Predictability Enables Parallelization of Nonlinear State Space Models

Xavier Gonzalez*
Stanford University
xavier18@stanford.edu

Leo Kozachkov*†
IBM Research
leokoz8@brown.edu

David M. Zoltowski Stanford University dzoltow@stanford.edu

Kenneth L. Clarkson IBM Research klclarks@us.ibm.com Scott Linderman Stanford University scott.linderman@stanford.edu

Abstract

The rise of parallel computing hardware has made it increasingly important to understand which nonlinear state space models can be efficiently parallelized. Recent advances like DEER [1] or DeepPCR [2] have shown that evaluating a state space model can be recast as solving a parallelizable optimization problem, and sometimes this approach can yield dramatic speed-ups in evaluation time. However, the factors that govern the difficulty of these optimization problems remain unclear, limiting the larger adoption of the technique. In this work, we establish a precise relationship between the dynamics of a nonlinear system and the conditioning of its corresponding optimization formulation. We show that the predictability of a system, defined as the degree to which small perturbations in state influence future behavior, impacts the number of optimization steps required for evaluation. In predictable systems, the state trajectory can be computed in $\mathcal{O}((\log T)^2)$ time, where T is the sequence length, a major improvement over the conventional sequential approach. In contrast, chaotic or unpredictable systems exhibit poor conditioning, with the consequence that parallel evaluation converges too slowly to be useful. Importantly, our theoretical analysis demonstrates that for predictable systems, the optimization problem is always well-conditioned, whereas for unpredictable systems, the conditioning degrades exponentially as a function of the sequence length. We validate our claims through extensive experiments, providing practical guidance on when nonlinear dynamical systems can be efficiently parallelized, and highlighting predictability as a key design principle for parallelizable models.

1 Introduction

Parallelization has been central to breakthroughs in deep learning, with GPUs enabling the fast training of large neural networks. In contrast, nonlinear state space models like recurrent neural networks (RNNs) have resisted efficient parallelization on GPUs due to their sequential nature.

Recent work addresses this mismatch by reformulating sequential dynamics into parallelizable optimization problems. Notably, the DEER/DeepPCR algorithm [1, 2] evaluates nonlinear state space dynamics by minimizing a residual-based merit function, facilitating efficient parallel computation

^{*}Equal contribution.

[†]Now at Brown University.

via the Gauss-Newton method.³ Gonzalez et al. [3] further developed these methods, including quasi-Newton methods and trust regions methods for parallel evaluation of nonlinear dynamical systems. These methods evaluate nonlinear dynamical systems by iteratively linearizing the nonlinear system and evaluating the resulting linear dynamical system (LDS) with a parallel (a.k.a. associative) scan [4, 5]. Each parallel evaluation of an LDS implements one optimization step [1–3, 6].

The usefulness of this optimization-based reformulation depends on two key factors: (a) the computational time per optimization step, and (b) the number of optimization steps required. The computational time per optimization step is only logarithmic in the sequence length, thanks to its parallel structure. How-

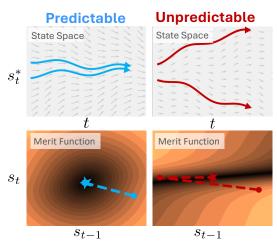


Figure 1: Predictable nonlinear state space models can be recast as well-conditioned, parallelizable optimization problems.

ever, the number of steps is governed by the conditioning of the merit function, and that remains poorly understood. In this paper, we characterize the merit function's conditioning, allowing us to draw a sharp distinction between systems that are amenable to efficient parallelization via merit function minimization and those that are not (see Figure 1, which is generated from trajectories of an RNN). Geometrically, we show that unpredictable systems lead to merit functions that have regions of extreme flatness, which can lead to very slow convergence.

Drawing from nonlinear dynamical systems theory—particularly contraction analysis [7] and Lyapunov exponent methods [8]—we formalize the relationship between system predictability and the conditioning of the merit function. **Unpredictable systems** are dynamical systems whose future behavior is highly sensitive to small perturbations. A common example is a chaotic system, like the weather: a butterfly flapping its wings in Tokyo today can lead to a thunderstorm in Manhattan next month [9, 10]. By contrast, **predictable systems** are those in which small perturbations are "forgotten." A familiar example is aviation: a patch of choppy air rarely makes an airplane land at the wrong airport. A more formal definition of (un)predictability is given in Definition 1. Our results establish key theoretical principles, make connections between optimization theory and dynamical systems, and demonstrate the practical applicability of parallel computations across a wide range of nonlinear state space modeling tasks.

Contributions & Outline Our central finding is that predictable systems give rise to well-conditioned merit functions, making them amenable to efficient parallelization. Unpredictable (e.g., chaotic) systems produce poorly conditioned merit functions and are not easily parallelizable.

The paper is organized as follows. Section 2 provides background, with formal definitions of predictable and unpredictable nonlinear state space models. Section 3 presents two key theoretical results that characterize the conditioning of the merit function, showing that the Polyak–Łojasiewicz (PL) constant μ of the merit function is controlled by the predictability of the dynamics (Theorem 2), and that the Lipschitz constant of the residual function Jacobian is governed by the nonlinearity of the dynamics (Theorem 3). Section 4 then uses the results about the conditioning of the merit function to prove results about Gauss-Newton in particular. We prove global linear rates of convergence for Gauss-Newton, with the precise rate scaling with the unpredictability of the problem (Theorem 4), and we characterize the basin of quadratic convergence in terms of the predictability and nonlinearity of the underlying dynamics (Theorem 5). In Section 5 we illustrate our results with experiments, and in Section 6 we conclude by summarizing context, implications, limitations, and future directions.

³DEER [1] and DeepPCR [2] were concurrent works that both proposed to use the Gauss-Newton method for optimizing nonlinear sum of squares to parallelize sequential processes. In this paper, we therefore use DEER, DeepPCR, and Gauss-Newton interchangeably.

2 Problem Statement & Background

Notation Throughout the paper, we use T to denote the length of a sequence and D to represent the dimensionality of a nonlinear state space model. Elements in \mathbb{R} , \mathbb{R}^D or $\mathbb{R}^{D \times D}$ are written using non-bold symbols, while elements in \mathbb{R}^{TD} or $\mathbb{R}^{TD \times TD}$ are denoted with bold symbols.

Sequential Evaluation vs. Merit Function Optimization We consider the *D*-dimensional non-linear state space model

$$s_t = f_t(s_{t-1}) \in \mathbb{R}^D. \tag{1}$$

A simple example is an input-driven nonlinear RNN, $s_t = \tanh(W s_{t-1} + B u_t)$, where W and B are weight matrices and u_t is the input into the network at time t. We want to compute the state trajectory, (s_1, \ldots, s_T) , starting from an initial condition s_0 , for a given sequence of functions f_1, \cdots, f_T .

Systems of the form (1) are widespread across science and engineering. Examples include physics (numerical weather prediction, molecular dynamics), biology (gene regulatory networks, population dynamics), engineering (control, robotics), and economics (macroeconomic forecasting, asset pricing). In machine learning, sequential operations arise in recurrent neural networks, iterative optimization, and the sampling pass of a diffusion model [2, 11]. Sequential operations even appear in the problem of evaluating transformer blocks over depth [12–17]. In probabilistic modeling, sequential operations arise in Markov Chain Monte Carlo [18]. In all of these cases, the state evolves through nonlinear transformations that capture the system's underlying dynamics.

The obvious approach is to sequentially compute the states according to eq. (1), taking T steps. Alternatively, one can cast state evaluation as an optimization problem. While less intuitive, an advantage of this approach is that it admits parallel computation [1-3]. Depending on the properties of the nonlinear state space model, the optimization algorithm, and the available hardware, the latter approach can be significantly faster than sequential evaluation.

We define the residual and corresponding merit⁴ function \mathcal{L} by stacking the elements $s_t \in \mathbb{R}^D$ of a trajectory into a TD-dimensional vector \mathbf{s} and considering the vector of temporal differences,

$$\mathbf{r}(\mathbf{s}) := \operatorname{vec}\left(\left[s_1 - f_1(s_0), \ldots, s_T - f_T(s_{T-1})\right]\right) \in \mathbb{R}^{TD}, \qquad \mathcal{L}(\mathbf{s}) := \frac{1}{2} \left\|\mathbf{r}(\mathbf{s})\right\|_2^2, \quad (2)$$

where $\text{vec}(\cdot)$ denotes the flattening of a sequence of vectors into a single column vector. The true trajectory \mathbf{s}^* is then obtained by minimizing $\mathcal{L}(\mathbf{s})$. Note that the residual is zero only at the true trajectory, i.e., when s_1, s_2, \dots, s_T satisfy (1) at every time point, so \mathbf{s}^* is the unique global minimum of $\mathcal{L}(\mathbf{s})$.

DeepPCR [2] and DEER [1] minimize the merit function using Gauss–Newton updates. Each update takes the form

$$\mathbf{s}^{(i+1)} = \mathbf{s}^{(i)} - \mathbf{J}(\mathbf{s}^{(i)})^{-1} \mathbf{r}(\mathbf{s}^{(i)}). \tag{3}$$

where $\mathbf{J}(\mathbf{s}^{(i)})$ denotes the Jacobian of the residual function, evaluated at the current iterate $\mathbf{s}^{(i)}$. The Jacobian is a $TD \times TD$ matrix with $D \times D$ block bidiagonal structure

$$\mathbf{J}(\mathbf{s}^{(i)}) := \frac{\partial \mathbf{r}}{\partial \mathbf{s}}(\mathbf{s}^{(i)}) = \begin{pmatrix} I_D & 0 & \dots & 0 & 0 \\ -J_2^{(i)} & I_D & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & I_D & 0 \\ 0 & 0 & \dots & -J_T^{(i)} & I_D \end{pmatrix} \quad \text{where} \quad J_t^{(i)} := \frac{\partial f_t}{\partial s_{t-1}}(s_{t-1}^{(i)}). \quad (4)$$

Due to this block bidiagonal structure, solving $\mathbf{J}(\mathbf{s}^{(i)})^{-1}\mathbf{r}(\mathbf{s}^{(i)})$ amounts to solving a linear recursion, which can be done in $\mathcal{O}(\log T)$ time with a parallel scan [5, 20, 21, 1, 3]. Further details are given in Appendix A.

This sublinear time complexity per step is only useful if the number of optimization steps required to minimize the merit function is small, otherwise it would be more efficient to evaluate the recursion sequentially. Thus, we seek to characterize the conditioning of the merit function — determining

⁴While minimizing a "merit function" is admittedly counterintuitive, we follow Nocedal and Wright [19, see eq. 11.35] in this convention.

when it is well-conditioned and when it is not — since this affects the difficulty of finding its minimum. Equation (4) already offers an important clue. The presence of the nonlinear state-space model Jacobians J_t , which measure the local stability and predictability of the nonlinear dynamics, foreshadows our central finding: the system's predictability dictates the conditioning of the merit function.

Predictable Systems: Lyapunov Exponents and Contraction Predictability is usually defined through its antonym: *un*predictability [9, 10]. In an unpredictable system, the system's intrinsic sensitivity amplifies small perturbations and leads to massive divergence of trajectories. Predictable systems show the opposite behavior: small perturbations are diminished over time, rather than amplified. The notion of (un)predictability can be formalized through various routes such as chaos theory [22, 23] and contraction analysis [7, 24].

The definition of predictability comes from the Largest Lyapunov Exponent (LLE) [8, 10]:

Definition 1 (Predictability and Unpredictability). Consider a sequence of Jacobians, $J_1, J_2, \cdots J_T$. We define the associated Largest Lyapunov Exponent (LLE) to be

LLE :=
$$\lim_{T \to \infty} \frac{1}{T} \log \left(\|J_T J_{T-1} \cdots J_1\| \right) = \lambda, \tag{5}$$

where $\|\cdot\|$ is an induced operator norm. If $\lambda < 0$, we say that the nonlinear state space model is *predictable* at s_0 . Otherwise, we say it is *unpredictable*.

Suppose we wish to evaluate the nonlinear state space model (1) from an initial condition s_0 , but we only have access to an approximate measurement s_0' that differs slightly from the true initial state. If the system is unpredictable ($\lambda > 0$), then the distance between nearby trajectories grows as

$$||s_t - s_t'|| \sim e^{\lambda t} ||s_0 - s_0'||.$$
 (6)

Letting Δ denote the maximum acceptable deviation beyond which we consider the prediction to have failed, the time horizon over which the prediction remains reliable scales as

Time to degrade to
$$\Delta$$
 prediction error $\sim \frac{1}{\lambda} \log \left(\frac{\Delta}{\|s_0 - s_0'\|} \right)$. (7)

This relationship highlights a key limitation in unpredictable systems: even significant improvements in the accuracy of the initial state estimate yield only logarithmic gains in prediction time. The system's inherent sensitivity to initial conditions overwhelms any such improvements. Predictable systems, such as contracting systems, have the opposite property: trajectories initially separated by some distance will eventually converge towards one another (Figure 1), *improving* prediction accuracy over time.

3 Conditioning of Merit Function Depends on Predictability of Model

The number of optimization steps required to minimize the merit function (2) is impacted by its conditioning, which in our setting is determined by the smallest singular value of the residual function Jacobian. As we will see, what determines the smallest singular value of the residual function Jacobian is the stability, or predictability, of the underlying nonlinear state space model (1).

3.1 The Merit Function is PL

To begin, we show that the merit function (2) satisfies the Polyak-Łojasiewicz (PL) condition [25], also known as the gradient dominance condition [26]. A function $\mathcal{L}(\mathbf{s})$ is μ -PL if it satisfies, for $\mu > 0$,

$$\frac{1}{2}||\nabla \mathcal{L}(\mathbf{s})||^2 \ge \mu \left(\mathcal{L}(\mathbf{s}) - \mathcal{L}(\mathbf{s}^*)\right) \tag{8}$$

for all s. The largest μ for which eq. (8) holds for all s is called the PL constant of $\mathcal{L}(s)$.

Proposition 1. The merit function $\mathcal{L}(\mathbf{s})$ defined in eq. (2) satisfies eq. (8) for

$$\mu := \inf_{\mathbf{s}} \sigma_{\min}^2(\mathbf{J}(\mathbf{s})). \tag{9}$$

Proof. See Appendix B. This result, known in the literature for general sum-of-squares [27], is included here for context and completeness. \Box

Proposition 1 is important as it characterizes the *flatness* of the merit function. If μ is very small in a certain region, this indicates that the norm of the gradient can be very small in that region, which can make gradient-based optimization inefficient. Proposition 1 also links $\sigma_{\min}(\mathbf{J})$ —important for characterizing the conditioning of \mathbf{J} —to the geometry of the merit function landscape.

3.2 Merit Function PL Constant is Controlled by the Largest Lyapunov Exponent of Model

As stated earlier, the Largest Lyapunov Exponent is a commonly used way to define the (un)predictability of a nonlinear state space model. In order to proceed, we need to control more carefully how the product of Jacobian matrices in (5) behaves for finite-time products. We will assume that there exists a "burn-in" period where the norm of Jacobian products can transiently differ from the LLE. In particular, we assume that

$$\forall t > 1, \ \forall k \ge 0, \ \forall \mathbf{s}, \qquad b e^{\lambda k} \le \|J_{t+k-1}J_{t+k-2}\cdots J_t\| \le a e^{\lambda k}, \tag{10}$$

where $a \ge 1$ and $b \le 1$. The constant a quantifies the potential for transient growth—or over-shoot—in the norm of Jacobian products before their long-term behavior emerges, while b quantifies the potential for undershoot.

Theorem 2. Assume that the LLE regularity condition (10) holds. Then the PL constant μ satisfies

$$\frac{1}{a} \cdot \frac{e^{\lambda} - 1}{e^{\lambda T} - 1} \le \sqrt{\mu} \le \frac{1}{b} \cdot \frac{1}{e^{\lambda(T - 1)}}.$$
 (11)

Proof. See Appendix C for the full proof and discussion. We provide a brief sketch. Because $\sigma_{\min}(\mathbf{J}) = {}^{1}/\sigma_{\max}(\mathbf{J}^{-1})$, it suffices to control $\|\mathbf{J}^{-1}\|_{2}$. We can write $\mathbf{J} = \mathbf{I} - \mathbf{N}$ where \mathbf{N} is a nilpotent matrix. Thus, it follows that $\mathbf{J}^{-1} = \sum_{k=0}^{T-1} \mathbf{N}^{k}$. As we discuss further in Appendix C, the matrix powers \mathbf{N}^{k} are intimately related to the dynamics of the system. The upper bound on $\|\mathbf{J}^{-1}\|_{2}$ follows after applying the triangle inequality and the formula for a geometric sum. The lower bound follows from considering $\|\mathbf{N}^{T-1}\|_{2}$.

Theorem 2 is our main result, offering a novel connection between the predictability λ of a nonlinear state space model and the conditioning μ of the corresponding merit function, which affects whether the system can be effectively parallelized. If the underlying dynamics are unpredictable $(\lambda>0)$, then the merit function quickly becomes poorly conditioned with increasing T, because the denominators of both the lower and upper bounds explode due to the exponentially growing factor. Predictable dynamics $\lambda<0$ lead to good conditioning of the optimization problem, and parallel methods based on merit function minimization can be expected to perform well in these cases.

The proof mechanism we have sketched upper and lower bounds $\|\mathbf{J}^{-1}\|_2$ in terms of norms of Jacobian products. We only use the assumption in eq. (10) to express those bounds in terms of λ . As we discuss at length in Appendix C, we can use different assumptions from eq. (10) to get similar results. Theorem 2 and its proof should be thought of as a framework, where different assumptions (which may be more or less relevant in different settings) can be plugged in to yield specific results.

Why Unpredictable Systems have Excessively Flat Merit Functions Theorem 2 demonstrates that the merit function becomes extremely flat for unpredictable systems and long trajectories. This flatness poses a fundamental challenge for *any* method that seeks to compute state trajectories by minimizing the merit function. We now provide further intuition to explain why unpredictability in the system naturally leads to a flat merit landscape.

Suppose that we use an optimizer to minimize the merit function (2) for an unpredictable system until it halts with some precision. Let us further assume that the first state of the output of this optimizer following the initial condition is ϵ -close to the true first state, $||s_1 - s_1^*|| = \epsilon$. Suppose also that the residuals for all times greater than one are precisely zero—in other words, the optimizer starts with a "true" trajectory starting from initial condition s_1 . Then the overall residual norm is at most ϵ ,

$$\|\mathbf{r}(\mathbf{s})\|^2 = \|s_1 - f(s_0)\|^2 \le (\|s_1 - s_1^*\| + \|s_1^* - f(s_0)\|)^2 = \|s_1 - s_1^*\|^2 = \epsilon^2.$$

However, since s_t and s_t^* are by construction both trajectories of an unpredictable system starting from slightly different initial conditions s_1 and s_1^* , the distance between them will grow exponentially as a consequence of eq. (7). By contrast, predictable systems will have errors that shrink exponentially. This shows that changing the initial state s_1 by a small amount can lead to a massive change in the trajectory of an unpredictable system, but a *tiny change* in the merit function. Geometrically, this corresponds to the merit function landscape for unpredictable systems having excessive flatness around the true solution (Figure 1, bottom right panel). Predictable systems do not exhibit such flatness, since small residuals imply small errors. Theorem 2 formalizes this idea.

3.3 Residual function Jacobian Inherits the Lipschitzness of the Nonlinear State Space Model

In addition to the parameter μ , which measures the conditioning of the merit function, the difficulty of minimizing the merit function is also influenced by the Lipschitz continuity of its Jacobian J. The following theorem establishes how the Lipschitz continuity of the underlying sequence model induces Lipschitz continuity in J.

Theorem 3. If the dynamics of the underlying nonlinear state space model have L-Lipschitz Jacobians, i.e.,

$$\forall t > 1, \quad s, s' \in \mathbb{R}^D : \quad ||J_t(s) - J_t(s')|| \le L||s - s'||,$$

then the residual function Jacobian ${f J}$ is also L-Lipschitz, with the same L.

Proof. See Appendix D.
$$\Box$$

Theorem 3 will be important for the analysis in Section 4, where we consider convergence rates. Because Gauss-Newton methods rely on iteratively linearizing the dynamics (or equivalently the residual), they converge in a single step for linear dynamics L=0, and converge more quickly if the system is close to linear (L is closer to 0).

4 Rates of Convergence for Optimizing the Merit Function

In Section 3, we established that the predictability of the nonlinear state space model directly influences the conditioning of the merit function. This insight is critical for analyzing *any* optimization method used to compute trajectories via minimization of the merit function.

In this section, we apply those results to study the convergence behavior of the Gauss-Newton (DEER) algorithm for the merit function defined in eq. (2). See Appendix A for a brief overview of DEER. We derive worst-case bounds on the number of optimization steps required for convergence. In addition, we present an average-case analysis of DEER that is less conservative than the worst-case bounds and more consistent with empirical observations.

DEER Always Converges Globally at a Linear Rate Although DEER is based on the Gauss-Newton method, which generally lacks global convergence guarantees, we prove that DEER always converges globally at a linear rate. This result relies on the problem's specific hierarchical structure, which ensures that both the residual function Jacobian **J** and its inverse are lower block-triangular. In particular we prove the following theorem

Theorem 4. Let the DEER (Gauss–Newton) updates be given by eq. (3), and let $\mathbf{s}^{(i)}$ denote the *i*-th iterate. Let $\mathbf{e}^{(i)} := \mathbf{s}^{(i)} - \mathbf{s}^*$ denote the error at iteration i. Then the error converges to zero at a linear rate:

$$\|\mathbf{e}^{(i)}\|_{2} \le \chi_{w} \beta^{i} \|\mathbf{e}^{(0)}\|_{2},$$

for some constant $\chi_w \geq 1$ independent of i, and a convergence rate $0 < \beta < 1$.

Proof. See Appendix E.
$$\Box$$

Theorem 4 is unexpected since, in general, Gauss-Newton methods do not enjoy global convergence. The key caveat of this theorem is the multiplicative factor χ_w , which can grow exponentially with the sequence length T. This factor governs the extent of transient error growth before the decay term β^i eventually dominates.

Theorem 4 has several useful, practical consequences. First, when the nonlinear state space model is sufficiently contracting (λ is sufficiently negative), then χ_w in Theorem 4 can be made small, implying that in this case DEER converges with little-to-no overshoot (Appendix F).

Theorem 4 also lets us establish key worst-case and average-case bounds on the number of steps needed for Gauss-Newton to converge to within a given distance of the solution. In particular, when χ_w does not depend on the sequence length T, then Theorem 4 implies Gauss-Newton will only require $\mathcal{O}\left((\log T)^2\right)$ total computational time, with one log factor coming from the parallel scan at each optimization step and the other coming from the total number of optimization steps needed. We elaborate on these points in Appendix G.

Size of DEER Basin of Quadratic Convergence It is natural that DEER depends on the Lipschitzness of J since Gauss-Newton converges in one step for linear problems, where L=0. In Section 3, we showed that the conditioning of the merit function, as measured by the PL-constant μ , depends on the stability, or predictability, of the nonlinear dynamics. Thus, the performance of DEER depends on the ratio of the nonlinearity and stability of the underlying nonlinear state space model. Note that once s is inside the basin of quadratic convergence, it takes $O(\log\log(1/\epsilon))$ steps to reach ϵ residual (effectively a constant number of steps).

Theorem 5. Let μ denote the PL-constant of the merit function, which Theorem 2 relates to the LLE λ . Let L denote the Lipschitz constant of the Jacobian of the dynamics function J(s). Then, μ/L lower bounds the radius of the basin of quadratic convergence of DEER; that is, if

$$||\mathbf{r}(\mathbf{s}^{(i)})||_2 \le \frac{\mu}{L},\tag{12}$$

then $\mathbf{s}^{(i)}$ is inside the basin of quadratic convergence. In terms of the LLE λ , it follows that if

$$||\mathbf{r}(\mathbf{s}^{(i)})||_2 \le \frac{1}{a^2 L} \cdot \left(\frac{e^{\lambda} - 1}{e^{\lambda T} - 1}\right)^2,$$

then $\mathbf{s}^{(i)}$ is inside the basin of quadratic convergence.

Proof. See Appendix H. We make no claim about the originality of lower bounding the size of the basin of quadratic convergence in Gauss-Newton. In fact, our proof of Theorem 5 closely follows the convergence analysis of *Newton's* method in Section 9.5.3 of Boyd and Vandenberghe [28]. Our contribution is we highlight the elegant way the predictability λ and non-linearity L of a dynamical system influence an important feature of its merit function's landscape.

5 Experiments

We conduct experiments to support the theory developed above, demonstrating that predictability enables parallelization of nonlinear SSMs. To illustrate this point, we use Gauss-Newton optimization (aka DEER [1, 2]). We provide more experimental details in Appendix J. Our code is at https://github.com/lindermanlab/predictability_enables_parallelization

The Convergence Rate Exhibits a Threshold between Predictable and Chaotic Dynamics Theorem 2 predicts a sharp phase transition in the conditioning of the merit function at $\lambda=0$, which should be reflected in the number of optimization steps required for convergence. To empirically validate this prediction, we vary both the LLE and sequence length T within a parametric family of recurrent neural networks (RNNs), and measure the number of steps DEER takes to converge. We generate mean-field RNNs following Engelken et al. [29], scaling standard normal weight matrices by a single parameter that controls their variance and therefore the expected LLE. In Figure 2, we observe a striking correspondence between the conditioning of the optimization problem (represented by $-\log \tilde{\mu}$, where $\tilde{\mu}$ is the lower bound for μ from Theorem 2) and the number of steps DEER takes to convergence. This relationship holds across the range of LLEs, λ , and sequence lengths, T. There is a rapid threshold phenomenon around $\lambda=0$, which divides predictable from unpredictable dynamics, precisely as expected from Theorem 2. As we discuss in Appendix J.1, the correspondence between $-\log \tilde{\mu}$ and the number of optimization steps needed for convergence can be explained by DEER iterates approaching the basin of quadratic convergence with linear rate.

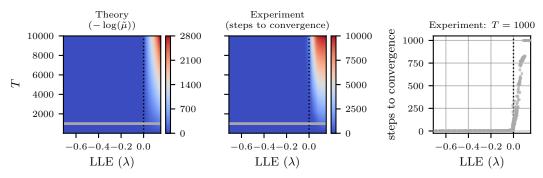


Figure 2: Threshold phenomenon in DEER convergence based on system predictability. In a family of RNNs, DEER has fast convergence for predictable systems and prohibitively slow convergence for chaotic systems. Left (Theory): We depict Theorem 2, illustrating how the conditioning of the optimization problem degrades as T and the LLE (λ) increase. Center (Experiment): We vary λ across the family of RNNs, and observe a striking concordance in the number of DEER optimization steps empirically needed for convergence with our theoretical characterization of the conditioning of the optimization problem. Right: For 20 seeds, each with 50 different values of λ , we plot the relationship between λ and the number of DEER steps needed for convergence for the sequence length T=1000 (gray line in left and center panels). We observe a sharp increase in the number of optimization steps at precisely the transition between predictability and unpredictability.

In Appendix J.3, we provide additional experiments in this setting. We parallelize the sequential rollout with other optimizers like quasi-Newton and gradient descent, and observe that the number of steps these optimizers take to converge also scales with the LLE. We also record wallclock times on an H100, and observe that DEER is faster than sequential by an order of magnitude in predictable settings, but slower by an order of magnitude in unpredictable settings.

DEER can converge quickly for predictable trajectories passing through unpredictable regions DEER may still converge quickly even if the system is unpredictable in certain regions. As long as the system is predictable on average, as indicated by a negative LLE, DEER can still converge quickly. This phenomenon is why we framed Theorem 2 in terms of the LLE λ and burn-in constants a, as opposed to a weaker result that assumes the system Jacobians have singular values less than one over the entire state space (see our discussion of condition (10) vs. condition (21) in Appendix C).

To illustrate, we apply DEER to Langevin dynamics in a two-well potential (visualized in Figure 3 for D=2). The dynamics are stable within each well but unstable in the region between them. Despite this local instability, the system's overall behavior is governed by time spent in the wells, resulting in a negative LLE and sublinear growth in DEER's convergence steps with sequence length T (Figure 3, right subplot). Additional details and discussion are in Appendix J.4.

Notably, prior works such as Lim et al. [1] and Gonzalez et al. [3] initialized optimization from $\mathbf{s}^{(0)} = \mathbf{0}$, which lies entirely in the unstable region. Thus, our theoretical insights into predictability and parallelizability suggest practical improvements for initialization.

Application: Chaotic Observers Finally, we demonstrate a practical application of our theory in the efficient parallelization of chaotic observers. Observers are commonly used to reconstruct the full state of a system from partial measurements [30, 31]. On nine chaotic flows from the dysts benchmark dataset [32], Table 1 shows that while DEER converges prohibitively slowly on chaotic systems, it converges rapidly on stable observers of these systems, in accordance with our theory that predictability implies parallelizability. For more details, see Appendix J.5.

6 Conclusion

Recent work has demonstrated that parallel computing hardware like GPUs can be used to rapidly compute state trajectories of nonlinear state space models. The central idea underlying these works is to reconceive the state trajectory as the solution to an optimization problem. Here, we provided a precise characterization of the optimization problem's inherent difficulty, which determines if

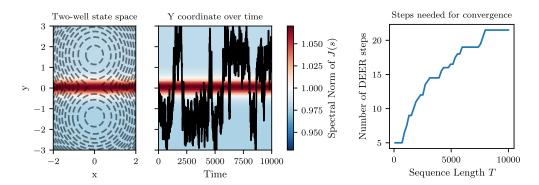


Figure 3: DEER converges quickly for Langevin dynamics in a two-well potential. (Left) An illustration of the two-well potential state space in D=2. We superimpose a contour plot of the potential on a color scheme showing the spectral norm of the dynamics Jacobian (blue indicates stability, red instability). (Center) A trace plot for the y-coordinate. The LLE of the system is -1.45. (Right) We observe that this system, which has negative LLE, enjoys sublinear scaling in the sequence length T in the number of DEER iterations needed to converge. We plot the median number of DEER steps to convergence over 20 random seeds.

Table 1: Comparison of system and observer LLEs and number of DEER steps for T = 30,000.

System	LLE (System)	LLE (Observer)	DEER Steps (System)	DEER Steps (Observer)
ABC	0.16	-0.08	4243	3
Chua's Circuit	0.02	-1.37	697	14
Kawczynski-Strizhak	0.01	-3.08	29396	2
Lorenz	1.02	-6.28	30000	3
Nosé-Hoover Thermostat	0.02	-0.13	29765	3
Rössler	0.01	-0.07	29288	7
SprottB	0.20	-0.39	29486	2
Thomas	0.01	-3.07	12747	7
Vallis El Niño	0.58	-2.48	30000	3

parallelization will be faster in practice than sequential evaluation. We show that the conditioning of the optimization problem is governed by the predictability of the underlying nonlinear system. We then translate this insight into worst-case performance guarantees for specific optimizers, including Gauss–Newton (DEER). Our main contribution can be summarized as: *Predictable dynamics yield well-conditioned merit functions, enabling rapid convergence. Unpredictable dynamics produce flat or ill-conditioned merit landscapes, resulting in slow convergence or, worse, numerical failure.*

Related Work The DeepPCR algorithm was introduced by Danieli et al. [2], who investigated its convergence rates empirically but not theoretically. Concurrently, Lim et al. [1] proposed an essentially identical method under the name DEER, proving local quadratic convergence for Gauss–Newton under standard assumptions but leaving the question of global convergence unresolved. Gonzalez et al. [3] proved the global convergence of DEER and other variants, though only with worst-case bounds of T optimization steps. None of these prior works addressed the relationship between system dynamics and conditioning, or established global linear convergence rates.

Global convergence rates for Gauss-Newton are rare, despite the breadth of optimization literature [19, 28, 33]. Theorem 4 establishes global convergence with linear rate for Gauss-Newton by leveraging our specific problem structure.

Fifty years ago, Hyafil and Kung [34] and Kung [35] showed that linear recursions enjoy speedups from parallel processors while nonlinear recursions of rational functions with degree larger than one cannot. These prescient works set the stage for our more general findings, which explicitly link the dynamical properties of the recursion to its parallelizability. Parallel-in-time methods for continuous

systems also have a long history [36, 37], with Chartier and Philippe [38] showing that dissipative systems can be parallelized using multiple shooting. Furthermore, Danieli and MacLachlan [39] and De Sterck et al. [40] study the CFL number for determining the usefulness of multigrid systems. Connecting this work with our paper is an interesting direction for future research.

More recently, several works have parallelized diffusion models via fixed-point iteration [2, 11, 41, 42], again with T-step worst-case guarantees. Anari et al. [43] proved $\log(T)$ rates for a particular dynamical system (Langevin dynamics) and a particular fixed-point iteration (Picard iteration). Crucially, prior work has not focused on the merit function, which we can define for any discrete-time dynamical system and optimizer.

To our knowledge, no prior work connects the LLE of a dynamical system to the conditioning of the corresponding optimization landscape, as established in Theorem 2. In particular, we showed that systems with high unpredictability yield poorly conditioned (i.e., flat) merit functions, linking dynamical instability to optimization difficulty in a geometrically appealing way.

The centrality of parallel sequence modeling architectures like transformers [44], deep SSMs [45, 21, 46], and linear RNNs [47] in modern machine learning underscores the need for our theoretical work. Merrill et al. [48] explored the question of parallelizability through the lens of circuit complexity, analyzing when deep-learning models can solve structured tasks in constant depth. Their focus complements ours, and suggests an opportunity for synthesis in future work.

Implications Nonlinear state space models are ubiquitous. Our results offer two key contributions.

First, they provide a principled way to determine, *a priori*, whether optimization-based parallelization of a given model is practical. In many robotic or control systems, particularly ones that are strongly dissipative, this insight can enable orders-of-magnitude speed-ups on GPUs [49–56].

In concurrent work, Zoltowski et al. [18] developed and leveraged quasi-Newton methods to parallelize Markov Chain Monte Carlo over the sequence length, attaining order-of-magnitude speedups. These speed-ups were made possible by the fast convergence of the quasi-Newton methods in the settings considered. Suggestively, there is an abundance of research studying the contractivity of MCMC in different settings [57–59]. In fact, the empirical results in this paper showing that Langevin dynamics can have negative LLE (cf. Figure 3) are suggestive that the Metropolisadjusted Langevin algorithm (MALA), a workhorse of MCMC, may also be predictable in settings of interest. A precise characterization of what makes an MCMC algorithm and target distribution contractive (i.e. predictable, in the language of our paper) would provide useful guidance for when one should aim to parallelize MCMC over the sequence length. Thus, providing precise theoretical justification for parallelizing MCMC over the sequence length is an exciting avenue for future work.

Second, our results have direct implications for system *design*. When constructing nonlinear dynamical systems in machine learning—such as novel RNNs—parallelization benefits are maximized when the system is made predictable. Given the large body of work on training stable RNNs [60–71], many effective techniques already exist for enforcing stability or predictability during training. A common approach is to *parameterize* the model's weights so that the model is always stable. For example, we provide a simple parameterization of a contractive SSM in Appendix I.

Notably, the concurrent work of Farsang et al. [71] develops a nonlinear SSM and trains it in parallel using quasi-DEER [3]. Farsang et al. [71] explicitly parameterizes their LrcSSM to be contractive (see their Appendix A.1). While they do not identify it as such, this contractivity is precisely what enables DEER to converge quickly throughout training. Ensuring a negative largest Lyapunov exponent through parameterization guarantees parallelizability for the entire training process, enabling faster and more scalable learning. Our contribution provides a theoretical foundation for why stability is essential in designing efficiently parallelizable nonlinear SSMs.

Limitations and Future Work While this work focuses on establishing the fundamental concepts and theoretical foundations, several practical considerations arise for scaling to large systems. Notably, DEER incurs a significant memory footprint. While this issue can be alleviated through quasi-Newton methods [3, 18], these approaches require more optimization steps to converge. Studying quasi-Newton methods with our theory could provide new insight into the efficacy of these methods.

Overall, the theoretical tools developed here have immediate implications for parallelizing nonlinear systems, and they open several exciting avenues for future work.

Acknowledgments

We thank members of the Linderman Lab for helpful feedback, particularly Noah Cowan, Henry Smith, Skyler Wu, and Etaash Katiyar. We thank Federico Danieli for helpful discussions.

X.G. would like to acknowledge support from the Walter Byers Graduate Scholarship from the NCAA. L.K. was a Goldstine Fellow at IBM Research while conducting this research. S.W.L. was supported by fellowships from the Simons Collaboration on the Global Brain, the Alfred P. Sloan Foundation, and the McKnight Foundation.

We thank Stanford University and the Stanford Research Computing Center for providing computational resources and support. Additional computations were performed on Marlowe [72], Stanford University's GPU-based Computational Instrument, supported by Stanford Data Science and Stanford Research Computing.

The authors have no competing interests to declare.

References

- [1] Yi Heng Lim, Qi Zhu, Joshua Selfridge, and Muhammad Firmansyah Kasim. Parallelizing non-linear sequential models over the sequence length. In *The Twelfth International Conference on Learning Representations*, 2024.
- [2] Federico Danieli, Miguel Sarabia, Xavier Suau Cuadros, Pau Rodriguez, and Luca Zappella. DeepPCR: Parallelizing sequential operations in neural networks. *Advances in Neural Information Processing Systems*, 36:47598–47625, 2023.
- [3] Xavier Gonzalez, Andrew Warrington, Jimmy T.H. Smith, and Scott W. Linderman. Towards Scalable and Stable Parallelization of Nonlinear RNNs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [4] Harold S. Stone. An efficient parallel algorithm for the solution of a tridiagonal linear system of equations. *Journal of the ACM*, 20(1):27–38, 1973.
- [5] Guy E. Blelloch. Prefix sums and their applications. Technical Report CMU-CS-90-190, School of Computer Science, Carnegie Mellon University, November 1990.
- [6] Xavier Gonzalez, E. Kelly Buchanan, Hyun Dong Lee, Jerry Weihong Liu, Ke Alexander Wang, David M. Zoltowski, Christopher Ré, and Scott W. Linderman. A unifying framework for parallelizing sequential models with linear dynamical systems. *arXiv*, 2025.
- [7] Winfried Lohmiller and Jean-Jacques E Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998.
- [8] Arkady Pikovsky and Antonio Politi. *Lyapunov exponents: a tool to explore complex dynamics*. Cambridge University Press, 2016.
- [9] Michael James Lighthill. The recently recognized failure of predictability in Newtonian dynamics. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 407(1832):35–50, 1986.
- [10] Steven H Strogatz. Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering. CRC press, 2018.
- [11] Andy Shih, Suneel Belkhale, Stefano Ermon, Dorsa Sadigh, and Nima Anari. Parallel sampling of diffusion models. *37th Conference on Neural Information Processing Systems*, 2023. 37th Conference on Neural Information Processing Systems.
- [12] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. In *International Conference on Learning Representations (ICLR)*, 2019. URL https://arxiv.org/abs/1807.03819.

- [13] Mark Schöne, Babak Rahmani, Heiner Kremer, Fabian Falck, Hitesh Ballani, and Jannes Gladrow. Implicit language models are rnns: Balancing parallelization and expressivity. In ICML, 2025. doi: 10.48550/arXiv.2502.07827. URL https://arxiv.org/abs/2502. 07827.
- [14] Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A recurrent depth approach, 2025. URL https://arxiv.org/abs/2502.05171.
- [15] Ramón Calvo-González, Daniele Paliotta, Matteo Pagliardini, Martin Jaggi, and François Fleuret. Leveraging the true depth of llms, 2025. URL https://arxiv.org/abs/2502.02790. Introduces Layer Parallelism for parallelizing adjacent Transformer layers.
- [16] Guan Wang, Jin Li, Yuhao Sun, Xing Chen, Changling Liu, Yue Wu, Meng Lu, Sen Song, and Yasin Abbasi Yadkori. Hierarchical reasoning model, 2025. URL https://arxiv.org/abs/2506.21734.
- [17] Alexia Jolicoeur-Martineau. Less is more: Recursive reasoning with tiny networks, 2025.
- [18] David M. Zoltowski, Skyler Wu, Xavier Gonzalez, Leo Kozachkov, and Scott W. Linderman. Parallelizing MCMC Across the Sequence Length. In Advances in Neural Information Processing Systems (NeurIPS), 2025.
- [19] Jorge Nocedal and Stephen J. Wright. Numerical Optimization. Springer, 2 edition, 2006.
- [20] Eric Martin and Chris Cundy. Parallelizing linear recurrent neural nets over sequence length. In *International Conference on Learning Representations*, 2018.
- [21] Jimmy T.H. Smith, Andrew Warrington, and Scott W. Linderman. Simplified state space layers for sequence modeling. In *International Conference on Learning Representations (ICLR)*, 2023.
- [22] James Gleick. Chaos: Making a new science. Penguin, 2008.
- [23] Heinz Georg Schuster and Wolfram Just. Deterministic chaos: an introduction. John Wiley & Sons, 2006.
- [24] F. Bullo. Contraction Theory for Dynamical Systems. Kindle Direct Publishing, 1.2 edition, 2024. ISBN 979-8836646806.
- [25] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part I 16*, pages 795–811. Springer, 2016.
- [26] Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476. PMLR, 2018.
- [27] Yurii Nesterov and B. T. Polyak. Cubic regularization of Newton method and its global performance. *Mathematical Programming, Series A*, 108(1):177–205, 2006. doi: 10.1007/s10107-006-0706-8. URL https://doi.org/10.1007/s10107-006-0706-8.
- [28] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004. ISBN 9780521833783.
- [29] Rainer Engelken, Fred Wolf, and Larry F Abbott. Lyapunov spectra of chaotic recurrent neural networks. *Physical Review Research*, 5(4):043044, 2023.
- [30] David G Luenberger. *Introduction to dynamic systems: theory, models, and applications.* John Wiley & Sons, 1979.
- [31] Dan Simon. Optimal state estimation: Kalman, H infinity, and nonlinear approaches. John Wiley & Sons, 2006.

- [32] William Gilpin. Chaos as an interpretable benchmark for forecasting and datadriven modelling. In Joaquin Vanschoren and Sai-Kit Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1 (NeurIPS Datasets and Benchmarks 2021)*, *December 2021*, *virtual*, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/ec5decca5ed3d6b8079e2e7e7bacc9f2-Abstract-round2.html.
- [33] Yurii Nesterov. Lectures on Convex Optimization, volume 137 of Springer Optimization and Its Applications. Springer, 2nd edition, 2018. ISBN 978-3-319-91577-4. doi: 10.1007/978-3-319-91578-1.
- [34] L Hyafil and HT Kung. *Bounds on the speed-up of parallel evaluation of recurrences*. Carnegie Mellon University, Department of Computer Science, 1975.
- [35] HT Kung. New algorithms and lower bounds for the parallel evaluation of certain rational expressions and recurrences. *Journal of the ACM (JACM)*, 23(2):252–261, 1976.
- [36] Martin J Gander. 50 years of time parallel time integration. In Multiple Shooting and Time Domain Decomposition Methods: MuS-TDD, Heidelberg, May 6-8, 2013, pages 69–113. Springer, 2015.
- [37] Benjamin W Ong and Jacob B Schroder. Applications of time parallelization. *Computing and Visualization in Science*, 23:1–15, 2020.
- [38] Philippe Chartier and Bernard Philippe. Eine parallele "shooting" technik zur lösung dissipativer gewöhnlicher differentialgleichungen. *Computing*, 51:209–236, 1993.
- [39] Federico Danieli and Scott MacLachlan. Multigrid reduction in time for non-linear hyperbolic equations. *arXiv preprint arXiv:2104.09404*, 2021.
- [40] Hans De Sterck, Stephanie Friedhoff, Oliver A Krzysik, and Scott P MacLachlan. Multigrid Reduction-In-Time Convergence for Advection Problems: A Fourier Analysis Perspective. *Numerical Linear Algebra with Applications*, 32(1):e2593, 2025.
- [41] Zhiwei Tang, Jiasheng Tang, Hao Luo, Fan Wang, and Tsung-Hui Chang. Accelerating parallel sampling of diffusion models. In Forty-first International Conference on Machine Learning, 2024.
- [42] Nikil Selvam, Amil Merchant, and Stefano Ermon. Self-Refining Diffusion Samplers: Enabling Parallelization via Parareal Iterations. In *Advances in Neural Information Processing Systems*, volume 37, pages 5429–5453, 2024.
- [43] Nima Anari, Sinho Chewi, and Thuy-Duong Vuong. Fast parallel sampling under isoperimetry. In *The Thirty Seventh Annual Conference on Learning Theory*, pages 161–185. PMLR, 2024.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All You Need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [45] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations (ICLR)*, 2021.
- [46] Albert Gu and Tri Dao. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. In *Conference on Language Modeling (COLM)*, 2024.
- [47] Songlin Yang, Bailin Wang, Yu Zhang, Yikang Shen, and Yoon Kim. Parallelizing linear transformers with the delta rule over sequence length. In *Proceedings of NeurIPS*, 2024.
- [48] William Merrill, Jackson Petty, and Ashish Sabharwal. The illusion of state in state-space models. In *Forty-first International Conference on Machine Learning*, 2024.
- [49] J Zico Kolter and Gaurav Manek. Learning stable deep dynamics models. *Advances in neural information processing systems*, 32, 2019.

- [50] Hadi Beik-Mohammadi, Søren Hauberg, Georgios Arvanitidis, Nadia Figueroa, Gerhard Neumann, and Leonel Rozo. Neural contractive dynamical systems. arXiv preprint arXiv:2401.09352, 2024.
- [51] Sean Jaffe, Alexander Davydov, Deniz Lapsekili, Ambuj K Singh, and Francesco Bullo. Learning neural contracting dynamics: Extended linearization and global guarantees. Advances in Neural Information Processing Systems, 37:66204–66225, 2024.
- [52] Fletcher Fan, Bowen Yi, David Rye, Guodong Shi, and Ian R Manchester. Learning stable koopman embeddings. In 2022 American Control Conference (ACC), pages 2742–2747. IEEE, 2022.
- [53] Vikas Sindhwani, Stephen Tu, and Mohi Khansari. Learning contracting vector fields for stable imitation learning. *arXiv preprint arXiv:1804.04878*, 2018.
- [54] Dawei Sun, Susmit Jha, and Chuchu Fan. Learning certified control using contraction metric. In *conference on Robot Learning*, pages 1519–1539. PMLR, 2021.
- [55] Hiroyasu Tsukamoto, Soon-Jo Chung, and Jean-Jacques E Slotine. Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview. *Annual Reviews in Control*, 52:135–169, 2021.
- [56] Max Revay, Ruigang Wang, and Ian R Manchester. Recurrent equilibrium networks: Flexible dynamic models with guaranteed stability and robustness. *IEEE Transactions on Automatic Control*, 69(5):2855–2870, 2023.
- [57] Nawaf Bou-Rabee, Andreas Eberle, and Raphael Zimmer. Coupling and convergence for Hamiltonian Monte Carlo. The Annals of Applied Probability, 30(3):1209-1250, June 2020. doi: 10.1214/19-AAP1528. URL https://projecteuclid.org/journals/annals-of-applied-probability/volume-30/issue-3/Coupling-and-convergence-for-Hamiltonian-Monte-Carlo/10.1214/19-AAP1528.
- [58] Oren Mangoubi and Aaron Smith. Mixing of Hamiltonian Monte Carlo on strongly log-concave distributions: Continuous dynamics. *The Annals of Applied Probability*, 31(5): 2019–2045, October 2021. doi: 10.1214/20-AAP1640. URL https://projecteuclid.org/journals/annals-of-applied-probability/volume-31/issue-5/Mixing-of-Hamiltonian-Monte-Carlo-on-strongly-log-concave-distributions/10.1214/20-AAP1640.
- [59] Persi Diaconis and David Freedman. Iterated random functions. SIAM Review, 41(1):45–76, 1999.
- [60] Sepp Hochreiter. Untersuchungen zu dynamischen neuronalen Netzen. Diploma thesis, Technische Universität München, Munich, Germany, 1991.
- [61] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9 (8):1735–1780, 1997.
- [62] Ilya Sutskever. Training Recurrent Neural Networks. PhD thesis, University of Toronto, Graduate Department of Computer Science, Toronto, Canada, 2013. PhD thesis.
- [63] John Miller and Moritz Hardt. Stable recurrent models. In *International Conference on Learn-ing Representations*, 2019.
- [64] N Benjamin Erichson, Omri Azencot, Alejandro Queiruga, Liam Hodgkinson, and Michael W Mahoney. Lipschitz recurrent neural networks. *arXiv preprint arXiv:2006.12070*, 2020.
- [65] Leo Kozachkov, Michaela Ennis, and Jean-Jacques Slotine. Rnns of rnns: Recursive construction of stable assemblies of recurrent neural networks. *Advances in neural information processing systems*, 35:30512–30527, 2022.
- [66] Karan Goel, Albert Gu, Chris Donahue, and Christopher Re. It's raw! audio generation with state-space models. In *International Conference on Machine Learning*, 2022.

- [67] Dmitry Krotov. A new frontier for Hopfield networks. Nature Reviews Physics, 5(7):366–367, 2023.
- [68] Rainer Engelken. Gradient flossing: Improving gradient descent through dynamic control of jacobians. *Advances in Neural Information Processing Systems*, 36:10412–10439, 2023.
- [69] Antonio Orvieto, Samuel L Smith, Albert Gu, Anushan Fernando, Caglar Gulcehre, Razvan Pascanu, and Soham De. Resurrecting recurrent neural networks for long sequences. In *International Conference on Machine Learning*, pages 26670–26698. PMLR, 2023.
- [70] Nicolas Zucchet and Antonio Orvieto. Recurrent neural networks: vanishing and exploding gradients are not the end of the story. In *Neural Information Processing Systems (NeurIPS)*, 2024.
- [71] Mónika Farsang, Ramin Hasani, Daniela Rus, and Radu Grosu. Scaling Up Liquid-Resistance Liquid-Capacitance Networks for Efficient Sequence Modeling. In Advances in Neural Information Processing Systems (NeurIPS), 2025.
- [72] Craig Kapfer, Kurt Stine, Balasubramanian Narasimhan, Christopher Mentzel, and Emmanuel Candès. Marlowe: Stanford's GPU-based Computational Instrument. Zenodo, 2025. Version 0.1.
- [73] Xavier Gonzalez. Parallelizing Nonlinear RNNs with the Ungulates: DEER and ELK, December 2 2024. URL https://lindermanlab.github.io/hackathons/. Linderman Lab Blog.
- [74] Sinho Chewi and Austin J. Stromme. The ballistic limit of the log-sobolev constant equals the polyak-łojasiewicz constant. *Annales de l'Institut Henri Poincaré (B) Probabilités et Statistiques*, 2025. URL https://arxiv.org/abs/2411.11415.
- [75] Chaoyue Liu, Libin Zhu, and Mikhail Belkin. Loss landscapes and optimization in over-parameterized non-linear systems and neural networks. *Applied and Computational Harmonic Analysis*, 59:85–116, 2022.
- [76] Mark Konstantinovich Gavurin. Nonlinear functional equations and continuous analogues of iteration methods. *Izvestiya Vysshikh Uchebnykh Zavedenii. Matematika*, pages 18–31, 1958.
- [77] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In 3rd International Conference on Learning Representations (ICLR), 2015. URL https://arxiv.org/abs/1412.6980.
- [78] Vineet Gupta, Tomer Koren, and Yoram Singer. Shampoo: Preconditioned stochastic tensor optimization. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1842–1850. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/gupta18a.html.
- [79] Nikhil Vyas, Depen Morwani, Rosie Zhao, Itai Shapira, David Brandfonbrener, Lucas Janson, and Sham M. Kakade. SOAP: Improving and Stabilizing Shampoo using Adam for Language Modeling. In *The Thirteenth International Conference on Learning Representations (ICLR 2025)*, 2025. URL https://arxiv.org/abs/2409.11321.
- [80] Paul Langevin. On the Theory of Brownian Motion. American Journal of Physics, 65 (11):1079–1081, 1908. English translation, introduced by D. S. Lemons and translated by A. Gythiel. Original: C. R. Acad. Sci. 146, 530–533 (1908).
- [81] Roy Friedman. A simplified overview of Langevin dynamics. Blog post, https://friedmanroy.github.io/blog/2022/Langevin/, 2022.
- [82] Desmond J. Higham. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review*, 43(3):525–546, 2001.
- [83] Matthew MacKay. Stats 305c lecture 17: Stochastic differential equations. *Course slides*, STATS 305C: Applied Statistics III, Stanford University, 2022.

- [84] Peter E. Holderrieth. Langevin dynamics: An introduction for machine learning engineers, 2023. URL https://www.peterholderrieth.com/blog/2023/Langevin-Dynamics-An-introduction-for-Machine-Learning-Engineers/. Blog post.
- [85] William Gilpin. Chaos as an interpretable benchmark for forecasting and data-driven modelling. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [86] Louis M Pecora and Thomas L Carroll. Synchronization in chaotic systems. *Physical review letters*, 64(8):821, 1990.
- [87] Ali Zemouche and Mohamed Boutayeb. Observer design for Lipschitz nonlinear systems: the discrete-time case. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 53(8): 777–781, 2006.
- [88] Bertrand Legras and Robert Vautard. A guide to Liapunov vectors. In *Proceedings of the ECMWF Seminar on Predictability, 4–8 September 1995*, volume 1, pages 143–156, Reading, UK, 1996. ECMWF.

Appendix

A	Brief Overview of DEER/DeepPCR				
В	Merit Function is PL				
C	Merit Function PL Constant is Controlled by Largest Lyapunov Exponent of Model C.1 A Generalized Proof that the Largest Lyapunov Exponent Controls the PL Constant				
D	DEI	ER Merit Function Inherits Lipschitzness of Dynamics Jacobians	26		
E	E DEER Always Converges Linearly				
F	DEER Converges Globally with Small Overshoot for Sufficiently Strongly Contracting Systems				
G	G Alternative Descent Techniques & Worst/Average Complexity				
Н	I Proof of Size of Basin of Quadratic Convergence				
I	Parameterizing Nonlinear SSMs to be contractive				
J	Exp	erimental Details and Discussion	32		
	J.1	Deriving the Empirical Scaling of DEER	32		
	J.2	Details and Discussion for mean-field RNN experiment	32		
	J.3	Additional experiment for the mean-field RNN: other optimizers and wallclock time	34		
	J.4	Additional details for the two-well potential	36		
	J.5	Building Stable Observers for Chaotic Systems	37		
	J.6	Numerical computation of the discrete-time LLE	37		
K	Disc	rete and Continuous Time LLE	38		

Brief Overview of DEER/DeepPCR

This section provides background on DEER/DeepPCR needed to support section 4 of the main text. Other options for further background on DEER are sections 2-4 of Gonzalez et al. [3] and the corresponding blog post [73].

We begin with a brief review of DEER/DeepPCR [2, 1, 3]. As mentioned in the introduction, the choice of optimizer is crucial for this procedure to outperform sequential evaluation in terms of wall clock time. Indeed, for this reason DEER uses the Gauss-Newton method (GN) to minimize the residual loss, since GN exhibits quadratic convergence rates near the optimum [19]. Recall from eq. (3) that the i-th step of the DEER algorithm is,

$$\mathbf{s}^{(i+1)} = \mathbf{s}^{(i)} - \mathbf{J}(\mathbf{s}^{(i)})^{-1} \mathbf{r}(\mathbf{s}^{(i)}).$$

This step requires inverting the $TD \times TD$ matrix, $\mathbf{J}(\mathbf{s}^{(i)})$. Rather than explicitly inverting it, which is generally infeasible, DEER solves for the updates by running a linear time-varying recursion [3]:

$$\Delta s_t^{(i+1)} = J_t^{(i)} \, \Delta s_{t-1}^{(i+1)} - r_t(s^{(i)}), \qquad \text{where} \qquad \Delta s_t^{(i+1)} \; := s_t^{(i+1)} - s_t^{(i)} \tag{13}$$

Unlike the standard sequential rollout, this recursion can be parallelized and computed in $O(\log T)$ time using a parallel associative scan [5]. When the number of optimization steps needed for DEER to converge to the true trajectory is relatively small, DEER can yield faster overall evaluation than the sequential approach. Since Gauss-Newton converges quadratically when the initial guess is sufficiently close to the true optimum [1, 19], DEER potentially only requires a tiny number of iterations to converge. Our first key result is to prove that DEER always converges globally with linear rate, and will thus always reach this basin of quadratic convergence after sufficient time.

A note about notation The DEER quantities:

- residual $\mathbf{r}(\mathbf{s}) \in \mathbb{R}^{TD}$ Jacobian $\mathbf{J}(\mathbf{s}) \in \mathbb{R}^{TD \times TD}$
- merit function $\mathcal{L}(\mathbf{s}) \in \mathbb{R}$

are functions of the current guess for the trajectory $\mathbf{s} = \text{vec}(s_1, \dots, s_T) \in \mathbb{R}^{TD}$. As much as possible, we try to emphasize the dependence on the current guess for the trajectory, but sometimes we will drop the dependence for notational compactness.

В **Merit Function is PL**

This section provides a proof of main text Proposition 1. We first note that Proposition 1 applies to optimizing any nonlinear sum of squares problem where $\mathcal{L}(\mathbf{s}) = \frac{1}{2} ||\mathbf{r}(\mathbf{s})||_2^2$, not just the r we consider in this paper (defined in eq. (2)).

Proposition (Proposition 1). The merit function $\mathcal{L}(\mathbf{s})$ defined in eq. (2) satisfies eq. (8) for

$$\mu := \inf_{\mathbf{s}} \sigma_{\min}^2(\mathbf{J}(\mathbf{s})).$$

Proof. Observe that

$$\nabla \mathcal{L}(\mathbf{s}) = \mathbf{J}(\mathbf{s})^{\top} \mathbf{r}(\mathbf{s})$$
 and $\mathcal{L}(\mathbf{s}^*) = 0$.

Substituting these expressions into the PL inequality in eq. (8) we obtain

$$\mathbf{r}^{\top} \mathbf{J}(\mathbf{s}) \mathbf{J}(\mathbf{s})^{\top} \mathbf{r} \geq \mu \mathbf{r}^{\top} \mathbf{r}.$$

Therefore, if **J** is full rank, then the merit function \mathcal{L} is μ -PL, where

$$\mu = \inf_{\mathbf{s}} \lambda_{\min} \left(\mathbf{J}(\mathbf{s}) \mathbf{J}(\mathbf{s})^{\top} \right)$$
$$= \inf_{\mathbf{s}} \sigma_{\min}^{2} \left(\mathbf{J}(\mathbf{s}) \right)$$

To be precise, we must have $\mu > 0$ for \mathcal{L} to satisfy the definition of PL. Therefore, a condition that must apply for \mathcal{L} to be PL is that we must have $\inf_{\mathbf{s}} \sigma_{\min}(\mathbf{J}(\mathbf{s})) > 0$. We note that the proof strategy of Theorem 2 ensures that $\inf_{\mathbf{s} \in \mathbb{R}^{TD}} \sigma_{\min}(\mathbf{J}(\mathbf{s})) > 0$ if we assume eq. (21), which holds for dynamical systems that are globally contracting.

By the chain rule, eq. (21) also holds for functions of the form $f(s) = \phi(Ws)$, where $W \in \mathbb{R}^{D \times D}$ and ϕ is a scalar function with bounded derivative that is applied elementwise. In particular, such a function $\phi(Ws)$ satisfies eq. (21) whether or not it is globally contracting. This function class is extremely common in deep learning (nonlinearities with bounded derivatives include \tanh , the logistic function and ReLU).

In our statement and proof of Proposition 1, we deliberately do not specify the set over which we take the infimum. The result is true regardless of what this set is taken to be. The largest such set would be \mathbb{R}^{TD} , but other sets that could be of interest are the optimization trajectory $\{\mathbf{s}^{(i)}, i \in \mathbb{N}\}$, or alternatively a neighborhood of the solution \mathbf{s}^* . We discuss further in Appendix C.

Some more general notes on the PL inequality The PL inequality or gradient dominance condition is stated differently in different texts [27, 26, 74]. We follow the presentation of Karimi et al. [25]. Karimi et al. [25] emphasizes that PL is often weaker than many other conditions that had been assumed in the literature to prove linear convergence rates.

We note that the PL inequality as stated in eq. (8) is not invariant to the scaling of \mathcal{L} . However, in Definition 3 of Nesterov and Polyak [27], they broaden the definition to be gradient dominant of degree $p \in [1, 2]$. The PL inequality we state in eq. (8) corresponds to gradient dominance of degree 2. Note that gradient dominance of degree 1 is scale-invariant.

C Merit Function PL Constant is Controlled by Largest Lyapunov Exponent of Model

This section provides the proof of main text Theorem 2.

Theorem (Theorem 2). Assume that the LLE regularity condition (10) holds. Then if $\lambda \neq 0$ the PL constant μ of the merit function in (8) satisfies

$$\frac{1}{a} \cdot \frac{e^{\lambda} - 1}{e^{\lambda T} - 1} \le \sqrt{\mu} \le \frac{1}{b} \cdot \frac{1}{e^{\lambda (T - 1)}}.$$
 (14)

By L'Hôpital's rule, if $\lambda = 0$, then the bounds are instead

$$\frac{1}{aT} \leq \sqrt{\mu} \leq \frac{1}{b}$$
.

Proof. We present two proofs. A shorter, direct proof of (14) assuming $\|\cdot\|$ is the standard Euclidean norm, and then a more general version in Appendix C.1, which will be useful later on.

Notice that the residual function Jacobian J (4) in can be written as the difference of the identity and a T-nilpotent matrix N, as

$$\mathbf{J} = \mathbf{I}_{TD} - \mathbf{N}$$
 with $\mathbf{N}^T = \mathbf{0}_{TD}$

Because ${\bf N}$ is nilpotent, the Neumann series for ${\bf J}^{-1}$ is a finite sum:

$$\mathbf{J}^{-1} = (\mathbf{I}_{TD} - \mathbf{N})^{-1} = \sum_{k=0}^{T-1} \mathbf{N}^k.$$
 (15)

Straightforward linear algebra also shows that the norms of the powers of this nilpotent matrix are bounded, which enables one to upper bound the inverse of the Jacobian

$$\|\mathbf{N}^k\|_2 \le a e^{\lambda k}$$
 and therefore $\|\mathbf{J}^{-1}\|_2 \le \sum_{k=0}^{T-1} \|\mathbf{N}^k\|_2 \le \sum_{k=0}^{T-1} a e^{\lambda k} = a \frac{1 - e^{\lambda T}}{1 - e^{\lambda}}.$ (16)

The powers of N are closely related to the dynamics of the nonlinear state space model. We provide a dynamical interpretation below, in the paragraph "The dynamical interpretation of N and its powers".

To lower bound $\|\mathbf{J}^{-1}\|_2$, we observe that by the SVD, a property of the spectral norm is that

$$\|\mathbf{J}^{-1}\|_{2} = \sup_{\substack{\|x\|_{2}=1\\\|y\|_{2}=1}} x^{\top} \mathbf{J}^{-1} y. \tag{17}$$

We pick two unit vectors u and v, both in \mathbb{R}^{TD} , that are zero everywhere other than where they need to be to pull out the bottom-left block of \mathbf{J}^{-1} (i.e., the only non-zero block in \mathbf{N}^{T-1} , which is equal to $J_T J_{T-1} \dots J_2$). Doing so, we get

$$u^T \mathbf{J}^{-1} v = \tilde{u}^T (J_T J_{T-1} \dots J_2) \tilde{v},$$

where \tilde{u} and \tilde{v} are unit vectors in \mathbb{R}^D , and are equal to the nonzero entries of u and v.

Note, therefore, that because of eq. (17), it follows that

$$\tilde{u}^T (J_T J_{T-1} \dots J_2) \tilde{v} \le \|\mathbf{J}^{-1}\|_2,$$
(18)

i.e. we also have a **lower bound** on $\|\mathbf{J}^{-1}\|_2$.

Furthermore, choosing \tilde{u} and \tilde{v} to make

$$\tilde{u}^T (J_T J_{T-1} \dots J_2) \, \tilde{v} = \|J_T J_{T-1} \dots J_2\|_2,$$

we can plug in this choice of \tilde{u} and \tilde{v} into eq. (18), to obtain

$$||J_T J_{T-1} \dots J_2||_2 \le ||\mathbf{J}^{-1}||_2.$$

Applying the regularity conditions (10) for k = T - 1 and t = 2 we obtain

$$b e^{\lambda(T-1)} \le \|\mathbf{J}^{-1}\|_2.$$
 (19)

Because

$$\lambda_{\min} \left(\mathbf{J} \mathbf{J}^{\top} \right) = \frac{1}{\| \mathbf{J}^{-1} \|_2^2},$$

the result follows by applying eq. (16) and eq. (19) at all $s^{(i)}$ along the optimization trajectory.

The above proof sheds light on how many dynamical system properties fall out of the structure of J(s), which we now discuss further.

Discussion of why small $\sigma_{\min}(\mathbf{J}(\mathbf{s}))$ **leads to ill-conditioned optimization** Recall that our goal is to find a lower bound on the smallest singular value of $\mathbf{J}(\mathbf{s})$, which we denote by $\sigma_{\min}(\mathbf{J}(\mathbf{s}))$. This quantity controls the difficulty of optimizing \mathcal{L} . For example, the Gauss-Newton update is given by $\mathbf{J}(\mathbf{s})^{-1}\mathbf{r}(\mathbf{s})$. Recall that

$$\sigma_{\max} \left(\mathbf{J}(\mathbf{s})^{-1} \right) = \frac{1}{\sigma_{\min}(\mathbf{J}(\mathbf{s}))}$$
$$= \|\mathbf{J}(\mathbf{s})^{-1}\|_{2}.$$

Recall that an interpretation of the spectral norm $\|\mathbf{J}(\mathbf{s})\|_2$ is how much multiplication by $\mathbf{J}(\mathbf{s})$ can increase the length of a vector. Therefore, we see that very small values of $\sigma_{\min}(\mathbf{J}(\mathbf{s}))$ result in large values of $\|\mathbf{J}(\mathbf{s})^{-1}\|_2$, which means that $\|\mathbf{J}(\mathbf{s})^{-1}\mathbf{r}(\mathbf{s})\|_2$ can become extremely large as well, and small perturbations in \mathbf{r} can lead to very different Gauss-Newton updates (i.e. the problem is ill-conditioned, cf. Nocedal and Wright [19] Appendix A.1).

Furthermore, we observe that in the $\lambda>0$ (unpredictable) setting and the large T limit, the upper and lower bounds in (14) are tight, as they are both $\mathcal{O}(e^{\lambda(T-1)})$. Thus, the upper and lower bounds together ensure that unpredictable dynamics will suffer from degrading conditioning.

In contrast, in the $\lambda < 0$ (predictable) setting, the lower bound on $\sqrt{\mu}$ converges to $\frac{1-e^{\lambda}}{a}$, which is bounded away from zero and *independent of the sequence length*. Thus, in predictable dynamics, there is a lower bound on $\sigma_{\min}(\mathbf{J})$ or, equivalently, an upper bound on $\sigma_{\max}(\mathbf{J}^{-1})$.

The dynamical interpretation of N and its powers As shown in the above proof,

$$\mathbf{J}(\mathbf{s})^{-1} = (\mathbf{I}_{TD} - \mathbf{N}(\mathbf{s}))^{-1} = \sum_{k=0}^{T-1} \mathbf{N}(\mathbf{s})^k.$$

It is worth noting explicitly that

$$\mathbf{N}(\mathbf{s}) = \begin{pmatrix} 0 & 0 & \dots & 0 & 0 \\ J_2 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & J_T & 0 \end{pmatrix} \quad \text{where} \quad J_t := \frac{\partial f_t}{\partial s_{t-1}}(s_{t-1}), \tag{20}$$

i.e. N(s) collects the Jacobians of the dynamics function along the first lower diagonal. Each matrix power N^k therefore collects length k products along the kth lower diagonal. Thus, multiplication by $J(s)^{-1} = \sum_{k=0}^{T-1} N(s)^k$ recovers running forward a linearized form of the dynamics, which is one of the core insights of DeepPCR and DEER [2, 1].

Concretely, in the setting where T=4, we have

A framing of Theorem 2 based on global bounds on $||J_t||_2$ We chose to prove Theorem 2 using condition (10) in order to highlight the natural connection between the smallest singular value of **J** and system stability (as measured by its LLE). However, an assumption with a different framing would be to impose a uniform bound on the spectral norm of the Jacobian over the entire state space:

$$\sup_{s \in \mathbb{R}^D} ||J(s)||_2 \le \rho. \tag{21}$$

For $\rho < 1$, this assumption corresponds to global contraction of the dynamics [7].

If we replace the LLE regularity condition (10) with the global spectral norm bound (21) in the proof of Theorem 2, we obtain that the PL constant is bounded away from zero, i.e.

$$\frac{1}{a} \cdot \frac{\rho - 1}{\rho^T - 1} \le \sqrt{\inf_{\mathbf{s} \in \mathbb{R}^{T_D}} \sigma_{\min}^2(\mathbf{J}(\mathbf{s}))}.$$

In particular, if the dynamics are contracting everywhere (i.e., $\rho < 1$), the condition (21) guarantees good conditioning of **J** throughout the entire state space.

Discussion of the LLE regularity conditions The LLE regularity conditions in eq. (10) highlight the more natural "average case" behavior experienced along actual trajectories $\mathbf{s} \in \mathbb{R}^{TD}$. This "average case" behavior is highlighted, for example, by our experiments with the two-well system (cf. Section 5 and Appendix J.4), where even though a global upper bound on $\|J_t(s_t)\|_2$ over all of state space would be greater than 1 (i.e., there are unstable regions of state space), we observe fast convergence of DEER because the system as a whole has negative LLE (its trajectories are stable on average).

We also note the pleasing relationship the LLE regularity conditions have with the definition of the LLE given in eq. (5). Note that in the LLE regularity conditions in eq. (10), the variable k denotes the sequence length under consideration. Taking logs and dividing by k, we therefore obtain

$$\frac{\log b}{k} + \lambda \le \frac{1}{k} \log \left(\|J_{t+k-1} J_{t+k-2} \cdots J_t\| \right) \le \frac{\log a}{k} + \lambda.$$

Therefore, as $k \to T$, and as $T \to \infty$ (i.e., we consider longer and longer sequences), we observe that the finite-time estimates of the LLE converge to the true LLE λ .

We observe that as $s^{(i)}$ approaches the true solution s^* , the regularity conditions in eq. (10) become increasingly reasonable. Since any successful optimization trajectory must eventually enter a neighborhood of s^* , it is natural to expect these conditions to hold there. In fact, rather than requiring the regularity conditions over all of state space or along the entire optimization trajectory, one could alternatively assume that they hold within a neighborhood of s^* , and prove a corresponding version of Theorem 2.

We now do so, using the additional assumption that J is L-Lipschitz.

Theorem 6. If **J** is L-Lipschitz, then there exists a ball of radius R around the solution s^* , denoted $B(s^*, R)$, such that

$$\forall \mathbf{s} \in B(\mathbf{s}^*, R) \qquad |\sigma_{\min}(\mathbf{J}(\mathbf{s})) - \sigma_{\min}(\mathbf{J}(\mathbf{s}^*))| \leq LR$$

Proof. The argument parallels the proof of Theorem 2 in Liu et al. [75].

A fact stemming from the reverse triangle inequality is that for any two matrices A and B,

$$\sigma_{\min}(\mathbf{A}) \geq \sigma_{\min}(\mathbf{B}) - \|\mathbf{A} - \mathbf{B}\|.$$

Applying this with A = J(s) and $B = J(s^*)$, we obtain

$$\sigma_{\min}(\mathbf{J}(\mathbf{s})) \geq \sigma_{\min}(\mathbf{J}(\mathbf{s}^*)) - \|\mathbf{J}(\mathbf{s}) - \mathbf{J}(\mathbf{s}^*)\|.$$

If the Jacobian $J(\cdot)$ is L-Lipschitz, then

$$\|\mathbf{J}(\mathbf{s}) - \mathbf{J}(\mathbf{s}^*)\| \le L\|\mathbf{s} - \mathbf{s}^*\|.$$

Combining, we get

$$\sigma_{\min}(\mathbf{J}(\mathbf{s})) \geq \sigma_{\min}(\mathbf{J}(\mathbf{s}^*)) - L \|\mathbf{s} - \mathbf{s}^*\|$$

and

$$\sigma_{\min}(\mathbf{J}(\mathbf{s}^*)) \geq \sigma_{\min}(\mathbf{J}(\mathbf{s})) - L \|\mathbf{s} - \mathbf{s}^*\|,$$

which gives

$$\sigma_{\min}(\mathbf{J}(\mathbf{s}^*)) - L \|\mathbf{s} - \mathbf{s}^*\| \le \sigma_{\min}(\mathbf{J}(\mathbf{s})) \le \sigma_{\min}(\mathbf{J}(\mathbf{s}^*)) + L \|\mathbf{s} - \mathbf{s}^*\|.$$

Ensuring that $\|\mathbf{s} - \mathbf{s}^*\| \le R$ completes the proof.

A consequence of Theorem 6 is that if the system is unpredictable, then there exists a finite ball around s* where the conditioning of the merit function landscape is provably bad.

As a concrete example, suppose that $\sigma_{\min}(\mathbf{J}(\mathbf{s}^*)) = \epsilon$ and L = 1. Then *at best*, the PL constant of the loss function inside the ball $B(\mathbf{s}^*, R)$ is $\epsilon + R$. If ϵ is small (bad conditioning) then R can be chosen such that the PL constant inside the ball $B(\mathbf{s}^*, R)$ is also small.

Controlling $\sigma_{\max}(\mathbf{J})$ In our proof of Theorem 2, we proved upper and lower bounds for $\sigma_{\min}(\mathbf{J}(\mathbf{s}))$ that depended on the sequence length T. We can also prove upper and lower bounds for $\sigma_{\max}(\mathbf{J}(\mathbf{s}))$, but these do not depend on the sequence length.

Assuming condition (21), an upper bound on $\sigma_{\text{max}}(\mathbf{J})$ is straightforward to compute via the triangle inequality,

$$\sigma_{\max}(\mathbf{J}) = \|\mathbf{J}\|_2$$

$$= \|\mathbf{I} - \mathbf{N}\|_2$$

$$\leq 1 + \|\mathbf{N}\|_2.$$

Recalling the definition of N in (20), we observe that it is composed of $\{J_t\}$ along its lower block diagonal, and so we have

$$\|\mathbf{N}(\mathbf{s})\|_2 = \sup_t \|J_t(s_t)\|$$
$$\sup_{\mathbf{s} \in \mathbb{R}^{TD}} \|\mathbf{N}(\mathbf{s})\|_2 = \sup_{s \in \mathbb{R}^D} \|J(s)\|$$

Elaborating, for a particular choice of trajectory $\mathbf{s} \in \mathbb{R}^{TD}$, $\|\mathbf{N}(\mathbf{s})\|_2$ is controlled by the maximum spectral norm of the Jacobians $J_t(s_t)$ along this trajectory. Analogously, $\sup_{\mathbf{s} \in \mathbb{R}^{TD}} \|\mathbf{N}(\mathbf{s})\|_2$ —i.e., the supremum of the spectral norm of $\mathbf{N}(\mathbf{s})$ over all possible trajectories $\mathbf{s} \in \mathbb{R}^{TD}$, i.e. the optimization space—is upper bounded by $\sup_{s \in \mathbb{R}^D} \|J(s)\|_2$, i.e. the supremum of the spectral norm of the system Jacobians over the state space \mathbb{R}^D .

Thus, it follows that

$$\sigma_{\max}(\mathbf{J}) \le 1 + \rho. \tag{22}$$

Importantly, the upper bound on $\sigma_{\max}(\mathbf{J})$ does not scale with the sequence length T.

To obtain the lower bound on $\sigma_{\max}(\mathbf{J})$, we notice that it has all ones along its main diagonal, and so simply by using the unit vector \mathbf{e}_1 , we obtain

$$\mathbf{e}_1^{\mathsf{T}} \mathbf{J} \mathbf{e}_1 = 1 \le \sigma_{\max}(\mathbf{J}). \tag{23}$$

Condition number of J Note that the condition number κ of a matrix is defined as the ratio of its maximum and minimum singular values, i.e.

$$\kappa(\mathbf{J}) = \frac{\sigma_{\max}(\mathbf{J})}{\sigma_{\min}(\mathbf{J})}.$$

However, because our bounds in eq. (22) and eq. (23) on $\sigma_{\max}(\mathbf{J})$ do not scale with the sequence length T, it follows that the scaling with T of an upper bound on $\kappa(\mathbf{J})$ —the conditioning of the optimization problem—is controlled solely by the bounds on $\sigma_{\min}(\mathbf{J})$ that we provided in Theorem 2. The importance of studying how the conditioning scales with T stems from the fact that we would like to understand if there are regimes—particularly involving large sequence lengths and parallel computers—where parallel evaluation can be faster than sequential evaluation.

C.1 A Generalized Proof that the Largest Lyapunov Exponent Controls the PL Constant

Lower Singular Value Bound Recall the following sequence of observations.

$$\lambda_{\min}(\mathbf{J}\mathbf{J}^{\top}) = \sigma_{\min}^{2}(\mathbf{J}) = \frac{1}{\sigma_{\max}^{2}(\mathbf{J}^{-1})} = \frac{1}{\|\mathbf{J}^{-1}\|_{2}^{2}}$$

Thus, to lower bound the eigenvalues of \mathbf{JJ}^{\top} as desired, we can *upper bound* the spectral norm of \mathbf{J}^{-1} .

General Bound As discussed in the main text, the predictability of the nonlinear state space model is characterized by the products of its Jacobians along a trajectory. We will need to control how this product behaves. To reduce notational burden, we will drop the DEER iteration superscript i. In particular, we will assume that there exists a function $g_J : \mathbb{N}_0 \to \mathbb{R}$ such that

$$\left\| J_{k-1} J_{k-2} \cdots J_i \right\|_{\xi} \le g_J(k-i)$$

holds for all products $J_{k-1}\cdots J_i$ with k>i, where $\|\cdot\|_{\xi}$ is the matrix operator norm induced by the vector norm $\|\cdot\|_{\xi}$. Intuitively, the function g_J measures the stability of the nonlinear state space model. For example, suppose the model is contracting with rate $\rho<1$. Then the product of Jacobians exponentially decreases, which we can write as

$$g_J(j) = a \rho^j$$
,

for some $a \ge 1$. The larger the value of a, the larger the potential "overshoot", before exponential shrinkage begins.

Lemma 1. Let $\|\cdot\|_{\xi}$ be the matrix operator norm induced by the vector norm $\|\cdot\|_{\xi}$. Suppose there is a function $g_J: \mathbb{N}_0 \to \mathbb{R}$ such that

$$\|J_{k-1}J_{k-2}\cdots J_i\|_{\varepsilon} \le g_J(k-i)$$

holds for all products $J_{k-1} \cdots J_i$ with k > i. Define

$$G_J(T) = \sum_{0 < j < T} g_J(j).$$

Then

$$\|\mathbf{J}^{-1}\|_{\xi} \leq G_J(T).$$

Proof. Let $y = J^{-1}x$. By backward substitution for the blockwise entries of $J^{-1}x$, we have

$$y_k = \sum_{i \in [k]} (J_{k-1}J_{k-2}\cdots J_i) x_i.$$

Omitting the subscript ξ in the norms for brevity and applying the triangle inequality and the induced-norm property,

$$\|\mathbf{y}\| \le \sum_{k \in [T]} \|y_k\| \le \sum_{k \in [T]} \sum_{i \in [k]} \|J_{k-1} \cdots J_i\| \|x_i\|.$$

By assumption, $||J_{k-1}\cdots J_i|| \leq g_J(k-i)$. Hence,

$$\|\mathbf{y}\| \le \sum_{k \in [T]} \sum_{i \in [k]} g_J(k-i) \|x_i\| = \sum_{i \in [T]} \|x_i\| \sum_{k=i}^T g_J(k-i) = \sum_{i \in [T]} \|x_i\| G_J(T-i+1).$$

Since $G_J(t)$ is nondecreasing in t, the largest multiplier in these sums is $G_J(T)$. In the worst case, $\|\mathbf{x}\| = \|x_1\|$. Thus,

$$\|\mathbf{J}^{-1}\| \le \frac{\|\mathbf{J}^{-1}\mathbf{x}\|}{\|\mathbf{x}\|} = \frac{\|\mathbf{y}\|}{\|\mathbf{x}\|} \le G_J(T).$$

This completes the proof.

Remark 1 (Contraction in The Identity Metric). Recall that a system is contracting in the identity metric when the system Jacobians have singular values less than one:

$$\forall i, \quad \|J_i\| \le \rho \quad \iff \quad J_i^\top J_i \le \rho^2 I$$

In this case, we can take

$$g_J(j) = \rho^j$$
.

Then, by Lemma 1,

$$\|\mathbf{J}^{-1}\| \le \sum_{j=0}^{T-1} \rho^j = \begin{cases} \frac{\rho^T - 1}{\rho - 1}, & \rho \ne 1, \\ T, & \rho = 1, \end{cases}$$
 (24)

where in the case $\rho = 1$, there are T summands and each term equals 1.

Remark 2 (Contraction in Time-Varying, State-Dependent Metrics). Recall that a system is contracting in metric $M_i = M(s_i, i)$ if the following linear matrix inequality is satisfied

$$\forall i \in [T-1], \quad J_i^\top M_{i+1} J_i \leq e^{2\lambda} M_i.$$

Equivalently, this condition can be written as a norm constraint

$$\forall i \in [T-1], \quad ||M_{i+1}^{1/2} J_i M_i^{-1/2}|| \leq \rho$$

Using these metrics, we define the block-diagonal, symmetric, positive-definite matrix

$$\mathbf{M} = \operatorname{diag}(M_1, M_2, \dots, M_T)$$

as well as the similarity transform of the residual function Jacobian, based on this matrix

$$\mathbf{J}_{\mathbf{M}}\coloneqq\mathbf{M}^{1/2}\,\mathbf{J}\,\mathbf{M}^{-1/2}.$$

Then the off-diagonal block entries of $J_{\mathbf{M}}$ are

$$M_{i+1}^{1/2} J_i M_i^{-1/2}$$
 for $i \in [T-1]$,

while its diagonal block entries are the identity matrix. If the off-diagonal blocks of J_M satisfy a product bound function $g_{\mathbf{J_M}}(j)$ as in Lemma 1, then $\mathbf{J_M}$ has norm bounded by $G_{\mathbf{J_M}}(T)$. Hence,

$$\begin{split} \|\mathbf{J}^{-1}\| &= \left\| \mathbf{M}^{-1/2} \ \mathbf{M}^{1/2} \mathbf{J}^{-1} \mathbf{M}^{-1/2} \ \mathbf{M}^{1/2} \right\| \\ &\leq \|\mathbf{M}^{-1/2}\| \ \left\| \mathbf{M}^{1/2} \mathbf{J}^{-1} \mathbf{M}^{-1/2} \right\| \ \|\mathbf{M}^{1/2}\| \\ &= \|\mathbf{M}^{-1/2}\| \ \|\mathbf{J}_{\mathbf{M}}^{-1}\| \ \|\mathbf{M}^{1/2}\| \\ &\leq \|\mathbf{M}^{-1/2}\| \ G_{\mathbf{J}_{\mathbf{M}}}(T) \ \|\mathbf{M}^{1/2}\| \\ &= \kappa_M \ G_{\mathbf{J}_{\mathbf{M}}}(T), \end{split}$$

where

$$\kappa_M \coloneqq \sqrt{\frac{\lambda_{\max}(\mathbf{M})}{\lambda_{\min}(\mathbf{M})}}.$$

In this case, we may again take $g_{J_M}(j)=\rho^j$, and we obtain the bound

$$\|\mathbf{J}^{-1}\| \le \kappa_M \sum_{0 \le j < T} \rho^j = \begin{cases} \kappa_M \frac{\rho^T - 1}{\rho - 1}, & \rho \ne 1, \\ \kappa_M T, & \rho = 1, \end{cases}$$

Remark 3 (Contraction After Burn-In). Suppose that

$$g_J(j) \leq ae^{-\lambda j}$$

where $a \ge 1$ and measures the degree of "overshoot" the system can undergo before eventually converging, and $\lambda > 0$. In particular, assume for concreteness that

$$||J_t|| < 1$$

Then, the product of two Jacobians can grow, if $a > e^{\lambda}$, since

$$||J_{t+1} J_t|| \le ae^{-\lambda}$$
.

In general, the product of Jacobians can transiently grow (i.e., overshoot) for

$$k_{\text{overshoot}} = \frac{1}{\lambda} \log a$$

time steps, at which point the product of $k > k_{\text{overshoot}}$ Jacobians will remain less than 1, and will in fact decay to zero exponentially with rate λ .

In this case, by Lemma 1:

$$\|\mathbf{J}^{-1}\| \le a \sum_{j=0}^{T-1} e^{-\lambda j} = a \frac{e^{-\lambda T} - 1}{e^{-\lambda} - 1}.$$

D DEER Merit Function Inherits Lipschitzness of Dynamics Jacobians

This section provides a proof of main text Theorem 3.

Theorem (Theorem 3). If the dynamics of the underlying nonlinear state space model have L-Lipschitz Jacobians, i.e.,

$$\forall t > 1, \quad s, s' \in \mathbb{R}^D : \quad ||J_t(s) - J_t(s')|| \le L||s - s'||,$$

then the residual function Jacobian J is also L-Lipschitz, with the same L.

Proof. By assumption, for each t,

$$\forall s, s' \in \mathbb{R}^D : \|J_t(s_t) - J_t(s_t')\|_2 \le L \|s_t - s_t'\|_2.$$

Define $D_t := J_t(s'_t) - J_t(s_t)$ and

$$\mathbf{D} := \mathbf{J}(\mathbf{s}') - \mathbf{J}(\mathbf{s}).$$

Since **D** places the blocks D_t along one subdiagonal, we have

$$\|\mathbf{D}\|_2 = \max_t \|D_t\|_2.$$

But each block D_t satisfies the Lipschitz bound

$$||D_t||_2 \leq L ||s_t' - s_t||_2,$$

so

$$\|\mathbf{D}\|_{2} = \max_{t} \|D_{t}\|_{2} \leq L \max_{t} \|s'_{t} - s_{t}\|_{2} \leq L \|\mathbf{s}' - \mathbf{s}\|_{2}.$$

Hence, it follows that

$$\|\mathbf{J}(\mathbf{s}') - \mathbf{J}(\mathbf{s})\|_2 = \|\mathbf{D}\|_2 \le L \|\mathbf{s}' - \mathbf{s}\|_2.$$

 \Box

Thus **J** is *L*-Lipschitz.

E DEER Always Converges Linearly

This section provides a proof of Theorem 4.

While proofs of global convergence are challenging in general for GN, DEER is highly structured, and this can be exploited to provide a global proof of convergence. In particular, we will exploit the *hierarchical* nature of DEER, which is reflected in the fact that $\bf J$ and $\bf J^{-1}$ are lower block-triangular.

Theorem (Theorem 4). Let the DEER (Gauss–Newton) updates be given by eq. (3), and let $\mathbf{s}^{(i)}$ denote the *i*-th iterate. Let $\mathbf{e}^{(i)} := \mathbf{s}^{(i)} - \mathbf{s}^*$ denote the error at iteration *i*. Then the error converges to zero at a linear rate:

$$\|\mathbf{e}^{(i)}\|_{2} \le \chi_{w} \beta^{i} \|\mathbf{e}^{(0)}\|_{2},$$

for some constant $\chi_w \geq 1$ independent of i, and a convergence rate $0 < \beta < 1$.

Proof. Our general strategy for deriving DEER convergence bounds will be to fix some weighted norm $\|\cdot\|_W \coloneqq \|\mathbf{W}^{1/2}\cdot\mathbf{W}^{-1/2}\|_2$ such that each DEER step is a contraction, with contraction factor $\beta \in [0,1)$. This will imply that the DEER error iterates decay to zero with linear rate, as

$$\|\mathbf{e}^{(i)}\|_W \le \beta^i \|\mathbf{e}^{(0)}\|_W.$$

To convert this bound back to standard Euclidean space, we incur an additional multiplicative factor that depends on the conditioning of W:

$$\|\mathbf{e}^{(i)}\|_{2} \le \chi_{w} \,\beta^{i} \|\mathbf{e}^{(0)}\|_{2}, \quad \text{where} \quad \chi_{w} \coloneqq \sqrt{\frac{\lambda_{\max}(\mathbf{W})}{\lambda_{\min}(\mathbf{W})}}.$$
 (25)

DEER as a Contraction Mapping Recall that the DEER (Gauss-Newton) updates are given by

$$\mathbf{s}^{(i+1)} = \mathbf{s}^{(i)} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{r}(\mathbf{s}^{(i)})$$

Recalling that $\mathbf{r}(\mathbf{s}^*) = \mathbf{0}$ and subtracting the fixed point \mathbf{s}^* from both sides, we have that

$$\mathbf{e}^{(i+1)} \ = \ \mathbf{e}^{(i)} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{r}^{(i)} + \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{r}(\mathbf{s}^*) \ = \ \mathbf{e}^{(i)} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\left(\mathbf{r}(\mathbf{s}^{(i)}) - \mathbf{r}(\mathbf{s}^*)\right).$$

This equation can be written using the mean value theorem as

$$\mathbf{e}^{(i+1)} = \left(\mathbf{I} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{B}^{(i)}\right)\mathbf{e}^{(i)} \quad \text{where} \quad \mathbf{B}^{(i)} \coloneqq \int_0^1 \mathbf{J}(\mathbf{s}^* + \tau \mathbf{e}^{(i)}) d\tau$$

From this, we can conclude that the DEER iterates will converge (i.e., the error shrinks to zero) if

$$\|\mathbf{I} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{B}^{(i)}\|_{W} = \|\mathbf{J}^{-1}(\mathbf{s}^{(i)})\left(\mathbf{J}(\mathbf{s}^{(i)}) - \mathbf{B}^{(i)}\right)\|_{W} \le \beta < 1.$$
 (26)

Constructing the Weighted Norm We will choose a diagonal weighted norm, given by

$$\mathbf{W} := \operatorname{Diag}(I_D, \ w^2 I_D, \ \dots, \ w^{2T} I_D) \in \mathbb{R}^{TD \times TD}, \qquad w > 0.$$
 (27)

Under the norm induced by (27) we have

$$\|\mathbf{J}(\mathbf{s}^{(i)}) - \mathbf{B}^{(i)}\|_{W} \le 2w\rho,\tag{28}$$

$$\|\mathbf{J}^{-1}(\mathbf{s}^{(i)})\|_{W} \le a \frac{1 - (we^{\lambda})^{T}}{1 - we^{\lambda}},\tag{29}$$

where ρ upper bounds $||J||_2$ over all states in the DEER optimization trajectory.

Multiplying (28) and (29) yields

$$\|\mathbf{J}^{-1}(\mathbf{s}^{(i)})\|_{W} \|\mathbf{J}(\mathbf{s}^{(i)}) - \mathbf{B}^{(i)}\|_{W} \le 2aw\rho \frac{1 - (we^{\lambda})^{T}}{1 - we^{\lambda}}.$$
 (30)

To ensure the right-hand side of (30) does not exceed a prescribed $\beta \in [0, 1)$, choose

$$w = \frac{\beta}{2oa + \beta e^{\lambda}}. (31)$$

With this choice,

$$we^{\lambda} < 1, \quad \text{and} \quad \frac{2aw\rho}{1 - we^{\lambda}} = \beta,$$
 (32)

so the geometric series in (29) is convergent and the bound in (30) holds for all T, because

$$\|\mathbf{J}^{-1}(\mathbf{s}^{(i)})\|_{W} \|\mathbf{J}(\mathbf{s}^{(i)}) - \mathbf{B}^{(i)}\|_{W} \le 2aw\rho \frac{1 - (we^{\lambda})^{T}}{1 - we^{\lambda}} = \beta (1 - (we^{\lambda})^{T}) \le \beta.$$

This shows that we can always pick a weighted norm so that DEER converges with linear rate *in that norm*. Converting back into the standard Euclidean norm using (25) and substituting in the condition number of **W** one finds that

$$\|\mathbf{e}^{(i)}\|_{2} \le \left(\frac{2\rho a + \beta e^{\lambda}}{\beta}\right)^{T} \beta^{i} \|\mathbf{e}^{(0)}\|_{2}.$$
 (33)

Thus, the DEER error converges with linear rate towards zero.

Remark 4. The multiplicative overshoot factor arising from the conditioning of W grows exponentially in the sequence length T, leading potentially to long convergence times. Indeed, a quick calculation shows that the number of steps needed to bring the DEER error to ϵ is upper bounded as O(T) because of this multiplicative constant.

Remark 5. One can ask under what conditions choosing w = 1 in (31) is possible, which eliminates the overshoot. We will address this in more detail in the next section. To provide a simple result here, we can assume that the system is contracting at every time step so that

$$\rho = e^{\lambda}$$

and a = 1. Then we have that

$$1 = \frac{\beta}{2\rho a + \beta e^{\lambda}} = \frac{\beta}{2e^{\lambda} + \beta e^{\lambda}}.$$

Solving for λ , we have that if

$$\lambda \le \log\left(\frac{\beta}{2+\beta}\right) < -\log(3),$$

then w can be chosen to be equal to one, meaning the DEER converges globally with rate β and no overshoot.

F DEER Converges Globally with Small Overshoot for Sufficiently Strongly Contracting Systems

In this section we show that DEER converges globally to the optimum s^* when the nonlinear state space model (1) is sufficiently strongly contracting. To do so, we first briefly recall the assumptions of Lemma 1. Let $\|\cdot\|_{\xi}$ be the matrix operator norm induced by the vector norm $\|\cdot\|_{\xi}$. Suppose there is a function $g_J: \mathbb{N}_0 \to \mathbb{R}$ such that

$$\|J_{k-1}J_{k-2}\cdots J_i\|_{\xi} \le g_J(k-i)$$

holds for all products $J_{k-1} \cdots J_i$ with k > i. Define

$$G_J(T) = \sum_{0 \le j \le T} g_J(j).$$

Then

$$\|\mathbf{J}^{-1}\|_{\mathcal{E}} \leq G_J(T).$$

For example, if there is no structure which can be exploited in the products of Jacobians J_t , we may consider the "one-step" growth/decay factor

$$\forall t, \quad ||J_t|| \le e^{\lambda},$$

which yields

$$g_J(j) = e^{\lambda j}$$
 \Longrightarrow $G_J(T) = \sum_{0 \le j < T} g_J(j) = \frac{1 - e^{\lambda T}}{1 - e^{\lambda}}.$

Theorem. DEER exhibits linear, global convergence to the optimum \mathbf{s}^* with rate $\beta \in [0,1)$ in the matrix operator norm $\|\cdot\|_{\xi}$ if

$$2g_J(1)G_J(T) \le \beta$$

Proof. Recall that the DEER (Gauss-Newton) updates are given by

$$\mathbf{s}^{(i+1)} = \mathbf{s}^{(i)} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{r}(\mathbf{s}^{(i)})$$

Define the error at DEER iteration (i) as $\mathbf{e}^{(i)} = \mathbf{s}^{(i)} - \mathbf{s}^*$. Recalling that $\mathbf{r}(\mathbf{s}^*) = \mathbf{0}$ and subtracting the fixed point \mathbf{s}^* from both sides, we have that

$$\mathbf{e}^{(i+1)} \ = \ \mathbf{e}^{(i)} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{r}^{(i)} + \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{r}(\mathbf{s}^*) \ = \ \mathbf{e}^{(i)} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\left(\mathbf{r}(\mathbf{s}^{(i)}) - \mathbf{r}(\mathbf{s}^*)\right).$$

This equation can be written in terms of the mean value theorem as

$$\mathbf{e}^{(i+1)} = \left(\mathbf{I} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{B}^{(i)}\right)\mathbf{e}^{(i)} \quad \text{where} \quad \mathbf{B}^{(i)} \coloneqq \int_0^1 \mathbf{J}(\mathbf{s}^* + \tau \mathbf{e}^{(i)}) d\tau$$

This follows from the identity:

$$\mathbf{r}(\mathbf{s}^{(i)}) - \mathbf{r}(\mathbf{s}^*) \coloneqq \int_0^1 \mathbf{J}(\tau \mathbf{s}^{(i)} + (1 - \tau)\mathbf{s}^*) d\tau (\mathbf{s} - \mathbf{s}^*) = \left(\int_0^1 \mathbf{J}(\mathbf{s}^* + \tau \mathbf{e}^{(i)})\right) d\tau\right) \mathbf{e}^{(i)}$$

This identity can be proven by starting from the fundamental theorem of calculus, by letting

$$\mathbf{s}(\tau) = \mathbf{s}^* + \tau \mathbf{e}^{(i)} = \mathbf{s}^* + \tau (\mathbf{s}^{(i)} - \mathbf{s}^*), \quad \tau \in [0, 1],$$

which defines a straight-line path from s^* to $s^{(i)}$. The fundamental theorem of calculus then says that

$$\mathbf{r}(\mathbf{s}^{(i)}) - \mathbf{r}(\mathbf{s}^*) = \int_0^1 \frac{d}{d\tau} \mathbf{r}(\mathbf{s}(\tau)) d\tau.$$

Applying the chain rule inside the integral gives the result, because

$$\frac{d}{d\tau}\mathbf{r}(\mathbf{s}(\tau)) = \mathbf{J}(\mathbf{s}(\tau)) \cdot \frac{d}{d\tau}\mathbf{s}(\tau) = \mathbf{J}(\mathbf{s}^* + \tau\mathbf{e}^{(i)}) \cdot \mathbf{e}^{(i)}.$$

From this, we can conclude that the DEER iterates will converge (i.e., the error shrinks to zero) if

$$\|\mathbf{I} - \mathbf{J}^{-1}(\mathbf{s}^{(i)})\mathbf{B}^{(i)}\|_{\xi} = \|\mathbf{J}^{-1}(\mathbf{s}^{(i)})\left(\mathbf{J}(\mathbf{s}^{(i)}) - \mathbf{B}^{(i)}\right)\|_{\xi} \le \beta < 1.$$
 (34)

By Lemma 1 we have that

$$\|\mathbf{e}^{(i+1)}\|_{\xi} \le \|\mathbf{J}^{-1}(\mathbf{s}^{(i)})\|_{\xi} \|\mathbf{J}(\mathbf{s}^{(i)}) - \mathbf{B}^{(i)}\|_{\xi} \|\mathbf{e}^{(i)}\|_{\xi} \le 2g_J(1)G_J(T)\|\mathbf{e}^{(i)}\|_{\xi}$$

Thus, if there exists some $\beta \in [0, 1)$ such that

$$2g_J(1)G_J(T) \leq \beta$$
,

then the DEER error converges globally to zero in the weighted norm:

$$\|\mathbf{e}^{(i)}\|_{\xi} \le \beta^i \|\mathbf{e}^{(0)}\|_{\xi}.$$

Corollary. Suppose the state space model is contracting in constant metric M, i.e.,

$$||M^{1/2}J_tM^{-1/2}|| = ||J_t||_M \le e^{\lambda} < 1.$$

If e^{λ} is sufficiently small, in particular if

$$e^{\lambda} \le \frac{\beta}{2+\beta} < \frac{1}{3},$$

then the DEER errors converge to zero with rate β .

 ${\it Proof.}$ Suppose the state space model is contracting in constant metric M, so that

$$||M^{1/2}J_tM^{-1/2}|| = ||J_t||_M \le e^{\lambda} < 1$$

for all t. Then, by Lemma 1 we have that

$$\|\mathbf{e}^{(i+1)}\|_{M} \le \|\mathbf{J}^{-1}\|_{M} \|\mathbf{J}(\mathbf{s}^{(i)}) - \mathbf{B}^{(i)}\|_{M} \|\mathbf{e}^{(i)}\|_{M} \le \left(\frac{1 - e^{\lambda T}}{1 - e^{\lambda}} \cdot 2e^{\lambda}\right) \|\mathbf{e}^{(i)}\|_{M}$$

Thus, in order to achieve linear convergence of the DEER iterates with rate $\beta \in [0, 1)$,

$$||\mathbf{e}^{(i+1)}||_{M} \le \beta ||\mathbf{e}^{(i)}||_{M} \implies ||\mathbf{e}^{(i)}||_{M} \le \beta^{i} ||\mathbf{e}^{(0)}||_{M},$$

we require that

$$2e^{\lambda} \cdot \frac{1 - e^{\lambda T}}{1 - e^{\lambda}} \le \beta < 1.$$

A simple sufficient condition for satisfying this inequality is

$$e^{\lambda} \le \frac{\beta}{2+\beta} < \frac{1}{3},$$

or,

$$\lambda < -\log(3)$$
.

Number of Steps to reach basin of quadratic convergence Let us assume that there exists $\beta \in [0,1)$ such that

$$e^{\lambda} \leq \frac{\beta}{2+\beta}$$

then the number of steps to reach the basin of quadratic convergence is upper bounded as

$$k_Q \leq \log\left(\frac{1}{\beta}\right) \cdot \log\left(\frac{e^{\lambda}||\mathbf{e}^{(0)}||L}{\mu}\right)$$

G Alternative Descent Techniques & Worst/Average Complexity

DEER uses the Gauss-Newton algorithm, which converges quadratically near the optimum but can be slow outside this basin. This motivates *inexact* GN methods that guarantee a certain loss decrease per step, such as line-search and trust-region techniques. These trade increased computation and possibly more iterations for faster convergence guarantees.

In practice, we found that plain GN reliably converged quickly to the global optimum in contracting systems, so such safeguards were unnecessary. Still, it is useful to analyze DEER's worst-case path to the quadratic basin.

Many inexact GN variants achieve global convergence from any starting point. These include stepsize schemes that approximate a continuous flow [76], trust-regions that bound update size (yielding ELK when applied to DEER [3]), and backtracking line search ensuring loss reduction at each step [19].

One can also use a simpler algorithm outside of the basin of quadratic convergence, and then switch to GN when needed. We will consider this latter option, and choose gradient descent as our simpler algorithm. Because the merit function is PL (see section 3.1), the number of steps required for gradient descent to reach the quadratic convergence region scales as:

$$k_Q \sim \frac{1}{\mu} \cdot \log \frac{||\mathbf{r}^{(0)}||}{\mu},\tag{35}$$

where $||\mathbf{r}^{(0)}||$ is the residual at initialization. For unpredictable systems, μ may shrink arbitrarily with increasing sequence length T, leading to unbounded growth in the number of optimization steps k_Q . By contrast, for predictable systems, μ remains bounded, implying that the number of optimization steps does not increase with sequence length. Since the cost of sequential evaluation always increases with T, DEER can, even in the worst case, compute the true rollout faster than sequential evaluation for predictable systems—especially for long sequences. Indeed, assuming the system is contracting with rate $e^{\lambda} < 1$, then the number of steps needed the reach the basin of quadratic convergence is $\mathcal{O}(\log ||\mathbf{r}^{(0)}||)$.

Thus, if the initial error grows polynomial in T, i.e., $||\mathbf{r}^{(0)}|| \propto T^p$, then this implies that the number of gradient descent steps needed to reach the basin of quadratic convergence is only $\mathcal{O}(\log T)$, and thus the total computational time is $\mathcal{O}((\log T)^2)$. In practice, for randomly initialized DEER, we observe p=1.

In practice, we observe that DEER converges much faster than the worst-case analysis (35) would suggest. In particular, we observe that DEER converges in roughly $\log \frac{1}{\mu}$, steps, even for unpredictable systems. This behavior can be explained with a simple "two-phase" model, wherein the DEER iterates move towards the basin of quadratic convergence at a rate which is independent of the PL-constant μ (see Appendix J.1).

H Proof of Size of Basin of Quadratic Convergence

This section provides a proof of Theorem 5:

Theorem (Theorem 5). Let μ denote the PL-constant of the merit function, which Theorem 2 relates to the LLE λ . Let L denote the Lipschitz constant of the Jacobian of the dynamics function J(s). Then, μ/L lower bounds the radius of the basin of quadratic convergence of DEER; that is, if

$$||\mathbf{r}(\mathbf{s}^{(i)})||_2 \le \frac{\mu}{L},$$

then $\mathbf{s}^{(i)}$ is inside the basin of quadratic convergence. In terms of the LLE λ , it follows that if

$$||\mathbf{r}(\mathbf{s}^{(i)})||_2 \le \frac{1}{a^2 L} \cdot \left(\frac{e^{\lambda} - 1}{e^{\lambda T} - 1}\right)^2,$$

then $\mathbf{s}^{(i)}$ is inside the basin of quadratic convergence.

Suppose we are at a point $\mathbf{s}^{(i)} \in \mathbb{R}^{TD}$ (i.e. DEER iterate i), and we want to get to $\mathbf{s}^{(i+1)}$. The change in the trajectory obtained from eq. (3) is,

$$\Delta \mathbf{s}^{(i)} := -\mathbf{J}(\mathbf{s}^{(i)})^{-1}\mathbf{r}(\mathbf{s}^{(i)})$$

(where the iteration number will hopefully be clear from context). The merit function is $\mathcal{L}(\mathbf{s}) = \frac{1}{2} \|\mathbf{r}(\mathbf{s})\|_2^2$, so if we can get some control over $\|\mathbf{r}(\mathbf{s}^{(i)})\|_2$, we will be well on our way to proving a quadratic rate of convergence.

First, leveraging the form of the Gauss-Newton update, we can simply "add zero" to write

$$\begin{aligned} \mathbf{r}(\mathbf{s}^{(i+1)}) &= \mathbf{r}(\mathbf{s}^{(i)} + \Delta \mathbf{s}^{(i)}) \\ &= \mathbf{r}(\mathbf{s}^{(i)} + \Delta \mathbf{s}^{(i)}) - \mathbf{r}(\mathbf{s}^{(i)}) - \mathbf{J}(\mathbf{s}^{(i)}) \Delta \mathbf{s}^{(i)} \end{aligned}$$

Next, we can write the difference $\mathbf{r}(\mathbf{s}^{(i)} + \Delta \mathbf{s}^{(i)}) - \mathbf{r}(\mathbf{s}^{(i)})$ as the integral of the Jacobian, i.e.

$$\mathbf{r}(\mathbf{s}^{(i)} + \Delta \mathbf{s}^{(i)}) - \mathbf{r}(\mathbf{s}^{(i)}) = \int_0^1 \mathbf{J} \left(\mathbf{s}^{(i)} + \tau \Delta \mathbf{s}^{(i)} \right) \Delta \mathbf{s}^{(i)} d\tau.$$

Therefore,

$$\mathbf{r}(\mathbf{s}^{(i+1)}) = \int_0^1 \left(\mathbf{J} \left(\mathbf{s}^{(i)} + \tau \Delta \mathbf{s}^{(i)} \right) - \mathbf{J}(\mathbf{s}^{(i)}) \right) \Delta \mathbf{s}^{(i)} d\tau$$

Taking ℓ_2 -norms and using the triangle inequality, it follows that

$$\|\mathbf{r}(\mathbf{s}^{(i+1)})\|_2 \le \int_0^1 \left\| \left(\mathbf{J} \left(\mathbf{s}^{(i)} + \tau \Delta \mathbf{s}^{(i)} \right) - \mathbf{J}(\mathbf{s}^{(i)}) \right) \Delta \mathbf{s}^{(i)} \right\|_2 d\tau.$$

Now, if we assume that J is L-Lipschitz and use the definition of spectral norm, it follows that

$$\left\| \left(\mathbf{J} \left(\mathbf{s}^{(i)} + \tau \Delta \mathbf{s}^{(i)} \right) - \mathbf{J} (\mathbf{s}^{(i)}) \right) \Delta \mathbf{s}^{(i)} \right\|_{2} \le \tau L \|\Delta \mathbf{s}^{(i)}\|_{2}^{2},$$

and so taking the integral we obtain

$$\begin{split} \|\mathbf{r}(\mathbf{s}^{(i+1)})\|_2 &\leq \frac{L}{2} \|\Delta \mathbf{s}^{(i)}\|_2^2 \\ &= \frac{L}{2} \mathbf{r}(\mathbf{s}^{(i)})^{\top} \mathbf{J}(\mathbf{s}^{(i)})^{-\top} \mathbf{J}(\mathbf{s}^{(i)})^{-1} \mathbf{r}(\mathbf{s}^{(i)}). \end{split}$$

By definition, $\sqrt{\mu}$ is a lower bound on all singular values of $\mathbf{J}(\mathbf{s}(i))$, for all i. Therefore, $\|\mathbf{J}(\mathbf{s}^{(i)})^{-1}\|_2 \leq 1/\sqrt{\mu}$ for all i, and it follows that

$$\|\mathbf{r}(\mathbf{s}^{(i+1)})\|_{2} \le \frac{L}{2\mu} \|\mathbf{r}(\mathbf{s}^{(i)})\|_{2}^{2},$$
 (36)

which is the direct analogy of Boyd and Vandenberghe [28, 9.33]. To reiterate, here L is the Lipschitz constant of \mathbf{J} , while $\mu := \inf_{i \in \mathbb{N}} \sigma_{\min}^2 \left(\mathbf{J}(\mathbf{s}^{(i)}) \right)$.

While this is a quadratic convergence result for GN, this result is not useful unless $\mathbf{r}(\mathbf{s}^{(i+1)})\|_2 \le \|\mathbf{r}(\mathbf{s}^{(i)})\|_2$ (i.e. would backtracking line search accept this update). However, if we have $\|\mathbf{r}(\mathbf{s}^{(i)})\|_2 \le \frac{\mu}{L}$, then every step guarantees a reduction in \mathbf{r} because in this case

$$\|\mathbf{r}(\mathbf{s}^{(i+1)})\|_2 \le \frac{1}{2} \|\mathbf{r}(\mathbf{s}^{(i)})\|_2.$$

Therefore, we have $\|\mathbf{r}(\mathbf{s}^{(j)})\|_2 \leq \frac{\mu}{L}$ for all j > i. Thus, we have related the size of the basin of quadratic convergence of GN on the DEER objective to the properties of \mathbf{J} . Note that with linear dynamics, each J_t is constant in s, and so each J_t is 0-Lipschitz. Thus, the basin of quadratic convergence becomes infinite. Intuitively, if J_t doesn't change too quickly with s, then DEER becomes a more and more potent method.

I Parameterizing Nonlinear SSMs to be contractive

In this section, we highlight a practical strategy for speeding up the training of nonlinear state space models (SSMs) based on our theoretical findings.

Our results indicate that nonlinear SSMs with negative largest Lyapunov exponents (LLEs) are efficiently parallelizable. To exploit this during training, one must ensure that the model maintains negative LLEs throughout optimization. One straightforward and effective method to achieve this is by design, through *parameterization*. In particular, by introducing an auxiliary variable to enforce the desired constraint (in this case, negative LLE), and then performing unconstrained optimization on this variable.

This strategy is particularly well-suited to neural network-based SSMs. For example, consider the scalar nonlinear SSM:

$$x_t = \tanh(wx_{t-1} + u_t)$$

To guarantee negative LLE, it suffices to ensure that the Jacobian norm is strictly less than one:

$$|J_t| = |w \cdot \operatorname{sech}^2(wx_{t-1} + u_t)| \le |w|$$

Thus, enforcing |w| < 1 is sufficient. This can be achieved by reparameterizing $w = \tanh(b)$, where b is a trainable, unconstrained auxiliary variable. This guarantees that $w \in (-1,1)$ for all finite b, ensuring contractivity and, hence, negative LLE. A similar argument holds in the multivariate case, using the spectral norm.

J Experimental Details and Discussion

All of our experiments use FP64 to, as much as possible, focus on algorithmic factors controlling the rate of convergence of DEER, as opposed to numerical factors. As noted in [3], DEER can be prone to numerical overflow in lower precision. While such numerical overflow can be overcome by resetting NaNs to their initialized value, such an approach resets the optimization and leads to rates that are slower than what Gauss-Newton would achieve in infinite precision (exact values in \mathbb{R}).

J.1 Deriving the Empirical Scaling of DEER

In our experiments, we observed that DEER typically converges in $\mathcal{O}(\log(1/\mu))$ steps (see, for example, Figure 2). To understand this scaling behavior, we propose a simple two-phase model of DEER convergence. In the first phase, the iterates approach the basin of quadratic convergence at a linear rate, as guaranteed by Theorem 4. In the second phase, rapid quadratic convergence occurs, typically requiring only one or two steps to reach the true solution (up to floating point precision).

Although Theorem 4 shows that, in unpredictable systems, the overshoot factor may be exponentially large in the sequence length T, this reflects a worst-case analysis. In practice, DEER behaves as though the overshoot factor is negligible. To formalize this observation, recall from Theorem 4 that the residuals satisfy the linear convergence bound

$$\|\mathbf{r}_i\| \leq \chi_w \beta^i \|\mathbf{r}_0\|,$$

for some $\beta \in [0,1)$ and $\chi_w \ge 1$, where β is always independent of T. In our two-phase model, we assume that χ_w is also independent of T, even when the largest Lyapunov exponent λ is positive.

We now upper-bound the number of steps k required to enter the basin of quadratic convergence, whose size is μ/L (as given by (12)). Solving

$$\frac{\mu}{L} = \chi_w \beta^k \|\mathbf{r}_0\| \implies k = \frac{1}{\log \beta} \log \left(\frac{\chi_w L \|\mathbf{r}_0\|}{\mu}\right), \tag{37}$$

we recover the empirically observed logarithmic scaling.

J.2 Details and Discussion for mean-field RNN experiment

We rolled out trajectories from a mean-field RNN with step size 1 for 20 different random seeds. The dynamics equations follow the form

$$s_{t+1} = W \tanh(s_t) + u_t,$$

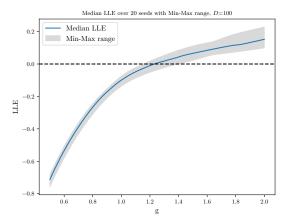


Figure 4: Robust relationship in mean field RNN between variance parameter g and LLE of the system. For 20 seeds, we observe a robust and monotonic relationship between the scalar parameter g and the LLE of the resulting mean-field RNN. The plot above is made for 50 different values of g from 0.5 to 2.0 (linearly spaced).

for mild sinusoidal inputs u_t . We have $s_t \in \mathbb{R}^D$, where in our experiments D=100. Note that because of the placement of the saturating nonlinearity, here s_t represents current, not voltage.

In the design of the weight matrix W, we follow Engelken et al. [29]. In particular, we draw each entry $W_{ij} \stackrel{\text{iid}}{\sim} \mathcal{N}(0,g^2/D)$, where g is a scalar parameter. We then set $W_{ii} = 0$ for all i (no self-coupling of the neurons). A key point of Engelken et al. [29] is that by scaling the single parameter g, the resulting RNN goes from predictable to chaotic behavior. While Engelken et al. [29] computes the full Lyapunov spectrum in the limit $D \to \infty$, for finite D we can compute a very accurate numerical approximation to the LLE (cf. Appendix J.6). In Figure 4, we verify numerically that there is a monotonic relationship between g and the LLE of the resulting system, and that the min-max range for 20 seeds is small. Accordingly, when making Figure 2 (Center), we use the monotonic relationship between g and the LLE from Figure 4 to map the average number of DEER steps (over 20 different seeds) needed for convergence for different values of g to the appropriate value of the LLE. We use 50 values of g from 9 to 9999 (log spaced) to make Figure 2 (Center). We highlight g 1000 in Figure 2 (Right).

For the purposes of Figure 2, we define

$$\tilde{\mu} := \left(\frac{e^{\lambda} - 1}{e^{\lambda T} - 1}\right)^2,$$

i.e. the lower bound on μ from Theorem 2, with a=1.

In Figure 4, we observe that around g=1.2, the RNNs have LLE around 0, which is the threshold between predictability and chaos. Working with chaotic dynamics in finite precision for long time series led to some interesting difficulties.

First, as discussed in Gonzalez et al. [3], DEER can experience numerical overflow when deployed on unstable systems. While we reset to the initialization (in this experiment we initialized $\mathbf{s}_{1:T}$ with iid draws from $\mathcal{U}[0,1]$), doing so slows convergence. Thus, many of our runs for $\lambda>0$ and large T take the maximum number of DEER iteration we allow (we do not allow more than T iterations, as this is the theoretical upper bound for number of DEER iterations before convergence, cf. Proposition 1 of [3]), which helps to explain the slight increase in red space for experiment (center plot of Figure 2) vs. theory (left plot of Figure 2). Note, however, that for T=1000 (the sequence length shown in the right plot), there is no numerical overflow for the DEER trajectories for any of the 20 random seeds or 50 values of q tried.

Second, we observe that for many values of λ in the chaotic range, even after the maximum number of DEER steps (T) was taken, there was still a large discrepancy between the true sequential rollout and the converged DEER iteration, even though the converged DEER iteration had numerically zero merit function. For example, in Figure 2 (Right), there are a series of points in the top right of the graph that all sit on the line T=1000, and while they have numerically zero merit function

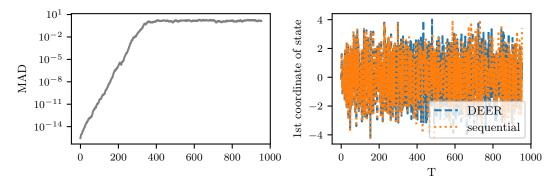


Figure 5: Chaotic behavior means numerically zero merit function can still be far from sequential trajectory. For g=1.85 and T=1000, we show the final DEER vs sequential trajectory. The DEER trajectory has merit function (2) numerically equal to zero. However: (Left) the mean absolute deviation (MAD) at each time point t between the final DEER iteration $\mathbf{s}_t^{(T)}$ and the sequential rollout \mathbf{s}_t^* grows exponentially. This exponential growth of error is a signature of chaos: compare, for example, with Figure 9.3.5 of Strogatz [10]. The saturation of the error eventually occurs because of the saturating nonlinearity present in the RNN. (Right) We visualize the first coordinate of both the final DEER iteration and the sequential trajectory, showing that while they initially coincide, they diverge around t=300.

value, the converged DEER trajectories are quite different from the true sequential trajectories. The reason for this behavior precisely stems from the fact that for large values of g (equivalently λ), these mean-field RNNs are chaotic. Even working in FP64, if slight numerical errors are introduced at any time point in the sequence (say t=1), then over the sequence length we can observe exponential divergence from the true trajectories, as illustrated in Figure 5. This experimental observation is complemented by our discussion of why unpredictable systems have excessively flat merit functions in Section 3.2, and provides a numerical perspective on why ill-conditioned landscapes are hard to optimize: if the landscape is extremely flat, many potential trajectories $\mathbf{s}_{1:T}$ can have numerically zero merit function, even in extremely high precision.

J.3 Additional experiment for the mean-field RNN: other optimizers and wallclock time

In this section, we provide further experiments in the setting of the mean-field RNN (Figure 2). In particular, we showcase the generality of our theory beyond DEER (Gauss-Newton optimization), and the practicality of our theory by reporting wallclock times. We consider the setting in the right most panel of Figure 2, where we evaluate a mean field RNN over a sequence length of length T=1000.

Quasi-Newton and Gradient Descent Instead of only using Gauss-Newton optimization (DEER) to parallelize the sequence length, we also consider other optimization algorithms (quasi-Newton and gradient descent) to showcase the generality of our theory.

We include a quasi-Newton algorithm proposed in Gonzalez et al. [3] called quasi-DEER. Quasi-DEER simply replaces the J_t defined in eq. (4) with $\mathrm{diag}(J_t)$, and so is also parallelizable over the sequence length with a parallel scan. Furthermore, we also include gradient descent on the merit function, which is embarrassingly parallel over the sequence length. In the top panel of Figure 6, we observe that the number of steps for gradient descent and quasi-DEER to converge also scales monotonically with the LLE, as we expect from Theorem 2. DEER (Gauss-Newton) converges in a small number of steps all the way up to the threshold between predictability and unpredictability ($\lambda=0$). Intuitively, the performance of the other optimizers degrades more quickly as unpredictability increases because quasi-Newton and gradient descent use less information about the curvature of the loss landscape.

Even though gradient descent was slower to converge in this setting, we only tried gradient descent with a fixed step size. An advantage of a first-order method like gradient descent over a second-order method like Gauss-Newton (DEER) is that the first-order method is embarrassingly parallel

(and so with sufficient parallel processors, the update runs in constant time), while DEER and quasi-DEER use parallel scans (and so the update runs in $O(\log T)$ time). Exploring accelerated first-order methods like Adam [77], or particularly Shampoo [78] or SOAP [79] (which are often preferred in recurrent settings like eq. (1))—or in general trying to remove the parallel scan—are therefore very interesting directions for future work.

We note that sequential evaluation of eq. (1) can also be thought of as block coordinate descent on the merit function $\mathcal{L}(\mathbf{s})$, where the block $s_t \in \mathbb{R}^D$ is optimized at optimization step (t). The optimization of each block is a convex problem: simply minimize $\|s_t - f(s_{t-1}^*)\|_2^2$, or equivalently set $s_t = f(s_{t-1}^*)$. As sequential evaluation will always take T steps to converge, we do not include it in the top panel of Figure 6.

Wallclock time In the bottom panel of Figure 6, we also report the wallclock times for these algorithms to run (our experiments are run on an H100 with 80 GB onboard memory). We observe that the run time of sequential evaluation (green) is effectively constant with respect to λ . We observe that in the predictable setting, DEER is an order of magnitude faster than sequential evaluation, while in the unpredictable regime, DEER is 1-2 orders of magnitude slower than sequential evaluation. This importance of using parallel evaluation only in predictable settings is a core practical takeaway from our theoretical contributions.

Further details We run the experiment in Figure 6 on a smaller scale than the experiment in Figure 2 (Right). In Figure 6, we consider 5 random seeds for 16 values of g equispaced between 0.5 and 2.0. Each wallclock time reported is the average of 5 runs for the same seed. We use a batch size of 1. While DEER (Gauss-Newton) and quasi-DEER effectively do not have a step size (they use a step size of 1 always). For each value of g, we ran gradient descent with the following set of step sizes α : 0.01, 0.1, 0.25, 0.5, 0.6, 0.7, 0.8, 0.9,and 1.0.For each value of g, we then pick the step size α that results in the fastest convergence of gradient descent. For the smallest value of

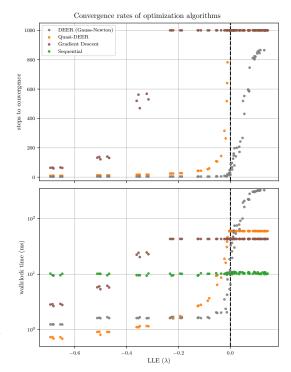


Figure 6: Convergence rates and wallclock time for many optimizers. We supplement the mean-field RNN experiment by also considering quasi-Newton and gradient descent methods (top), and recording wallclock time, including for sequential evaluation (bottom).

g=0.5, we use $\alpha=0.6$; for g=0.6, we use $\alpha=0.5$; and for all other values of g, we use $\alpha=0.25$. Future work may investigate more adaptive ways to tune the step size α , or to use a learning rate schedule.

We use a larger tolerance of $\mathcal{L}(\mathbf{s})/T \leq 10^{-4}$ to declare convergence than in the rest of the paper (where we use a tolerance of 10^{-10}) because gradient descent often did not converge to the same degree of numerical precision as sequential, quasi-DEER, or DEER. However, this is a per time-step average error on the order of 10^{-4} , in a system where D=100 and each state has current on the order of 1. Nonetheless, it is an interesting direction for future work to investigate how to get gradient descent to converge to greater degrees of numerical precision in these settings; and, in general, how to improve the performance of all of these parallel sequence evaluators in lower numerical precision.

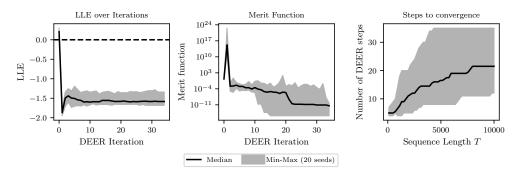


Figure 7: In this plot, we provide additional information about the behavior of DEER when rolling out Langevin dynamics on a two-well potential. (Left) We observe that across 20 random seeds (including different Langevin dynamics trajectories), the LLE for intermediate DEER iterations becomes negative after the first iteration. Consequently, we observe that the merit function (Center) experiences a spike on the very first DEER iteration (following initialization, which was the only trajectory with positive LLE), before trending towards convergence. As the system spends most of its time in contracting regions, we observe (Right) that the number of DEER iterations needed for convergence scales sublinearly with the sequence length T. We plot the min-max range for 20 seeds, and observe that even out of 20 seeds, the maximum number of DEER iterations needed to converge on a sequence length of T=10,000 is around 35.

J.4 Additional details for the two-well potential

We form the two-well potential for our experiment in Section 5 as a sum of two quadratic potentials. Concretely, we define the potential ϕ as the negative log probability of the mixture of two Gaussians, where one is centered at (0, -1.4) and the other is centered at (0, 1.6), and they both have diagonal covariance. In Langevin dynamics [80, 81] for a potential ϕ , the state s_t evolves according to

$$s_{t+1} = s_t - \epsilon \nabla \phi(s_t) + \sqrt{2\epsilon} w_t,$$

where ϵ is the step size and $w_t \stackrel{\mathrm{iid}}{\sim} \mathcal{N}(0, I_D)$. ⁵ Accordingly, the Jacobians of the dynamics (those used in DEER) take the form

$$J_t = I_D - \epsilon \nabla^2 \phi(s_t).$$

As a result, the dynamics are contracting in regions where ϕ has positive curvature (inside of the wells, where the dynamics are robustly oriented towards one of the two basins) and unstable in regions where ϕ has negative curvature (in the region between the two wells, where the stochastic inputs can strongly influence which basin the trajectory heads towards). We observe that even though there are regions in state space where the dynamics are not contracting, the resulting trajectories have negative LLE. Accordingly, in Figure 3 (Right), we observe that the number of DEER iterations needed for convergence scales sublinearly, as the LLE of all the intermediate DEER trajectories after initialization are negative. These results demonstrate that if the DEER optimization path remains in contractive regions on average, we can still attain fast convergence rates as the sequence length grows.

Moreover, a further added benefit of our theory is demonstrated by our choice of initialization of DEER. Both [1] and [3] exclusively initialized all entries of $\mathbf{s}^{(0)}$ to zero. However, such an initialization can be extremely pathological if the region of state space containing $\mathbf{0}$ is unstable, as is the case for the particular two well potential we consider. For this reason, we initialize $\mathbf{s}^{(0)}$ at random (as iid standard normals).

An important consequence of this experiment is that it shows that there are systems that are not globally contracting that nonetheless enjoy fast rates of convergence with DEER. This fact is important because a globally contractive neural network may not be so interesting/useful for classification, while a locally contracting network could be.

⁵Notice that this is a discretization (with time step ϵ) of the Langevin Diffusion SDE $ds(t) = -\nabla \phi(s(t))dt + \sqrt{2}dw(t)$, where w(t) is Brownian motion [82–84].

J.5 Building Stable Observers for Chaotic Systems

To further demonstrate the applicability of our results—and to validate them in the context of non-autonomous systems—we construct nonlinear observers. Observers are commonly used in science and engineering to reconstruct the full state of a system from partial measurements [30, 31]. As a benchmark, we consider nine chaotic flows from the dysts dataset [85]. According to Theorem (2), these systems exhibit poorly conditioned merit function landscapes and are thus not well-suited for parallelization via DEER. If the corresponding observers are stable, then they should be suitable for DEER.

We design observers for these systems using two standard approaches: (1) by directly substituting the observation into the observer dynamics, following Pecora and Carroll [86], or (2) by incorporating the observation as feedback through a gain matrix, as in Zemouche and Boutayeb [87]. We then apply DEER to compute the trajectories of both the original chaotic systems and their corresponding stable observers. As anticipated by Theorem (2), the chaotic systems exhibit slow convergence—often requiring the full sequence length—whereas the stable observers converge rapidly (Figure 8).

As with the two-well experiment, we initialize our guess for $s_t^{(0)}$ as iid standard normals.

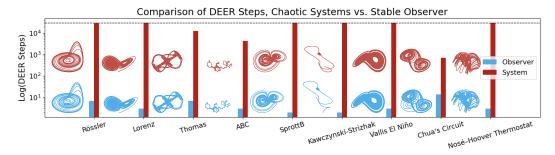


Figure 8: Comparison of DEER convergence behavior for original chaotic systems (red) and corresponding stable observers (blue) across nine flows taken from the dysts dataset. As predicted by Theorem (2), the chaotic systems converge slowly—often taking the whole sequence length T, denoted by the horizontal dashed line—due to poorly conditioned merit landscapes, while the stable observers achieve rapid convergence

J.6 Numerical computation of the discrete-time LLE

The Largest Lyapunov Exponent (LLE), which we often denote by λ , is defined in Definition 1. However, for long sequences T, naively computing it would be numerically unstable. Thus, we use Algorithm 1 to compute the LLE in a numerically stable way. Note that the algorithm nominally depends on the initial unit vector u_0 . For this reason, we choose 3 different unit vectors (initialized at random on the unit sphere) and average over the 3 stochastic estimates. However, in practice we observe that the estimate is very stable with respect to choice u_0 , and agrees with systems for which the true LLE is known, such as the Henon and logistics maps.

Algorithm 1 Numerically Stable Computation of Largest Lyapunov Exponent (LLE)

- 1: **Input:** Initial unit vector u_0 , total iterations T
- 2: **Initialize:** LLE $\leftarrow 0$
- 3: **for** t = 1 to T **do**
- 4: Compute evolved vector: $u_t \leftarrow J_t u_{t-1}$
- 5: Compute stretch factor: $\lambda_t \leftarrow ||u_t||$
- 6: Normalize vector: $u_t \leftarrow u_t/\lambda_t$
- 7: Accumulate logarithmic stretch: LLE \leftarrow LLE + log λ_t
- 8: **Output:** Estimated LLE $\lambda \leftarrow \text{LLE}/T$

K Discrete and Continuous Time LLE

We provide the definition of the LLE of a discrete-time dynamical system (often called a *map*) in Definition 1. As our paper studies discrete-time SSMs as in (1), this discrete-time definition of LLE makes sense for our setting. However, as many of our experiments involve the discretization of continuous time systems, we want to review how the LLEs of discrete and continuous time systems relate to each other. Helpful references on this topic include [7, 88].

The LLE quantifies what happens to a perturbation δx over time. Does its magnitude $\|\delta x\|$ grow or shrink over time? The LLE λ is that value that makes eq. (6) hold, i.e.

$$\|\delta x(t)\| \sim e^{\lambda t} \|\delta x(0)\|.$$

This notion of the change in the size of a perturbation $\delta x(t)$ over time—as quantified by the LLE—makes sense for both a discrete-time SSM $x_t = f_t(x_{t-1})$ as well as a continuous time system $\dot{x} = F(x,t)$.

Let us consider how a perturbation δx evolves in both a discrete time system $x_t = f_t(x)$ and a continuous time system $\dot{x} = F(x,t)$. An infinitesimal perturbation δx is intimately related to derivatives with respect to x. Therefore, their variational equations are:

discrete-time:
$$\delta x_t = \frac{\partial f_t(x_{t-1})}{\partial x} \delta x_{t-1}$$

continuous-time: $\dot{x}(t) = \frac{\partial F}{\partial x} \left(x(t), t \right) x(t).$

Case study: discretizing a continuous-time system

All of our experiments are the forward Euler discretizations of continuous-time systems. So, in this section, we we work out what happens to the LLE in such a setting. Let's consider running our continuous-time setting for a length of time T.

In this setting, if we discretize by timestep Δt , our resulting discrete-time map f is given by

$$f_{t+\Delta t}(x_t) = x_t + \Delta t F(x_t, t).$$

Therefore, it follows that

$$J_{t+\Delta t} := \frac{\partial f}{\partial x}(x_t) = I_D + \Delta t \frac{\partial F}{\partial x}(x_t).$$

We can naturally define the continuous-time LLE as the limit of the products of these J_t as $\Delta t \to 0$, i.e.

$$\lambda_c := \lim_{\Delta t \to 0} \frac{1}{T} \log \left\| \prod_{k=1}^N J_{k\Delta t} \right\|,$$

where the number of discrete time-steps N is given by $N := T/\Delta t$.

Notice that this is extremely similar to the definition of the discrete-time LLE λ_d we gave in Definition 1, except division is occurring by the length of the time window T instead of the number of discrete steps taken $N=T/\Delta t$. (Of course, N=T if $\Delta t=1$.)

Therefore, if we have a discretization of a continuous-time system, and naively plug into our discrete-time LLE defined in Definition 1 (i.e., divide by number of discrete-time steps N instead of the length of the time window T), the resulting $\lambda_d(\Delta t)$ will satisfy

$$\lambda_c = \frac{\lambda_d(\Delta t)}{\Delta t}.$$

Note that naively plugging in this discrete-time estimator would therefore result in a different estimate of the LLE depending on the size of the time-step Δt . For this reason, in our experiments we report an estimate of the continuous-time LLE λ_c given by $\lambda_d(\Delta t)/\Delta t$.

Note that dividing by Δt , which is positive, does not change the sign of λ , i.e. whether or not the system is predictable or unpredictable. However, estimating λ_c has the benefit of giving stability to the estimate even as we change Δt , and allowing for comparison to the known LLEs of continuous time systems.

⁶technically a virtual displacement, cf. Lohmiller and Slotine [7].

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Our main claim is that predictability implies parallelizability. We demonstrate this theoretically by analyzing the condition of the merit loss function, and illustrate these findings experimentally as well.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: See Limitations section in the conclusion.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide detailed, complete, and correct proofs of all of our claims in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide our code at https://github.com/lindermanlab/predictability_enables_parallelization

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide code the in supplement.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide experimental descriptions in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Our experiments are run with 20 random seeds, with the full range being reported in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide experimental details in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: Our paper is theoretical and so doesn't involve human subjects or proprietary data.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss the wide-ranging applicability of our theoretical contribution: that understanding the predictability implies parallelizability allows for the parallelizability of nonlinear systems across many domains.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point

out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not create any language models or generators or scraped datasets.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite original creators, such as the creator of the dysts benchmark dataset.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We provide our code and documentation for it at https://github.com/lindermanlab/predictability_enables_parallelization

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not involve research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: the paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core methods introduced in this paper do not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.