On the difficulty of learning chaotic dynamics with RNNs

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recurrent neural networks (RNNs) are wide-spread machine learning tools for modeling sequential and time series data. They are notoriously hard to train because their loss gradients backpropagated in time tend to saturate or diverge during training. This is known as the exploding and vanishing gradient problem. Previous solutions to this issue either built on rather complicated, purpose-engineered architectures with gated memory buffers, or - more recently - imposed constraints that ensure convergence to a fixed point or restrict (the eigenspectrum of) the recurrence matrix. Such constraints, however, convey severe limitations on the expressivity of the RNN. Essential intrinsic dynamics such as multistability or chaos are disabled. This is inherently at disaccord with the chaotic nature of many, if not most, time series encountered in nature and society. It is particularly problematic in scientific applications where one aims to reconstruct the underlying dynamical system. Here we offer a comprehensive theoretical treatment of this problem by relating the loss gradients during RNN training to the Lyapunov spectrum of RNN-generated orbits. We mathematically prove that RNNs producing stable equilibrium or cyclic behavior have bounded gradients, whereas the gradients of RNNs with chaotic dynamics always diverge. Based on these analyses and insights we suggest ways of how to optimize the training process on chaotic data according to the system's Lyapunov spectrum, regardless of the employed RNN architecture.

1 Introduction

Recurrent neural networks (RNNs) are widely used across various fields in engineering and science for learning sequential tasks or modeling and predicting time series [55]. Yet, they struggle when long-term temporal dependencies, very slow, or hugely varying time scales are involved [6, 35, 53, 73, 80]. Time series or sequential data with such properties are, however, very common in fields like climate physics [84], neuroscience [23, 75], ecology [87], or language processing [12]. Training RNNs on such data is hard because the loss gradients backpropagated in time easily saturate or diverge in this process. This is commonly referred to as the exploding and vanishing gradient problem (EVGP) [6, 35, 67].

One solution to the EVGP is based on specifically designed RNN architectures with gating mechanisms, such as long short-term memory (LSTM) [36] or gated recurrent units (GRU) [11]. These architectures allow states at earlier time steps to more easily influence activity much later through a kind of protected memory buffer, thus alleviating the EVGP by structural design. In practice, such models need to be backed up by further techniques like gradient clipping to keep the gradients in check [67]. The relatively complex architectural design of these networks impedes their mathematical analysis and requires reverse engineering after training [57, 63, 64, 80]. Partly to forego these complications, a variety of other solutions has been proposed recently, imposing restrictions on the recurrence matrix to bound the gradients [4, 10], or enforcing global stability by design or

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

regularization [20, 50]. Often these procedures dramatically curtail the expressivity of the RNN [46, 66, 80]; in particular, they rule out chaotic dynamics (for reasons discussed further below).

This is at odds with the plethora of chaotic phenomena in nature, engineering, and society. Chaotic dynamics are commonplace, almost default in any complex physical or biological system. This includes scientific areas as diverse as neuroscience [18, 93], physiology [47], geophysics [81], climate systems [88], astrophysics [52], ecology [15], chemical reactions [22], cell [65] or population [59] biology. Chaotic phenomena are also crucial for the understanding of societal and epidemiological processes, such as the spread of diseases [58], or in economics [21]. They are further relevant in purely technical contexts such as electrical engineering [42, 83] or laser optics [44]. They have even been suggested to play an up to now largely neglected, but potentially very significant role in speech recognition [77] and natural language processing [37]. Hence, in almost any practical setting, chaotic phenomena abound. They cannot, in general, be ignored when devising RNN training algorithms.

Here we offer a comprehensive theoretical treatment of the relation between RNN dynamics and the behavior of the loss gradients during training. We find a close connection between an RNN's loss gradients and the largest Lyapunov exponent of its freely generated orbits. We mathematically prove that RNNs producing stable fixed point or cyclic behavior have bounded gradients. Crucially, however, the loss gradients of RNNs producing chaotic dynamics always diverge. Hence, the chaotic nature of many time series data induces a *principle* problem, and, despite significant efforts in the past to solve the EVGP, training RNNs on such data remains an open issue. We illustrate the implications of our theory for RNN training on several simulated and empirical chaotic time series, and adapt the old idea of *sparsely forced BPTT* as a simple yet effective remedy that enables to learn the underlying dynamics despite exploding gradients.

2 Related works

Exploding and vanishing gradients. While 'classical' remedies of the EVGP [6, 35, 67] rest on purpose-tailored architectures with gating mechanisms, which safeguard information flow across longer temporal distances [11, 36], the focus has recently shifted to simpler RNNs that address the EVGP by restricting the recurrence matrix to be orthogonal [32, 33, 38], unitary [4], or antisymmetric [10], or by ensuring globally stable fixed point solutions [39, 41], for example through co-trained Lyapunov functions [50]. However, all these approaches impose strong limitations on the dynamical repertoire of the RNN, enforcing global convergence to fixed points or simple cycles.[1] In doing so, they drastically reduce the expressiveness of these models [46, 66]. To address this problem, Erichson et al. [20] somewhat relaxed the constraints on the recurrence matrix by introducing a skew-symmetric decomposition combined with a Lipschitz condition on the activation function. Another recent approach discretizes oscillator ODEs to arrive at a stable system of coupled [73] or independent [74] oscillators which increase the RNN's expressiveness while bounding its gradients. By design (and as acknowledged by the authors), neither of these architectures is capable of producing chaotic dynamics, however, as the underlying ODEs do not allow for exponential divergence of close-by trajectories (a prerequisite for chaos). Given these often principle limitations of parametrically or dynamically strongly constrained models, a fruitful direction may be to modify the training process itself, e.g. through modified or auxiliary loss functions [80, 85], or special procedures for parameter updating [40] or loss truncation [61, 95]. Our empirical evaluation will follow up on such ideas, but also highlight that simple loss truncation, windowing, or architectural solutions like LSTMs are not sufficient.

Learning dynamical systems. Surprisingly disconnected from the work on the EVGP and learning long-term dependencies, a huge and long-standing literature deals with training RNNs on nonlinear dynamical systems (DS) [70, 86, 89], including chaotic systems like the famous Lorenz equations [56] or chaotic turbulence in fluid dynamics [54, 69]. Of those, methods based on reservoir computing [69] are special in that they start with a large complex dynamical repertoire to begin with for which an optimal mapping onto the observations is learned, together with a feedback to the reservoir (see also [9] for issues associated with this strategy in the context of DS reconstruction). Teacher forcing (TF; [14, 70, 95], see also [29]) is one of the earliest techniques introduced to keep RNN trajectories on track while training. The idea behind TF is to simply replace RNN states by observations when available, thereby also effectively cutting off the gradients. TF essentially derives from ideas in dynamical control theory, and adaptive schemes that increasingly hand over control to the RNN

[1] We make this point more formal in Appx. A.1.6.

throughout training have been devised [1, 2, 5]. A related technique from the control theory literature is "multiple shooting" [92]: Here the whole observed time series is chopped into chunks, and for each chunk of trajectory a new initial condition is estimated. Explicit constraints ensure continuity between the separate trajectory bits during optimization. State space models and the Expectation-Maximization algorithm became popular particularly in the 90es for uncovering the latent dynamics underlying a set of time series observations [24], and remain an important tool until today [17, 51]. Most recently, approaches based on variational inference and the reparameterization trick [49], like sequential variational autoencoders (SVAE), gained in popularity for DS approximation [8, 34]. "Deterministic" RNNs (i.e., with latent states not treated as random variables), like conventional LSTMs [90], remain top choices for DS reconstruction, however.

Although connections between DS ideas and loss gradients have been drawn early on [6], so far only particular scenarios (like fixed point attractors) have been considered. Closest to our work is recent work by Schmidt et al. [80], where non-divergence of loss gradients is established when RNNs converge to fixed points or cycles. However, this was done only for the particular class of piecewise-linear RNNs (PLRNNs), more restrictive conditions for cycles were imposed than assumed here, and - most importantly - the chaotic case on which we focus here was not considered. Recent studies [19, 91] also point out the general connections between Lyapunov exponents and loss gradients that we develop in sect. 3.1, but do not provide any in-depth theoretical treatment, proofs, or empirical evaluation of methods to alleviate exploding gradients in training, as we do here. Thus, a systematic theoretical framework that relates RNN dynamics more generally, and across a range of different RNN architectures, to the behavior of its training gradients, is still lacking so far.

3 Theoretical analysis: Relation between RNN dynamics and loss gradients

In our analysis, we will cover all major types of system dynamics (fixed points, cycles, chaos, and quasi-periodicity), and mathematically investigate their implications for the loss gradients. We will do this for all major classes of RNNs, including standard RNNs with largely arbitrary activation function, LSTMs, GRUs, and PLRNNs. The next section will first develop and illustrate the basic intuition behind the relations between RNN dynamics and loss gradients.

3.1 Preliminaries: RNN dynamics and loss gradients

Formally, all popular RNN architectures, including LSTMs, GRUs, or PLRNNs, are discrete time DS, defined by a (first-order-Markovian) recursive prescription for the temporal evolution of the latent states $\boldsymbol{z}_t \in \mathbb{R}^M$ of the form

$$\boldsymbol{z}_t = F_{\boldsymbol{\theta}}(\boldsymbol{z}_{t-1}, \boldsymbol{s}_t), \tag{1}$$

where $\boldsymbol{s}_t \in \mathbb{R}^N$ is the input at time t and $\boldsymbol{\theta}$ are RNN parameters. For instance, for standard RNNs we have $F_{\boldsymbol{\theta}}(\boldsymbol{z}_{t-1}, \boldsymbol{s}_t) = f\big(\boldsymbol{W}\boldsymbol{z}_{t-1} + \boldsymbol{B}\boldsymbol{s}_t + \boldsymbol{h}\big)$, where f is an element-wise activation function like $\tanh$ or a rectified linear unit (ReLU).

Assuming we start at some initial value $\boldsymbol{z}_1 \in \mathbb{R}^M$, and given a sequence of external inputs $\boldsymbol{S} = \{\boldsymbol{s}_t\}$, we can recursively rewrite eq. (1) as

$$\boldsymbol{z}_T = F_{\boldsymbol{\theta}}(F_{\boldsymbol{\theta}}(F_{\boldsymbol{\theta}}(...F_{\boldsymbol{\theta}}(\boldsymbol{z}_1, \boldsymbol{s}_2)...))) \;=:\; F_{\boldsymbol{\theta}}^{T-1}(\boldsymbol{z}_1, \boldsymbol{s}_2). \tag{2}$$

In DS theory, we characterize the long-term behavior of such sequences by its spectrum of Lyapunov exponents. The Lyapunov exponents estimate the exponential growth rates in different local directions of the system's state space, and the largest Lyapunov exponent gives the dominant exponential behavior. Let us denote the system's Jacobian at time t by

$$\boldsymbol{J}_t \;:=\; \frac{\partial F_{\boldsymbol{\theta}}(\boldsymbol{z}_{t-1}, \boldsymbol{s}_t)}{\partial \boldsymbol{z}_{t-1}} \;=\; \frac{\partial \boldsymbol{z}_t}{\partial \boldsymbol{z}_{t-1}}. \tag{3}$$

Then, the maximum Lyapunov exponent along an RNN trajectory $\{\boldsymbol{z}_1, \boldsymbol{z}_2, \cdots, \boldsymbol{z}_T, \cdots\}$ is defined as

$$\lambda_{max} \;:=\; \lim_{T\to\infty} \frac{1}{T} \log \left\| \prod_{r=0}^{T-2} \boldsymbol{J}_{T-r} \right\|, \tag{4}$$

where $\|\cdot\|$ denotes the spectral norm (or any subordinate norm) of a matrix. If $\lambda_{max} < 0$ nearby trajectories will ultimately converge to a fixed point or cycle, while for $\lambda_{max} > 0$ (a necessary condition for chaos) initially nearby trajectories will exponentially separate, i.e. we will have divergence along one (or more) directions in state space. This accounts for the sensitive dependence on initial conditions in chaotic systems.

Now let $\mathcal{L}(\boldsymbol{W}, \boldsymbol{B}, \boldsymbol{h})$ be some loss function employed for RNN training that decomposes in time as $\mathcal{L} = \sum_{t=1}^{T} \mathcal{L}_t$. Suppose we fancy BPTT as our training algorithm (similar derivations could be performed for RTRL), we recursively develop the loss gradients w.r.t. some RNN parameter θ in time (i.e., across layers of the RNN unrolled in time) as

$$\frac{\partial \mathcal{L}}{\partial \theta} = \sum_{t=1}^{T} \frac{\partial \mathcal{L}_t}{\partial \theta} \quad \text{with} \quad \frac{\partial \mathcal{L}_t}{\partial \theta} = \sum_{r=1}^{t} \frac{\partial \mathcal{L}_t}{\partial \boldsymbol{z}_t} \frac{\partial \boldsymbol{z}_t}{\partial \boldsymbol{z}_r} \frac{\partial^+ \boldsymbol{z}_r}{\partial \theta}, \tag{5}$$

and

$$\begin{aligned} \frac{\partial \boldsymbol{z}_t}{\partial \boldsymbol{z}_r} &= \frac{\partial \boldsymbol{z}_t}{\partial \boldsymbol{z}_{t-1}} \frac{\partial \boldsymbol{z}_{t-1}}{\partial \boldsymbol{z}_{t-2}} \cdots \frac{\partial \boldsymbol{z}_{r+1}}{\partial \boldsymbol{z}_r} \\ &= \prod_{k=0}^{t-r-1} \frac{\partial \boldsymbol{z}_{t-k}}{\partial \boldsymbol{z}_{t-k-1}} = \prod_{k=0}^{t-r-1} \boldsymbol{J}_{t-k}, \end{aligned} \tag{6}$$

where ∂^+ denotes the immediate derivative. Now observe that the behavior of the loss gradients crucially depends on the *product series* of Jacobians in eqn. (6): If the maximum absolute eigenvalues of the Jacobians $\boldsymbol{J}_t$ will, in the geometric mean, be larger than 1 (i.e., $\left\| \prod_{r=0}^{T-2} \boldsymbol{J}_{T-r} \right\|^{1/T} > 1$), gradients will explode as $T \to \infty$, while they will saturate if $\left\| \prod_{r=0}^{T-2} \boldsymbol{J}_{T-r} \right\|^{1/T} < 1$. Thus, the key point to note is that the same terms that occur in the definition of the Lyapunov spectrum, eqn. (4), resurface in the loss gradients, eqn. (5) & (6). This accounts for the tight links between system dynamics and gradients.

3.2 Fixed points and cyclic dynamics

Let us start by considering the simplest types of dynamics that can occur in RNNs (or any discrete-time DS): fixed points and cycles. In fact, by far most of the literature on global stability in RNNs and on loss gradients focused on just fixed points [10, 20, 50], with only few authors who recently started to also connect cyclic behavior to loss gradients [73, 80]. Recall that a fixed point of a recursive map $\boldsymbol{z}_t = F(\boldsymbol{z}_{t-1})$ is defined as a point $\boldsymbol{z}^*$ for which we have $\boldsymbol{z}^* = F(\boldsymbol{z}^*)$.[2] Likewise, a k-cycle ($k > 1$) is a set of temporally consecutive periodic points $P_k := \{\boldsymbol{z}_{t_1}, \boldsymbol{z}_{t_2}, \ldots, \boldsymbol{z}_{t_k}\} = \{\boldsymbol{z}_{t_1}, F(\boldsymbol{z}_{t_1}), \ldots, F^{k-1}(\boldsymbol{z}_{t_1})\}$ that we obtain from recursive application of the map such that each of the cyclic points $\boldsymbol{z}_{t_r} \in P_k$ is a fixed point of the k times iterated map F^k (with k being the smallest positive integer for which this holds). To simplify the subsequent treatment, we will collectively refer to fixed points and cycles as k-cycles ($k \geq 1$). Further recall that a fixed point or k-cycle is called *stable* if the maximum absolute eigenvalue of the Jacobian evaluated at that point is smaller than 1, *neutrally stable* if exactly 1, and *unstable* otherwise. Although the results we develop in this and the following sections will hold more widely, we will restrict our attention to recursive maps $F_{\boldsymbol{\theta}}$ from the class of RNNs $\mathcal{R} = \{\texttt{standardRNN}, \texttt{LSTM}, \texttt{GRU}, \texttt{PLRNN}\}$ (see Appx. A.1 for details).

Based on the observations made in the previous sections we can state the following theorem that links RNN dynamics and loss gradients:

Theorem 1. *Consider an RNN $F_{\boldsymbol{\theta}} \in \mathcal{R}$ parameterized by $\boldsymbol{\theta}$, and assume that it converges to a stable fixed point or k-cycle Γ_k ($k \geq 1$) with $\mathcal{B}_{\Gamma_k}$ as its basin of attraction. Then for every $\boldsymbol{z}_1 \in \mathcal{B}_{\Gamma_k}$ (i) the Jacobian $\frac{\partial \boldsymbol{z}_T}{\partial \boldsymbol{z}_1}$ exponentially vanishes as $T \to \infty$; (ii) for Γ_k the tangent vectors $\frac{\partial \boldsymbol{z}_T}{\partial \theta}$ and thus the gradient of the loss function, $\frac{\partial \mathcal{L}_T}{\partial \boldsymbol{\theta}}$, will be bounded from above, i.e. will not diverge for $T \to \infty$; and (iii) for the PLRNN (27) both $\left\| \frac{\partial \boldsymbol{z}_T}{\partial \theta} \right\|$ and $\left\| \frac{\partial \mathcal{L}_T}{\partial \boldsymbol{\theta}} \right\|$ will remain bounded for every $\boldsymbol{z}_1 \in \mathcal{B}_{\Gamma_k}$ as $T \to \infty$.*

[2] From here on we will suppress the explicit dependence on external inputs $\boldsymbol{s}_t$ notation-wise, see Remark 2.

Proof. (i) Assume that Γ_k is a stable k-cycle $(k \geq 1)$ denoted by

$$\Gamma_k = \{\boldsymbol{z}_1, \boldsymbol{z}_2, \cdots, \boldsymbol{z}_T, \cdots\} = \{\boldsymbol{z}_{t^{*k}}, \boldsymbol{z}_{t^{*k}-1}, \cdots, \boldsymbol{z}_{t^{*k}-(k-1)}, \boldsymbol{z}_{t^{*k}}, \boldsymbol{z}_{t^{*k}-1}, \cdots, \boldsymbol{z}_{t^{*k}-(k-1)}, \cdots\}. \tag{7}$$

Then, the largest Lyapunov exponent of Γ_k is given by

$$\begin{aligned} \lambda_{\Gamma_k} &= \lim_{t\to\infty} \frac{1}{t} \ln \left\| J_t^* \, J_{t-1}^* \, \cdots \, J_2^* \right\| \\ &= \lim_{j\to\infty} \frac{1}{jk} \ln \left\| \left(\prod_{s=0}^{k-1} J_{t^{*k}-s} \right)^j \right\|. \end{aligned} \tag{8}$$

By assumption we have $\lambda_{\Gamma_k} < 0$ and also $\rho\big(\prod_{s=0}^{k-1} J_{t^{*k}-s}\big) < 1$ (the spectral radius), which implies

$$\lim_{t\to\infty} J_t^* \, J_{t-1}^* \, \cdots \, J_2^* = \lim_{j\to\infty} \left(\prod_{s=0}^{k-1} J_{t^{*k}-s} \right)^j = 0. \tag{9}$$

Now suppose that $\mathcal{O}_{\boldsymbol{z}_1}$ is an orbit of the map eqn. (1) converging to Γ_k, i.e. $\boldsymbol{z}_1 \in \mathcal{B}_{\Gamma_k}$. Since $\mathcal{O}_{\boldsymbol{z}_1}$ and Γ_k have the same largest Lyapunov exponent, we have

$$\lambda_{\mathcal{O}_{\boldsymbol{z}_1}} = \lim_{T\to\infty} \frac{1}{T} \ln \left\| J_T \, J_{T-1} \, \cdots \, J_2 \right\| = \lambda_{\Gamma_k} < 0, \tag{10}$$

and hence for $\boldsymbol{z}_1 \in \mathcal{B}_{\Gamma_k}$

$$\lim_{T\to\infty} \left\| \frac{\partial \boldsymbol{z}_T}{\partial \boldsymbol{z}_1} \right\| = \lim_{T\to\infty} \left\| J_T \, J_{T-1} \, \cdots \, J_2 \right\| = 0. \tag{11}$$

(ii) & (iii) See Appx. A.2.1. □

Remark 1. *The result of Theorem 1 part (i) will be generally true for any first-order-Markovian recursive map (1), but the conclusions in part (ii) may hinge on its specific definition.*

Remark 2. *None of the results above and throughout sect. 3 require the dynamics to be autonomous, the theory applies whether there is external input or not. In fact, mathematically, non-autonomous (externally forced) systems can always be rewritten as autonomous dynamical systems [3, 71, 97], see Appx. A.1.1 for details.*

The results above ensure that loss gradients will not diverge (explode) as $T \to \infty$ in RNNs that are "well-behaved" in the sense that they converge to a fixed point or cycle. This is a generalization of the results given in Theorem 1 in Schmidt et al. [80], where this was shown only a) for the specific class of PLRNNs and b) for specific constraints imposed on the eigenvalue spectrum of the RNN's Jacobians which were relaxed in our theorem above.

While our treatment above is centered on the "exploding-gradients" case, various architectural modifications or regularization techniques can ensure that gradients do not vanish either, i.e. remain bounded from below as well. This was established, for instance, in Schmidt et al. [80] for PLRNNs using 'manifold attractor regularization'. In Appx. A.2.1 we show that the results from Theorem 2 from Schmidt et al. [80] on *doubly* bounded (from below and above) loss gradients can indeed be extended to the more general case covered by Theorem 1 above.

3.3 Chaotic dynamics

We will now consider the all-important chaotic case. Let F be a recursive map and $\mathcal{O}_{\boldsymbol{z}_1} = \{\boldsymbol{z}_1, \boldsymbol{z}_2, \boldsymbol{z}_3, \cdots\}$ be an orbit of F. The orbit is chaotic if (i) it is not asymptotically periodic and (ii) has at least one positive Lyapunov exponent [27, 60]. If the system's invariant set is *bounded*, condition (ii) is considered a standard signature of chaos, as in this case two nearby orbits separate exponentially fast, but at the same time their mutual separation cannot go to infinity so that there are also folds. The following theorem states the sufficient condition for exploding gradients:

Theorem 2. *Suppose that an RNN $F_{\boldsymbol{\theta}} \in \mathcal{R}$ (parameterized by $\boldsymbol{\theta}$) has a chaotic attractor Γ^* with $\mathcal{B}_{\Gamma^*}$ as its basin of attraction. Then, for almost every orbit with $\boldsymbol{z}_1 \in \mathcal{B}_{\Gamma^*}$, (i) the Jacobians connecting temporally distal states $\boldsymbol{z}_T$ and $\boldsymbol{z}_t$ $(T \gg t)$, $\frac{\partial \boldsymbol{z}_T}{\partial \boldsymbol{z}_t}$, will exponentially explode for $T \to \infty$, and (ii) the tangent vector $\frac{\partial \boldsymbol{z}_T}{\partial \theta}$ and so the gradients of the loss function, $\frac{\partial \mathcal{L}_T}{\partial \theta}$, will diverge as $T \to \infty$.*

Proof. Let the RNN $F_{\boldsymbol{\theta}} \in \mathcal{R}$ have a chaotic orbit denoted by $\Gamma^* = \{\boldsymbol{z}_1^*, \boldsymbol{z}_2^*, \cdots, \boldsymbol{z}_T^*, \cdots\}$. Then, denoting by J_T^* the Jacobian of (1) at $\boldsymbol{z}_T^* \in \Gamma^*$, the largest Lyapunov exponent of Γ^* is given by

$$\lambda = \lim_{T\to\infty} \frac{1}{T} \ln \left\| J_T^* \, J_{T-1}^* \cdots J_2^* \right\| . \tag{12}$$

Since Γ^* is chaotic, so $\lambda > 0$. Hence, from (12), it is concluded that

$$\lim_{T\to\infty} \left\| J_T^* \, J_{T-1}^* \cdots J_2^* \right\| = \lim_{T\to\infty} \left\| \frac{\partial \boldsymbol{z}_T^*}{\partial \boldsymbol{z}_t^*} \right\| = \infty, \quad T \gg t. \tag{13}$$

Now, according to Oseledec's multiplicative ergodic Theorem, almost all the points in the basin of attraction of Γ^* have the same largest Lyapunov exponent λ. Thus, (13) holds for almost every $\boldsymbol{z}_1 \in \mathcal{B}_{\Gamma^*}$.

(ii) See Appx. A.2.2. □

Remark 3. *The first part of Theorem 2 holds for all first-order-Markovian recursive maps (1). Note that for LSTMs, $\frac{\partial \boldsymbol{z}_T}{\partial \boldsymbol{z}_t}$ ($\boldsymbol{z} := (\boldsymbol{h}, \boldsymbol{c})^\top$) denotes the full Jacobian of both hidden and cell states.*

We collect some further mathematical results and remarks related to Theorem 2 in Appx. A.3.1.

Hence, the essential result is that for all popular RNNs $\mathcal{R}$ and activation functions, loss gradients will inevitably diverge if the RNN latent states converge to a chaotic attractor.

3.4 Quasi-periodicity

Quasi-periodicity is a long-term behavior which occurs on a torus and, superficially, bears some similarity to chaos in the sense that, strictly speaking, orbits are also *aperiodic*. That is, as $T \to \infty$, trajectories will never close up with themselves. Moreover, every trajectory becomes arbitrarily close to any point on the torus, that is, it is dense. One important difference between quasi-periodic and chaotic systems is, however, that in a quasi-periodic system, as time passes, two close initial conditions are *linearly* diverging, while in a chaotic system the divergence is exponential.

Theorem 3. *Assume that an RNN $F_{\boldsymbol{\theta}} \in \mathcal{R}$ (parameterized by $\boldsymbol{\theta}$) has a quasi-periodic attractor Γ with $\mathcal{B}_\Gamma$ as its basin of attraction. Then, for every $\boldsymbol{z}_1 \in \mathcal{B}_\Gamma$*

$$\forall\, 0 < \epsilon < 1 \;\; \exists\, T_0 > 1 \;\; s.t. \;\; \forall\, T \geq T_0 \implies$$

$$(1-\epsilon)^{T-1} < \left\| \frac{\partial \boldsymbol{z}_T}{\partial \boldsymbol{z}_1} \right\| < (1+\epsilon)^{T-1}. \tag{14}$$

Proof. See Appx. A.2.3. □

According to Theorem 3, for every orbit converging to a quasi-periodic attractor, the Jacobians $\frac{\partial \boldsymbol{z}_T}{\partial \boldsymbol{z}_t}$ may diverge or vanish as $T \to \infty$, but this will *not* occur exponentially fast as $T \to \infty$. Thus, even for bounded non-chaotic RNNs we may sometimes stumble into the problem of diverging gradients. Although this may be a less common scenario, we point out it may occur if we train RNNs on real data from oscillatory systems with incommensurate frequencies, as for instance encountered in electronic engineering.

In Appx. A.3.2 we have collected further mathematical results on the connection between RNN dynamics and loss gradients that hold regardless of the RNN's limiting behavior.

4 Empirical evaluation

Our theoretical results imply that chaotic time series pose a principle challenge for RNN training that cannot easily be circumvented through specifically designed architectures, constraints, or regularization criteria. If the underlying DS we aim to capture is chaotic, loss gradients propagated back in time will inevitably explode. Hence we need to curtail gradients in an ideal way. The issue arises especially in scientific ML where time series from chaotic systems are ubiquitous and the aim is to *reconstruct* the generating DS with its limiting behavior. Thus, our exposition will focus on this area.

4.1 Training on systems with exploding gradients by sparse teacher forcing

To illustrate the connections between theory and RNN training, we revive the old idea of TF [95] as a mechanism for truncating error gradients while training. However, we would like to do this such that important information about the system dynamics does not get lost, for which Lyapunov theory offers some guidance. Specifically, we should not force the system back onto the true trajectory all or most of the time (as in "classical TF"), but should effectively "re-calibrate" it only at certain time points chosen wisely according to the system's local divergence rates. This procedure will be referred to as *sparsely forced BPTT* in the following. Assume we want to train an RNN with hidden states $\boldsymbol{z}_t \in \mathbb{R}^M$ and linear (or affine) output layer on a time-series $\{\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_T\}$ generated by a chaotic system.[3] The linear output layer $\hat{\boldsymbol{x}}_t = \boldsymbol{B}\boldsymbol{z}_t$, $\boldsymbol{B} \in \mathbb{R}^{N \times M}$, maps the RNN hidden states into the observation space. This allows us to modify the original TF procedure by constructing a control series $\{\tilde{\boldsymbol{z}}_1, \tilde{\boldsymbol{z}}_2, \cdots, \tilde{\boldsymbol{z}}_T\}$ from the observations by "inverting" the linear output mapping [4]

$$\tilde{\boldsymbol{z}}_t = (\boldsymbol{B}^{\mathsf{T}}\boldsymbol{B})^{-1}\boldsymbol{B}^{\mathsf{T}}\boldsymbol{x}_t. \tag{15}$$

The idea is to supply this control signal only sparsely, separated by the learning interval τ between consecutive forcings. Hence, defining $\mathcal{T} = \{n\tau + 1\}_{n \in \mathbb{N}_0}$ as the set of all time points at which we force the RNN onto the 'true' values, the RNN updates can be written as

$$\boldsymbol{z}_{t+1} = \begin{cases} RNN(\tilde{\boldsymbol{z}}_t) & \text{if } t \in \mathcal{T} \\ RNN(\boldsymbol{z}_t) & \text{else} \end{cases}. \tag{16}$$

This forcing is applied *after* calculation of the loss, such that $\mathcal{L}_t = \|\boldsymbol{x}_t - \boldsymbol{B}\boldsymbol{z}_t\|_2^2$ irrespective of whether t is in $\mathcal{T}$ or not (and of course it is applied only during training, not at test time!). Replacing hidden states $\boldsymbol{z}_t$ with their teacher-forced signals $\tilde{\boldsymbol{z}}_t$ simply breaks divergence between true and predicted trajectories at time points $t \in \mathcal{T}$, and also cuts off the Jacobians by breaking the temporal contingency (for details see Appx. A.7). The learning interval τ hence controls how many time steps are included in the gradient calculation and has to be chosen with care such as to balance the effects of exploding gradients vs. those of losing relevant time scales and long-term dependencies. While it is general wisdom that an optimal batch size will facilitate training, the point here is thus much more specific: Ideally τ should be chosen in accordance with the system's Lyapunov spectrum, for instance based on the *predictability time* [7]

$$\tau_{\text{pred}} = \frac{\ln 2}{\lambda_{\max}}. \tag{17}$$

Various open-source packages exist for calculating the maximal Lyapunov exponent $\lambda_{\max}$ from empirical time series data (e.g., Julia: DynamicalSystems.jl [13], C++: TISEAN [31]), potentially after delay-embedding the data (see Appx. A.4 for more details). Note that this needs to be done on the *empirical* data and only once *before* RNN training, as $\lambda_{\max}$ is an invariant characteristic of the empirical system that we are aiming to reconstruct by our RNN (i.e., after successful training the RNN should have the same invariant properties as the underlying DS). We also emphasize that such a simple recipe for addressing the exploding gradient problem is based on modifying the training routine, and is thus in principle applicable to any model architecture.

4.2 Example 1: Lorenz system and externally forced Duffing oscillator in chaotic regime

Let us illustrate these ideas on two classical textbook examples of chaotic DS, the chaotic Lorenz attractor as an autonomous system, and the chaotically forced Duffing oscillator as an example with explicit external input (see Appx. A.4 for details). Trajectories were repeatedly drawn from these systems, on which we trained a PLRNN, a vanilla RNN with $\tanh$ activation function, and a LSTM by stochastic gradient descent (SGD) to minimize the MSE loss between predicted and actual observations. As optimizer we used Adam [48] from PyTorch [68] with a learning rate of 0.001. For all models, training proceeded solely by *sparsely forced BPTT* and did not employ gradient clipping or any other technique that may interfere with optimal loss truncation.

In nonlinear DS reconstruction, we are mainly interested in reproducing *invariant* properties of the underlying system such as the attractor geometry (or topology; [78, 82]) or the frequency composition

[3] Note that in DS reconstruction one usually considers the data as observations (*unsupervised* problem).

[4] To ensure invertibility, one could add a regularizer $\lambda\mathbf{I}$ to $\boldsymbol{B}^{\mathsf{T}}\boldsymbol{B}$ in eqn. (15), as in ridge regression, but we did not find this necessary in any of our examples.

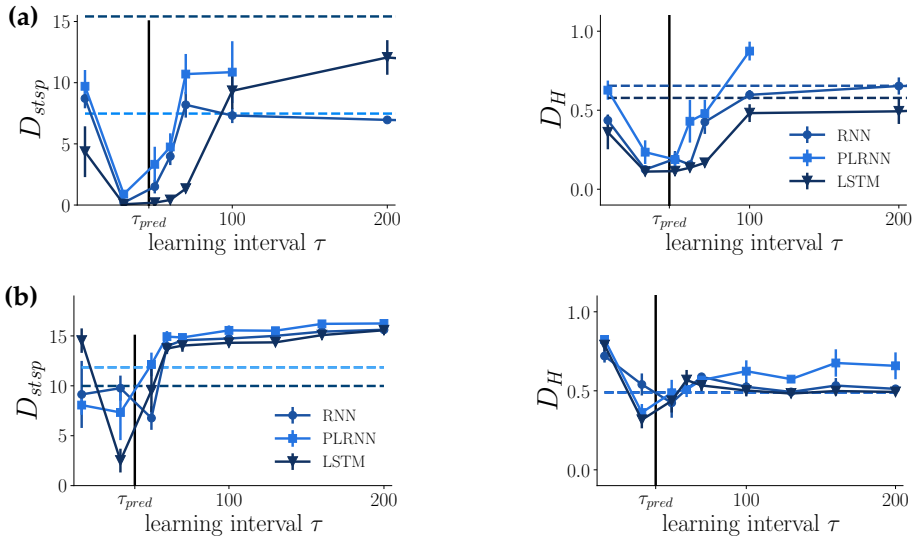

Figure 1: Overlap in attractor geometry (D_{stsp}, lower = better) and dimension-wise comparison of power-spectra (D_H, lower = better) against learning interval τ for (a) the Lorenz and (b) the chaotically forced Duffing oscillator. Continuous lines = sparsely forced BPTT. Dashed lines = classical BPTT with gradient clipping. Prediction time indicated vertically in black.

(i.e., time-averaged properties), while measures like ahead-prediction errors are less meaningful especially on chaotic time series [51, 96]. Thus, in evaluating training performance, here we follow Koppe et al. [51] in using a Kullback-Leibler divergence D_{stsp} to quantify the agreement between observed and generated probability distributions across state-space to asses the overlap in attractor geometry (Appx. A.5). Moreover, we calculate the dimension-wise Hellinger distance D_H between power spectra to quantify the temporal agreement of the observed and generated time-series.

Fig. 1 shows the dependence of the reconstruction quality on the learning interval τ for all RNN architectures on (a) the Lorenz and (b) the externally forced Duffing system. Fig. 2 provides particular examples of reconstructions for τ chosen too small, too large, or about right.

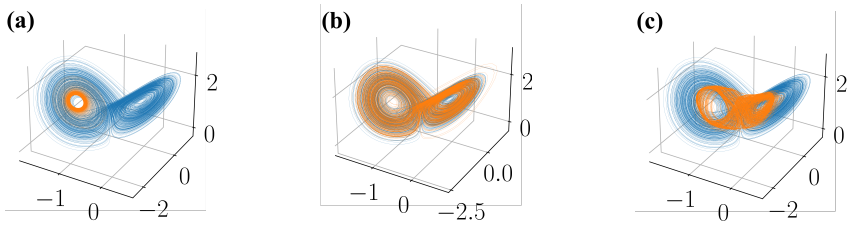

Figure 2: Lorenz attractor (blue) and example reconstructions by a LSTM (orange) trained with a learning interval (a) chosen too small ($\tau = 5$), (b) chosen optimally ($\tau = 30$), and (c) chosen too large ($\tau = 200$).

For all models we find a system-dependent range for the optimal learning interval that agrees well with the predictability time defined in eqn. (17), where estimates for the maximal Lyapunov exponent were taken from the literature [25, 72]. As a reference, dashed lines represent the reconstruction performance for all architectures when trained with classical BPTT and gradient clipping. The training procedure was the same as for sparsely forced BPTT, except that instead of supplying a control-signal, gradients were normalized to 1 prior to each parameter update (see Fig. 12 for a more systematic evaluation of different clipping procedures and thresholds). As evidenced by the much worse performance, gradient clipping does not effectively address the EVGP, *even for LSTMs*. As further shown in Fig. 9 in Appx. A.6.4, using the optimal (or any other) window length τ but resetting the initial condition to either zero or the last forward-iterated state (instead of its control value $\tilde{z}_t$) for each chunk equally destroys performance. This suggests that neither mere gradient normalization nor simple windowing are sufficient, but will wipe out essential information about the dynamics.
In Appx. A.6 we collect further results on the chaotic Rössler attractor (Fig. 5), high-dimensional Mackey-Glass equations (Fig. 7), and the Lorenz attractor with partial observations (Fig. 8).

4.3 Example 2: Chaotic weather data

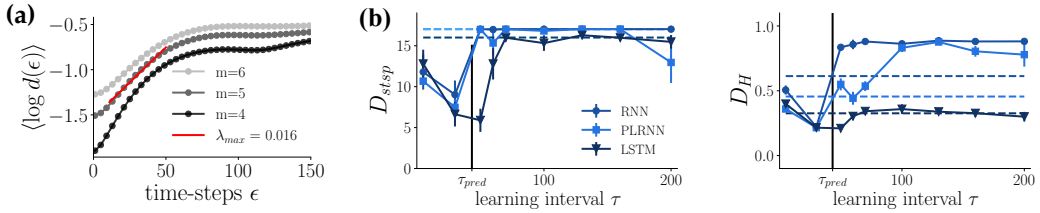

Figure 3: (a) The maximal Lyapunov exponent was determined as the slope of the average log-divergence of nearest neighbors in embedding space (m = embedding dimension). (b) Reconstruction quality assessed by attractor overlap (lower = better) and dimension-wise comparison of power-spectra (D_H, lower = better). Black vertical lines = τ_{pred}.

As for an empirical example, we trained all RNNs (vanilla RNN, PLRNN, LSTM) on a temperature time series recorded at the Weather Station at the Max Planck Institute for Biogeochemistry in Jena, Germany. To expose the chaotic behavior and obtain a robust estimate of the maximal Lyapunov exponent, trends and yearly cycles were removed, and nonlinear noise-reduction was performed ([44]; Appx. A.4). The maximal Lyapunov exponent was determined with the *TISEAN* package [31], as shown in Figure 3 (a). The value obtained is in close agreement with the literature ($\lambda_{max} \approx 0.017$ [62]).

Figure 3 shows that also for these empirical data the optimal training interval τ agrees well with the predictability time, eqn. (17), for all trained RNNs. Furthermore, as was the case for the DS benchmarks, gradient clipping was not able to satisfactorily tackle the EVGP, even when paired with architectures like LSTMs explicitly designed for alleviating this problem. Similar results are reported for another real-world dataset, electroencephalogram (EEG) recordings, in Appx. A.A.6.5.

5 Discussion and conclusions

In this paper we proved that RNN dynamics and loss gradients are intimately related for all major types of RNNs and activation functions. If the RNN is "well behaved" in the sense that its dynamics converges to a fixed point or cycle, loss gradients will remain bounded, and established remedies [36, 80] can be used to refrain them from vanishing. However, if the dynamics are chaotic, gradients will always explode. This constitutes a *principle* problem in RNN training that cannot easily be mastered through architectural design or gradient clipping. It is furthermore a practically highly relevant one, as most time series we encounter in nature, and many from man-made systems as well, are inherently chaotic. While we do not offer a full solution to this problem here, we suggest it might be tackled in training by taking a system's local divergence rates as measured through the Lyapunov spectrum into account. Hence, rather than conquering the EVGP by structural design or specific constraints or regularization terms, we recommend to put the focus more on the training process itself. We illustrated this point empirically using *sparsely forced BPTT*, a training technique that pulls trajectories back on track at times determined by the maximal Lyapunov exponent. Doing so leads to optimal reconstruction results for a variety of simulated and real-world benchmarks, regardless of the specific RNN architecture employed in training.

We stress that our goal here above all was to provide a mathematically grounded perspective on the problem, with the empirical section focused on elucidating the practical implications of the theoretical results. We believe that a more thorough theoretical understanding is important and needed for guiding future research into more powerful training procedures that avoid exploding gradients *without* compromising expressiveness. In our application examples, we developed the case from the perspective of scientific machine learning, which by now is a broad area in its own right with huge societal relevance (e.g., climate or epidemiological time series), and where the reconstruction of geometrical or topological (invariant) properties is important, beyond mere prediction. Nevertheless, we believe that our theoretical results will also have implications for other domains, like NLP [37]. While scientific time series problems traditionally have been extensively considered from a DS perspective (e.g., [43]), much more groundwork is needed, however, in areas like NLP, where, for instance, it may not even be immediately clear how to best define a Lyapunov spectrum.

References

[1] H. D. I. Abarbanel. *Predicting the Future: Completing Models of Observed Complex Systems*. en. Understanding Complex Systems. New York: Springer-Verlag, 2013. DOI: 10.1007/978-1-4614-7218-6. URL: https://www.springer.com/gp/book/9781461472179 (visited on 09/13/2021).

[2] H. D. I. Abarbanel, P. J. Rozdeba, and S. Shirman. "Machine Learning: Deepest Learning as Statistical Data Assimilation Problems". eng. In: *Neural Computation* 30.8 (Aug. 2018), pp. 2025–2055. DOI: 10.1162/neco_a_01094.

[3] K. T. Alligood, T. D. Sauer, and J. A. Yorke. *Chaos: An Introduction to Dynamical Systems*. Springer, New York, NY, 1996.

[4] M. Arjovsky, A. Shah, and Y. Bengio. "Unitary Evolution Recurrent Neural Networks". In: *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*. ICML'16. New York, NY, USA: JMLR.org, 2016, pp. 1120–1128.

[5] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer. "Scheduled sampling for sequence prediction with recurrent Neural networks". In: *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*. NIPS'15. Cambridge, MA, USA: MIT Press, Dec. 2015, pp. 1171–1179. (Visited on 09/07/2021).

[6] Y. Bengio, P. Simard, and P. Frasconi. "Learning long-term dependencies with gradient descent is difficult". In: *IEEE Transactions on Neural Networks* 5.2 (1994), pp. 157–166. DOI: 10.1109/72.279181.

[7] B. P. Bezruchko and D. A. Smirnov. *Extracting Knowledge From Time Series: An Introduction to Nonlinear Empirical Modeling*. en. Springer Series in Synergetics. Berlin Heidelberg: Springer-Verlag, 2010. DOI: 10.1007/978-3-642-12601-7. URL: https://www.springer.com/gp/book/9783642126000.

[8] P. L. Bommer, D. Kramer, C. Tombolini, G. Koppe, and D. Durstewitz. *Identifying nonlinear dynamical systems from multi-modal time series data*. 2021. arXiv: 2111.02922 [cs.LG].

[9] M. Brenner, F. Hess, J. M. Mikhaeil, L. F. Bereska, Z. Monfared, P.-C. Kuo, and D. Durstewitz. "Tractable Dendritic RNNs for Reconstructing Nonlinear Dynamical Systems". In: *Proceedings of the 39th International Conference on Machine Learning*. Ed. by K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato. Vol. 162. Proceedings of Machine Learning Research. PMLR, 2022, pp. 2292–2320.

[10] B. Chang, M. Chen, E. Haber, and E. H. Chi. "AntisymmetricRNN: A Dynamical System View on Recurrent Neural Networks". In: *International Conference on Learning Representations*. 2019. URL: https://openreview.net/forum?id=ryxepo0cFX.

[11] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio. "On the Properties of Neural Machine Translation: Encoder–Decoder Approaches". In: *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Doha, Qatar: Association for Computational Linguistics, Oct. 2014, pp. 103–111. DOI: 10.3115/v1/W14-4012. URL: https://aclanthology.org/W14-4012.

[12] K. Cho, B. van Merriënboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio. "Learning phrase representations using RNN encoder-decoder for statistical machine translation". English (US). In: *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*. 2014.

[13] G. Datseris. "DynamicalSystems.jl: A Julia software library for chaos and nonlinear dynamics". In: *Journal of Open Source Software* 3.23 (2018), p. 598. DOI: 10.21105/joss.00598. URL: https://doi.org/10.21105/joss.00598.

[14] K. Doya. "Bifurcations in the learning of recurrent neural networks". In: *[Proceedings] 1992 IEEE International Symposium on Circuits and Systems*. Vol. 6. May 1992, 2777–2780 vol.6. DOI: 10.1109/ISCAS.1992.230622.

[15] J. Duarte, C. Januário, N. Martins, and J. Sardanyés. "Quantifying chaos for ecological stoichiometry". In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 20.3 (Sept. 2010). Publisher: American Institute of Physics, p. 033105. DOI: 10.1063/1.3464327. URL: https://aip.scitation.org/doi/full/10.1063/1.3464327.

[16] G. Duffing. *Erzwungene Schwingungen bei Veränderlicher Eigenfrequenz*. 1918.

[17] D. Durstewitz. *Advanced Data Analysis in Neuroscience: Integrating Statistical and Computational Models*. en. Bernstein Series in Computational Neuroscience. Springer International Publishing, 2017. DOI: 10.1007/978-3-319-59976-2. URL: https://www.springer.com/de/book/9783319599748 (visited on 09/10/2021).

[18] D. Durstewitz and T. Gabriel. "Dynamical Basis of Irregular Spiking in NMDA-Driven Prefrontal Cortex Neurons". en. In: *Cerebral Cortex* 17.4 (Apr. 2007), pp. 894–908. DOI: 10.1093/cercor/bhk044. URL: https://academic.oup.com/cercor/article-lookup/doi/10.1093/cercor/bhk044.

[19] R. Engelken, F. Wolf, and L. F. Abbott. "Lyapunov spectra of chaotic recurrent neural networks". In: *arXiv:2006.02427 [nlin, q-bio]* (June 2020). arXiv: 2006.02427. URL: http://arxiv.org/abs/2006.02427 (visited on 09/26/2021).

[20] N. B. Erichson, O. Azencot, A. Queiruga, L. Hodgkinson, and M. W. Mahoney. "Lipschitz Recurrent Neural Networks". In: *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=-N7PBXqOUJZ.

[21] M. Faggini. "Chaotic time series analysis in economics: Balance and perspectives". In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 24.4 (Dec. 2014). Publisher: American Institute of Physics, p. 042101. DOI: 10.1063/1.4903797. URL: https://aip.scitation.org/doi/full/10.1063/1.4903797.

[22] R. J. Field and L. Györgyi. *Chaos in Chemistry and Biochemistry*. WORLD SCIENTIFIC, 1993. DOI: 10.1142/1706. eprint: https://www.worldscientific.com/doi/pdf/10.1142/1706. URL: https://www.worldscientific.com/doi/abs/10.1142/1706.

[23] S. Fusi, W. F. Asaad, E. K. Miller, and X.-J. Wang. "A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales". eng. In: *Neuron* 54.2 (Apr. 2007), pp. 319–333. DOI: 10.1016/j.neuron.2007.03.017.

[24] Z. Ghahramani and S. Roweis. "Learning Nonlinear Dynamical Systems Using an EM Algorithm". In: *Advances in Neural Information Processing Systems*. Ed. by M. Kearns, S. Solla, and D. Cohn. Vol. 11. MIT Press, 1999. URL: https://proceedings.neurips.cc/paper/1998/file/0ebcc77dc72360d0eb8e9504c78d38bd-Paper.pdf.

[25] W. 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. URL: https://openreview.net/forum?id=enYjtbjYJrf.

[26] L. Glass and M. C. Mackey. "Pathological conditions resulting from instabilities in physiological control systems". eng. In: *Annals of the New York Academy of Sciences* 316 (1979), pp. 214–235. DOI: 10.1111/j.1749-6632.1979.tb29471.x.

[27] P. A. Glendinning and D. J. W. Simpson. "A constructive approach to robust chaos using invariant manifolds and expanding cones". In: *Discrete & Continuous Dynamical Systems* 41.7 (2021), pp. 3367–3387.

[28] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley. "PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals". In: *Circulation* 101.23 (2000). Circulation Electronic Pages: http://circ.ahajournals.org/content/101/23/e215.full PMID:1085218; doi: 10.1161/01.CIR.101.23.e215, e215–e220.

[29] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. http://www.deeplearningbook.org. MIT Press, 2016.

[30] A. Graves, G. Wayne, M. Reynolds, and et al. "Hybrid computing using a neural network with dynamic external memory". In: *Nature* 538 (2016), 471—476. DOI: 10.1038/nature20101.

[31] R. Hegger, H. Kantz, and T. Schreiber. "Practical implementation of nonlinear time series methods: The TISEAN package". In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 9.2 (June 1999). Publisher: American Institute of Physics, pp. 413–435. DOI: 10.1063/1.166424. URL: https://aip.scitation.org/doi/citedby/10.1063/1.166424 (visited on 09/07/2021).

[32] K. Helfrich, D. Willmott, and Q. Ye. "Orthogonal Recurrent Neural Networks with Scaled Cayley Transform". en. In: *International Conference on Machine Learning*. ISSN: 2640-3498. PMLR, July 2018, pp. 1969–1978. URL: http://proceedings.mlr.press/v80/helfrich18a.html.

[33] M. Henaff, A. Szlam, and Y. LeCun. "Recurrent Orthogonal Networks and Long-Memory Tasks". en. In: *International Conference on Machine Learning*. ISSN: 1938-7228. PMLR, June 2016, pp. 2034–2042. URL: http://proceedings.mlr.press/v48/henaff16.html (visited on 07/21/2021).

[34] D. Hernandez, A. K. Moretti, Z. Wei, S. Saxena, J. Cunningham, and L. Paninski. "Nonlinear Evolution via Spatially-Dependent Linear Dynamics for Electrophysiology and Calcium Data". en. In: *arXiv preprint arXiv:1811.02459* (2020). URL: http://arxiv.org/abs/1811.02459 (visited on 07/23/2020).

[35] S. Hochreiter. *Untersuchungen zu dynamischen neuronalen Netzen*. 1991.

[36] S. Hochreiter and J. Schmidhuber. "Long Short-Term Memory". In: *Neural Computation* 9.8 (Nov. 1997), pp. 1735–1780. DOI: 10.1162/neco.1997.9.8.1735. URL: https://doi.org/10.1162/neco.1997.9.8.1735 (visited on 07/05/2021).

[37] K. Inoue, S. Ohara, Y. Kuniyoshi, and K. Nakajima. "Transient Chaos in BERT". In: *arXiv:2106.03181 [nlin]* (June 2021). arXiv: 2106.03181. URL: http://arxiv.org/abs/2106.03181.

[38] L. Jing, C. Gulcehre, J. Peurifoy, Y. Shen, M. Tegmark, M. Soljacic, and Y. Bengio. "Gated Orthogonal Recurrent Units: On Learning to Forget". eng. In: *Neural Computation* 31.4 (Apr. 2019), pp. 765–783. DOI: 10.1162/neco_a_01174.

[39] A. Kag and V. Saligrama. "Time Adaptive Recurrent Neural Network". en. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Nashville, TN, USA: IEEE, June 2021, pp. 15144–15153. DOI: 10.1109/CVPR46437.2021.01490. URL: https://ieeexplore.ieee.org/document/9578651/ (visited on 11/14/2021).

[40] A. Kag and V. Saligrama. "Training Recurrent Neural Networks via Forward Propagation Through Time". In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by M. Meila and T. Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 2021, pp. 5189–5200. URL: https://proceedings.mlr.press/v139/kag21a.html.

[41] A. Kag, Z. Zhang, and V. Saligrama. "RNNs Incrementally Evolving on an Equilibrium Manifold: A Panacea for Vanishing and Exploding Gradients?" In: *International Conference on Learning Representations*. 2020. URL: https://openreview.net/forum?id=HylpqA4FwS.

[42] L. Kamdjeu Kengne, J. R. Mboupda Pone, and H. B. Fotsin. "On the dynamics of chaotic circuits based on memristive diode-bridge with variable symmetry: A case study". en. In: *Chaos, Solitons & Fractals* 145 (Apr. 2021), p. 110795. DOI: 10.1016/j.chaos.2021.110795. URL: https://www.sciencedirect.com/science/article/pii/S0960077921001478.

[43] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. 2nd ed. Cambridge University Press, 2003. DOI: 10.1017/CBO9780511755798.

[44] H. Kantz, T. Schreiber, I. Hoffmann, T. Buzug, G. Pfister, L. G. Flepp, J. Simonet, R. Badii, and E. Brun. "Nonlinear noise reduction: A case study on experimental data". In: *Physical Review E* 48.2 (Aug. 1993). Publisher: American Physical Society, pp. 1529–1538. DOI: 10.1103/PhysRevE.48.1529. URL: https://link.aps.org/doi/10.1103/PhysRevE.48.1529 (visited on 09/10/2021).

[45] M. B. Kennel, R. Brown, and H. D. I. Abarbanel. "Determining embedding dimension for phase-space reconstruction using a geometrical construction". In: *Phys. Rev. A* 45 (6 1992), pp. 3403–3411. DOI: 10.1103/PhysRevA.45.3403. URL: https://link.aps.org/doi/10.1103/PhysRevA.45.3403.

[46] G. Kerg, K. Goyette, M. Puelma Touzel, G. Gidel, E. Vorontsov, Y. Bengio, and G. Lajoie. "Non-normal Recurrent Neural Network (nnRNN): learning long time dependencies while improving expressivity with transient dynamics". In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett. Vol. 32. Curran Associates, Inc., 2019. URL: https://proceedings.neurips.cc/paper/2019/file/9d7099d87947faa8d07a272dd6954b80-Paper.pdf.

[47] M. Kesmia, S. Boughaba, and S. Jacquir. "Control of continuous dynamical systems modeling physiological states". en. In: *Chaos, Solitons & Fractals* 136 (July 2020), p. 109805. DOI: 10.1016/j.chaos.2020.109805. URL: https://www.sciencedirect.com/science/article/pii/S096007792030206X.

[48] D. P. Kingma and J. Ba. "Adam: A Method for Stochastic Optimization". In: *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. Ed. by Y. Bengio and Y. LeCun. 2015. URL: http://arxiv.org/abs/1412.6980.

[49] D. P. Kingma and M. Welling. "Auto-Encoding Variational Bayes". en. In: *Proceedings of the 2nd International Conference on Learning Representations*. 2014. URL: http://arxiv.org/abs/1312.6114 (visited on 01/29/2020).

[50] J. Z. Kolter and G. Manek. "Learning Stable Deep Dynamics Models". In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett. Vol. 32. Curran Associates, Inc., 2019. URL: https://proceedings.neurips.cc/paper/2019/file/0a4bbceda17a6253386bc9eb45240e25-Paper.pdf.

[51] G. Koppe, H. Toutounji, P. Kirsch, S. Lis, and D. Durstewitz. "Identifying nonlinear dynamical systems via generative recurrent neural networks with applications to fMRI". In: *PLOS Computational Biology* 15.8 (Aug. 2019), pp. 1–35. DOI: 10.1371/journal.pcbi.1007263. URL: https://doi.org/10.1371/journal.pcbi.1007263.

[52] J. Laskar and P. Robutel. "The chaotic obliquity of the planets". en. In: *Nature* 361.6413 (Feb. 1993), pp. 608–612. DOI: 10.1038/361608a0. URL: https://www.nature.com/articles/361608a0 (visited on 09/10/2021).

[53] S. Li, W. Li, C. Cook, C. Zhu, and Y. Gao. "Independently Recurrent Neural Network (IndRNN): Building A Longer and Deeper RNN". In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2018, pp. 5457–5466. DOI: 10.1109/CVPR.2018.00572.

[54] Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. "Fourier Neural Operator for Parametric Partial Differential Equations". In: *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=c8P9NQVtmnO.

[55] Z. C. Lipton, J. Berkowitz, and C. Elkan. "A Critical Review of Recurrent Neural Networks for Sequence Learning". In: *arXiv:1506.00019 [cs]* (Oct. 2015). arXiv: 1506.00019. URL: http://arxiv.org/abs/1506.00019 (visited on 09/12/2021).

[56] E. N. Lorenz. "Deterministic Nonperiodic Flow". EN. In: *Journal of the Atmospheric Sciences* 20.2 (Mar. 1963). Publisher: American Meteorological Society Section: Journal of the Atmospheric Sciences, pp. 130–141. DOI: 10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2. URL: https://journals.ametsoc.org/view/journals/atsc/20/2/1520-0469_1963_020_0130_dnf_2_0_co_2.xml.

[57] N. Maheswaranathan, A. Williams, M. Golub, S. Ganguli, and D. Sussillo. "Reverse engineering recurrent networks for sentiment classification reveals line attractor dynamics". In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett. Vol. 32. Curran Associates, Inc., 2019. URL: https://proceedings.neurips.cc/paper/2019/file/d921c3c762b1522c475ac8fc0811bb0f-Paper.pdf.

[58] S. Mangiarotti, M. Peyre, Y. Zhang, M. Huc, F. Roger, and Y. Kerr. "Chaos theory applied to the outbreak of COVID-19: an ancillary approach to decision making in pandemic context". In: *Epidemiology and Infection* 148 (May 2020), e95. DOI: 10.1017/S0950268820000990. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7231667/.

[59] R. M. May. "Chaos and the Dynamics of Biological Populations". In: *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* 413.1844 (1987), pp. 27–44. URL: http://www.jstor.org/stable/2398225.

[60] J. D. Meiss. *Differential Dynamical Systems*. Society for Industrial and Applied Mathematics, 2007.

[61] J. Menick, E. Elsen, U. Evci, S. Osindero, K. Simonyan, and A. Graves. "Practical Real Time Recurrent Learning with a Sparse Approximation". In: *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=q3KSThy2GwB.

[62] H. Millán, B. Ghanbarian-Alavijeh, and I. García-Fornaris. "Nonlinear dynamics of mean daily temperature and dewpoint time series at Babolsar, Iran, 1961–2005". en. In: *Atmospheric Research*. Clouds, Aerosols and Radiation 98.1 (Oct. 2010), pp. 89–101. DOI: 10.1016/j.atmosres.2010.06.001. URL: https://www.sciencedirect.com/science/article/pii/S0169809510001419 (visited on 06/12/2021).

[63] Z. Monfared and D. Durstewitz. "Existence of n-cycles and border-collision bifurcations in piecewise-linear continuous maps with applications to recurrent neural networks". en. In: *Nonlinear Dynamics* 101.2 (July 2020), pp. 1037–1052. DOI: 10.1007/s11071-020-05841-x. URL: http://link.springer.com/10.1007/s11071-020-05841-x (visited on 07/15/2021).

[64] Z. Monfared and D. Durstewitz. "Transformation of ReLU-based recurrent neural networks from discrete-time to continuous-time". en. In: *International Conference on Machine Learning*. ISSN: 2640-3498. PMLR, Nov. 2020, pp. 6999–7009. URL: http://proceedings.mlr.press/v119/monfared20a.html (visited on 07/15/2021).

[65] L. F. Olsen and H. Degn. "Chaos in an enzyme reaction". en. In: *Nature* 267.5607 (May 1977), pp. 177–178. DOI: 10.1038/267177a0. URL: https://www.nature.com/articles/267177a0 (visited on 09/12/2021).

[66] E. Orhan and X. Pitkow. "Improved memory in recurrent neural networks with sequential non-normal dynamics". In: *International Conference on Learning Representations*. 2020. URL: https://openreview.net/forum?id=ryx1wRNFvB.

[67] R. Pascanu, T. Mikolov, and Y. Bengio. "On the Difficulty of Training Recurrent Neural Networks". In: *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*. ICML'13. Atlanta, GA, USA: JMLR.org, 2013, III–1310–III–1318.

[68] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. "Automatic differentiation in PyTorch". In: 2017. URL: https://openreview.net/forum?id=BJJsrmfCZ.

[69] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott. "Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach". In: *Phys. Rev. Lett.* 120 (2 2018), p. 024102. DOI: 10.1103/PhysRevLett.120.024102.

[70] B. Pearlmutter. *Dynamic recurrent neural networks*. 1990. DOI: 10.1184/R1/6605018.v1. URL: https://kilthub.cmu.edu/articles/journal_contribution/Dynamic_recurrent_neural_networks/6605018/1.

[71] L. Perko. *Differential Equations and Dynamical Systems*. Vol. 7. Springer, New York, NY, 2001.

[72] M. T. Rosenstein, J. J. Collins, and C. J. De Luca. "A practical method for calculating largest Lyapunov exponents from small data sets". en. In: *Physica D: Nonlinear Phenomena* 65.1 (May 1993), pp. 117–134. DOI: 10.1016/0167-2789(93)90009-P. URL: https://www.sciencedirect.com/science/article/pii/016727899390009P (visited on 06/11/2021).

[73] T. K. Rusch and S. Mishra. "Coupled Oscillatory Recurrent Neural Network (coRNN): An accurate and (gradient) stable architecture for learning long time dependencies". In: *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=F3s69XzWOia.

[74] T. K. Rusch and S. Mishra. "UnICORNN: A recurrent model for learning very long time dependencies". In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by M. Meila and T. Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 2021, pp. 9168–9178. URL: https://proceedings.mlr.press/v139/rusch21a.html.

[75] E. Russo and D. Durstewitz. "Cell assemblies at multiple time scales with arbitrary lag constellations". In: *eLife* 6 (Jan. 2017). Ed. by M. Howard. Publisher: eLife Sciences Publications, Ltd, e19428. DOI: 10.7554/eLife.19428. URL: https://doi.org/10.7554/eLife.19428 (visited on 09/16/2021).

[76] O. E. Rössler. "An equation for continuous chaos". en. In: *Physics Letters A* 57.5 (July 1976), pp. 397–398. DOI: 10.1016/0375-9601(76)90101-8. URL: https://www.sciencedirect.com/science/article/pii/0375960176901018 (visited on 09/14/2021).

[77] S. Sabanal and M. Nakagawa. "The fractal properties of vocal sounds and their application in the speech recognition model". en. In: *Chaos, Solitons & Fractals* 7.11 (Nov. 1996), pp. 1825–1843. DOI: 10.1016/S0960-0779(96)00043-4. URL: https://www.sciencedirect.com/science/article/pii/S0960077996000434.

[78] T. Sauer, J. A. Yorke, and M. Casdagli. "Embedology". en. In: *Journal of Statistical Physics* 65.3 (Nov. 1991), pp. 579–616. DOI: 10.1007/BF01053745. URL: https://doi.org/10.1007/BF01053745 (visited on 07/14/2021).

[79] G. Schalk, D. J. McFarland, T. Hinterberger, N. Birbaumer, and J. R. Wolpaw. "BCI2000: a general-purpose brain-computer interface (BCI) system". eng. In: *IEEE transactions on biomedical engineering* 51.6 (June 2004), pp. 1034–1043. DOI: 10.1109/TBME.2004.827072.

[80] D. Schmidt, G. Koppe, Z. Monfared, M. Beutelspacher, and D. Durstewitz. "Identifying nonlinear dynamical systems with multiple time scales and long-range dependencies". In: *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=_XYzwxPIQu6.

[81] B. Sivakumar. "Chaos theory in geophysics: past, present and future". en. In: *Chaos, Solitons & Fractals*. Fractals in Geophysics 19.2 (Jan. 2004), pp. 441–462. DOI: 10.1016/S0960-0779(03)00055-9. URL: https://www.sciencedirect.com/science/article/pii/S0960077903000559.

[82] F. Takens. "Detecting strange attractors in turbulence". In: *Dynamical Systems and Turbulence, Warwick 1980*. Ed. by D. Rand and L.-S. Young. Berlin, Heidelberg: Springer Berlin Heidelberg, 1981, pp. 366–381.

[83] R. Tchitnga, B. A. Mezatio, T. Fonzin Fozin, R. Kengne, P. H. Louodop Fotso, and A. Fomethe. "A novel hyperchaotic three-component oscillator operating at high frequency". en. In: *Chaos, Solitons & Fractals* 118 (Jan. 2019), pp. 166–180. DOI: 10.1016/j.chaos.2018.11.015. URL: https://www.sciencedirect.com/science/article/pii/S0960077918303047.

[84] D. J. Thomson. "Time series analysis of Holocene climate data". In: *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 330.1615 (Apr. 1990). Publisher: Royal Society, pp. 601–616. DOI: 10.1098/rsta.1990.0041. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rsta.1990.0041 (visited on 09/12/2021).

[85] T. H. Trinh, A. M. Dai, M.-T. Luong, and Q. V. Le. *Learning Longer-term Dependencies in RNNs with Auxiliary Losses*. 2018. URL: https://openreview.net/forum?id=Hy9xDwyPM.

[86] A. P. Trischler and G. M. D'Eleuterio. "Synthesis of recurrent neural networks for dynamical system simulation". en. In: *Neural Networks* 80 (2016), pp. 67–78. DOI: 10.1016/j.neunet.2016.04.001. URL: https://linkinghub.elsevier.com/retrieve/pii/S0893608016300314 (visited on 03/08/2019).

[87] P. Turchin and A. D. Taylor. "Complex Dynamics in ecological Time Series". en. In: *Ecology* 73.1 (1992), pp. 289–305. DOI: 10.2307/1938740. (Visited on 09/12/2021).

[88] E. Tziperman, H. Scher, S. E. Zebiak, and M. A. Cane. "Controlling Spatiotemporal Chaos in a Realistic El Niño Prediction Model". In: *Physical Review Letters* 79.6 (Aug. 1997). Publisher: American Physical Society, pp. 1034–1037. DOI: 10.1103/PhysRevLett.79.1034. URL: https://link.aps.org/doi/10.1103/PhysRevLett.79.1034.

[89] P. R. Vlachas, G. Arampatzis, C. Uhler, and P. Koumoutsakos. "Learning the Effective Dynamics of Complex Multiscale Systems". In: *arXiv:2006.13431 [nlin, physics:physics]* (July 2020). arXiv: 2006.13431. URL: http://arxiv.org/abs/2006.13431 (visited on 07/19/2021).

[90] P. R. Vlachas, W. Byeon, Z. Y. Wan, T. P. Sapsis, and P. Koumoutsakos. "Data-driven forecasting of high-dimensional chaotic systems with long short-term memory networks". en. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 474.2213 (2018), p. 20170844. DOI: 10.1098/rspa.2017.0844. URL: https://royalsocietypublishing.org/doi/10.1098/rspa.2017.0844 (visited on 11/17/2020).

[91] R. Vogt, M. Puelma Touzel, E. Shlizerman, and G. Lajoie. "On Lyapunov Exponents for RNNs: Understanding Information Propagation Using Dynamical Systems Tools". In: *Frontiers in Applied Mathematics and Statistics* 8 (2022). DOI: 10.3389/fams.2022.818799. URL: https://www.frontiersin.org/article/10.3389/fams.2022.818799.

[92] H. U. Voss, J. Timmer, and J. Kurths. "Nonlinear dynamical system identification from uncertain and indirect measurements". en. In: *International Journal of Bifurcation and Chaos* 14.06 (June 2004), pp. 1905–1933. DOI: 10.1142/S0218127404010345. URL: https://www.worldscientific.com/doi/abs/10.1142/S0218127404010345.

[93] C. van Vreeswijk and H. Sompolinsky. "Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity". en. In: *Science* 274.5293 (Dec. 1996). Publisher: American Association for the Advancement of Science Section: Reports, pp. 1724–1726. DOI: 10.1126/science.274.5293.1724. URL: https://science.sciencemag.org/content/274/5293/1724.

[94] J. H. M. Wedderburn. *Lectures on Matrices*. New York: American mathematical society, New York : Dover Publications, 1964.

[95] R. J. Williams and D. Zipser. "A Learning Algorithm for Continually Running Fully Recurrent Neural Networks". en. In: *Neural Computation* 1.2 (June 1989), pp. 270–280. DOI: 10.1162/neco.1989.1.2.270. URL: https://direct.mit.edu/neco/article/1/2/270-280/5490.

[96] S. N. Wood. "Statistical inference for noisy nonlinear ecological dynamic systems". en. In: *Nature* 466.7310 (Aug. 2010). DOI: 10.1038/nature09319. URL: https://www.nature.com/articles/nature09319 (visited on 07/14/2021).

[97] H. Zhang, D. Liu, and Z. Wang. *Controlling Chaos*. Springer, London, 2009.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Sec. 5.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] We provide comprehensive proofs of all presented theorems in the Appendix.
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code will be made public after publication.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]