_{(1)} \end{aligned} \quad (117)$$

discretized into timesteps of duration 0.01.

The dynamic programming space was approximated using the product space of 30 PSD matrix samples with eigenvalues following exponential distributions with parameter  $\lambda = 1$  and 120 state space samples. The state space samples were chosen by selecting 30 i.i.d. Gaussian vectors with each element independent with variance 25, then advancing each of the points by 4 timesteps, retaining each sample in the trajectory. The maximum distance for averaging was chosen as  $d_{max} = 2$ .

Observations of the system were made under i.i.d. Gaussian noise with variance  $\sigma^2 = 1$ . The action space for both the dynamic programming baseline approach and the proposed method was the set of vectors along the unit circle spaced by 0.1 radians between 0 and  $\pi$ .

The CRLB was initialized as

$$\Sigma = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}, \quad (118)$$

while the state was randomly initialized by sampling two i.i.d. Gaussian random variables with variance 4.

Value iteration was run for 100 iterations with a discount factor of  $\gamma = 0.99$  to ensure a focus on distant time horizons.

For the proposed technique, the limiting singular vector was approximated at 30 seconds of system time in the future.

## H.2.4 Panel D

Random sampling was completed through MersenneTwister in the Random library of Julia with the seed 1234, chosen arbitrarily.

The system followed a Hopf bifurcation dynamic

$$\begin{aligned} \dot{x}_{(1)} &= 10(x_{(2)} - x_{(1)}) \\ \dot{x}_{(2)} &= x_{(1)}(28 - x_{(3)}) - x_{(2)} \\ \dot{x}_{(3)} &= x_{(1)}x_{(2)} - \frac{8}{3}x_{(3)} \end{aligned} \quad (119)$$

discretized into timesteps of duration 0.01.

The dynamic programming space was approximated using the product space of 30 PSD matrix samples with eigenvalues following exponential distributions with parameter  $\lambda = 1$  and 120 state space samples. The state space samples were chosen by selecting 100 i.i.d. Gaussian vectors with each element independent with variance 25. The maximum distance for averaging was chosen as  $d_{max} = 2$ .

Observations of the system were made under i.i.d. Gaussian noise with variance  $\sigma^2 = 1$ . The action space for both the dynamic programming baseline approach and the proposed method was a set of 50 i.i.d. uniformly chosen unit vectors.

The CRLB was initialized as

$$\Sigma = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (120)$$

while the state was randomly initialized by sampling two i.i.d. Gaussian random variables with variance 4.

Value iteration was run for 100 iterations with a discount factor of  $\gamma = 0.9$ , intentionally lower due to the chaotic behavior of the system.

For the proposed technique, the limiting singular vector was approximated at 1 second of system time in the future, lower than the others to account for the lack of actual convergence. The time horizon was chosen based upon the time horizon being plotted.

## H.3 Figure 4

### H.3.1 Panels A and B

Simulations were identical to those in Figure 3 with the additional step of the extended Kalman filter. Observation decisions were made according to the estimated state.

The cost was normalized to the oracle solution for the dynamic programming approach in order to emphasize the difference between the performance of the four techniques.



### H.3.2 Panel C

Random sampling was completed through MersenneTwister in the Random library of Julia with the seed 1234, chosen arbitrarily.

The simulated system followed a Van Der Pol Oscillator with  $\mu = 1$  as the chosen parameter augmented with additional linear dimensions governed by  $\dot{x}_{(i)} = -x_{(i)}$ .

The CRLB was evaluated after 1000 observations, spaced by 0.05 units of system time. Observations were made under i.i.d. Gaussian noise with variance  $\sigma^2 = 1$ .

1000 steps of gradient ascent were completed for each observation with decaying step sizes  $\alpha_i = 10^4 i^{-2/3}$ , which empirically helped with ensuring quick convergence for all objective functions.

### H.4 Figure 6

Random sampling was completed through MersenneTwister in the Random library of Julia with the seed 4321, chosen arbitrarily.

For each system dimensionality  $M$ , 10000 i.i.d.  $M \times M$  positive definite matrices were generated with eigenvalues sampled from an exponential distribution with parameter  $\lambda = 1$ .

Then, for each given number of samples  $N$ , we choose  $N + 1$  points uniformly without replacement from the original dataset and compute the minimum distance from one of those selected points chosen randomly from the subset. The process is averaged 5000 times.