

# UnICORN: A recurrent model for learning very long time dependencies

T. K. Rusch and S. Mishra

Research Report No. 2021-10  
March 2021

Seminar für Angewandte Mathematik  
Eidgenössische Technische Hochschule  
CH-8092 Zürich  
Switzerland

# UnICORN: A recurrent model for learning *very* long time dependencies

T. Konstantin Rusch \*      Siddhartha Mishra \*

## Abstract

The design of recurrent neural networks (RNNs) to accurately process sequential inputs with long-time dependencies is very challenging on account of the exploding and vanishing gradient problem. To overcome this, we propose a novel RNN architecture which is based on a structure preserving discretization of a Hamiltonian system of second-order ordinary differential equations that models networks of oscillators. The resulting RNN is fast, invertible (in time), memory efficient and we derive rigorous bounds on the hidden state gradients to prove the mitigation of the exploding and vanishing gradient problem. A suite of experiments are presented to demonstrate that the proposed RNN provides state of the art performance on a variety of learning tasks with (very) long time-dependencies.

## 1 Introduction

Recurrent Neural Networks (RNNs) have been very successful in solving a diverse set of learning tasks involving sequential inputs [30]. These include text and speech recognition, time-series analysis and natural language processing. However, the well-known *Exploding and Vanishing Gradient Problem* (EVGP) [38] and references therein, impedes the efficiency of RNNs on tasks that require processing (very) long sequential inputs. The EVGP arises from the fact that the backpropagation through time (BPTT) algorithm for training RNNs entails computing products of hidden state gradients over a large number of steps and this product can either be exponentially small or large as the number of recurrent interactions increases.

Different approaches to solve the EVGP has been suggested in recent years. These include the use of gating mechanisms, such as in LSTMs [23] and GRUs [12], where the additive structure of the gates mitigates the vanishing gradient problem. However, gradients might still explode, impeding the efficiency of LSTMs and GRUs on problems with very long time dependencies (LTDs) [32]. The EVGP can also be mitigated by constraining the structure of the recurrent weight matrices, for instance requiring them to be orthogonal or unitary [22, 2, 46, 25]. Constraining recurrent weight matrices may lead to a loss of expressivity of the resulting RNN, reducing its efficiency in handling realistic learning tasks [25]. Finally, restricting weights of the RNN to lie within some prespecified bounds might lead to control over the norms of the recurrent weight matrices and alleviate the EVGP. Such an approach has been suggested in the context of *independent neurons* in each layer in [32], and using a coupled system of damped oscillators in [41], among others. However, ensuring that weights remain within a pre-defined range during training might be difficult. Furthermore, *weight clipping* could also reduce expressivity of the resulting RNN.

In addition to EVGP, the learning of sequential tasks with very long time dependencies can require significant computational resources, for training and evaluating the RNN. Moreover, as the BPTT training algorithms entail storing all hidden states at every time step, the overall memory requirements can be prohibitive. Thus, *the design of a fast and memory efficient RNN architecture that can mitigate the EVGP is highly desirable for the effective use of RNNs in realistic learning tasks with very long time dependencies.* The main objective of this article is to propose, analyze and test such an architecture.

The basis of our proposed RNN is the observation that a large class of dynamical systems in physics and engineering, the so-called *Hamiltonian systems* [3], allow for very precise control on the underlying

---

\*Seminar for Applied Mathematics (SAM), D-MATH  
ETH Zürich, Rämistrasse 101, Zürich-8092, Switzerland

states. Moreover, the fact that the phase space volume is preserved by the trajectories of a Hamiltonian system, makes such systems *invertible* and allows one to significantly reduce the storage requirements. Furthermore, if the resulting hidden state gradients also evolve according to a Hamiltonian dynamical system, one can obtain precise bounds on the hidden state gradients and alleviate the EVGP. We combine and extend these ideas into an RNN architecture that will allow us to prove rigorous bounds on the hidden states and their gradients, mitigating the EVGP. Moreover, our RNN architecture results in a fast implementation that attains state of the art performance on a variety of learning tasks with very long time dependencies.

## 2 The proposed RNN

Our proposed RNN is based on the time-discretization of the following system of *second-order ordinary differential equations* (ODEs),

$$\mathbf{y}'' = -[\sigma(\mathbf{w} \odot \mathbf{y} + \mathbf{V}\mathbf{u} + \mathbf{b}) + \alpha\mathbf{y}]. \quad (1)$$

Here,  $t \in [0, 1]$  is the (continuous) time variable,  $\mathbf{u} = \mathbf{u}(t) \in \mathbb{R}^d$  is the time-dependent *input signal*,  $\mathbf{y} = \mathbf{y}(t) \in \mathbb{R}^m$  is the *hidden state* of the RNN with  $\mathbf{w} \in \mathbb{R}^m$  is a weight vector,  $\mathbf{V} \in \mathbb{R}^{m \times d}$  a weight matrix,  $\mathbf{b} \in \mathbb{R}^m$  is the bias vector and  $\alpha \geq 0$  is a control parameter. The operation  $\odot$  is the Hadamard product and the function  $\sigma : \mathbb{R} \mapsto \mathbb{R}$  is the *activation function* and is applied component wise. For the rest of this paper, we set  $\sigma(u) = \tanh(u)$ .

By introducing the auxiliary variable  $\mathbf{z} = \mathbf{y}'$ , we can rewrite the second order ODE (1) as a first order ODE system:

$$\mathbf{y}' = \mathbf{z}, \quad \mathbf{z}' = -[\sigma(\mathbf{w} \odot \mathbf{y} + \mathbf{V}\mathbf{u} + \mathbf{b}) + \alpha\mathbf{y}]. \quad (2)$$

Assuming that  $\mathbf{w}_i \neq 0$ , for all  $1 \leq i \leq m$ , it is easy to see that the ODE system (2) is a *Hamiltonian system*,

$$\mathbf{y}' = \frac{\partial H}{\partial \mathbf{z}}, \quad \mathbf{z}' = -\frac{\partial H}{\partial \mathbf{y}}, \quad (3)$$

with the *time-dependent Hamiltonian*,

$$H(\mathbf{y}, \mathbf{z}, t) = \frac{\alpha}{2} \|\mathbf{y}\|^2 + \frac{1}{2} \|\mathbf{z}\|^2 + \sum_{i=1}^m \frac{1}{\mathbf{w}_i} \log(\cosh(\mathbf{w}_i \mathbf{y}_i + (\mathbf{V}\mathbf{u}(t))_i + \mathbf{b}_i)), \quad (4)$$

with  $\|\mathbf{x}\|^2 = \langle \mathbf{x}, \mathbf{x} \rangle$  denoting the Euclidean norm of the vector  $\mathbf{x} \in \mathbb{R}^m$  and  $\langle \cdot, \cdot \rangle$  the corresponding inner product.

The next step is to find a discretization of the ODE system (2). Given that it is highly desirable to ensure that the discretization respects the Hamiltonian structure of the underlying continuous ODE, the simplest such *structure preserving discretization* is the *symplectic Euler* method [42, 19]. Applying the symplectic Euler method to the ODE (2) results in the following discrete dynamical system,

$$\begin{aligned} \mathbf{y}_n &= \mathbf{y}_{n-1} + \Delta t \mathbf{z}_n, \\ \mathbf{z}_n &= \mathbf{z}_{n-1} - \Delta t [\sigma(\mathbf{w} \odot \mathbf{y}_{n-1} + \mathbf{V}\mathbf{u}_n + \mathbf{b}) + \alpha \mathbf{y}_{n-1}], \end{aligned} \quad (5)$$

for  $1 \leq n \leq N$ . Here,  $0 < \Delta t < 1$  is the time-step and  $\mathbf{u}_n \approx \mathbf{u}(t_n)$ , with  $t_n = n\Delta t$ , is the input signal. It is common to initialize with  $\mathbf{y}_0 = \mathbf{z}_0 = \mathbf{0}$ .

We see from the structure of the discrete dynamical system (5) that there is *no interaction* between the neurons in the hidden layer of (5). Such an RNN will have very limited expressivity. Hence, we *stack* more hidden layers to propose the following deep or *multilayer* RNN,

$$\begin{aligned} \mathbf{y}_n^\ell &= \mathbf{y}_{n-1}^\ell + \Delta t \hat{\sigma}(\mathbf{c}^\ell) \odot \mathbf{z}_n^\ell, \\ \mathbf{z}_n^\ell &= \mathbf{z}_{n-1}^\ell - \Delta t \hat{\sigma}(\mathbf{c}^\ell) \odot [\sigma(\mathbf{w}^\ell \odot \mathbf{y}_{n-1}^\ell + \mathbf{V}^\ell \mathbf{y}_{n-1}^{\ell-1} + \mathbf{b}^\ell) + \alpha \mathbf{y}_{n-1}^\ell]. \end{aligned} \quad (6)$$

Here  $\mathbf{y}_n^\ell, \mathbf{z}_n^\ell \in \mathbb{R}^m$  are hidden states and  $\mathbf{w}^\ell, \mathbf{V}^\ell, \mathbf{b}^\ell$  are weights and biases, corresponding to layer  $\ell = 1, \dots, L$ . We set  $\mathbf{y}_n^0 = \mathbf{u}_n$  in the multilayer RNN (6).

Observe that we use the same step-size  $\Delta t$  for every layer, while multiplying a trainable parameter vector  $\mathbf{c} \in \mathbb{R}^m$  to the time step. The action of  $\mathbf{c}$  is modulated with the sigmoidal activation function  $\hat{\sigma}(u) = 0.5 + 0.5 \tanh(u/2)$ , which ensures that the time-step  $\Delta t$  is multiplied by a value between 0 and 1. We remark that the presence of this trainable vector  $\mathbf{c}$  allows us to incorporate *multi-scale behavior* in the proposed RNN, as the effective time-step is learned during training and can be significantly different from the nominal time-step  $\Delta t$ . It is essential to point out that including this multi-scale time stepping is only possible, as each neuron (within the same hidden layer) is independent of the others and can be integrated with a different effective time step. Finally, we also share the control hyperparameter  $\alpha$  across the different layers, which results in a memory unit of  $L$  layers with a total of only 2 hyperparameters.

## 2.1 Motivation and background

The ODE system (2) is a model for a nonlinear system of uncoupled driven oscillators [18]. To see this, we denote  $\mathbf{y}_i(t)$  as the displacement and  $\mathbf{z}_i(t)$  as the velocity. Then, the dynamics of the  $i$ -th oscillator is determined by the frequency  $\alpha$  and also by the *forcing* or *driving* term in the second equation of (2), where the forcing acts through the input signal  $\mathbf{u}$  and is modulated by the weight  $\mathbf{V}$  and bias  $\mathbf{b}$ . Finally, the weight  $\mathbf{w}$  modulates the frequency  $\alpha$  and allows each neuron to oscillate with its own frequency, rather than the common frequency  $\alpha$  of the system. The structure of  $\mathbf{w}$  implies that each neuron is independent of the others. A key element of the oscillator system (2) is the absence of any damping or friction term. This allows the system to possess a Hamiltonian structure, with desirable long time behavior. Thus, we term the resulting RNN (6), based on the ODE system (2) as **Undamped Independent Controlled Oscillatory RNN** or **UnICORN**. We remark that networks of oscillators are very common in science and engineering [18, 43] with prominent examples being pendulums in mechanics, electrical circuits in engineering, business cycles in economics and functional brain circuits such as cortical columns in neurobiology.

## 2.2 Comparison with related work.

UnICORN lies firmly in the class of ODE-based or ODE-inspired RNNs, which have received considerable amount of attention in the machine learning literature in recent years. Neural ODEs, first proposed in [10], are a prominent example of using ODEs to construct neural networks. In this architecture, the continuous ODE serves as the learning model and gradients are computed from a sensitivity equation, which allows one to trade accuracy with computing time. Moreover, it is argued that these neural ODEs are invertible and hence, memory efficient. However, it is unclear if a general neural ODE, without any additional structure, can be invertible. Other RNN architectures that are based on discretized ODEs include those proposed in [14] and [8], where the authors proposed an *anti-symmetric* RNN, based on the discretization of a stable ODE resulting from a skew-symmetric hidden weight matrix, thus constraining the gradient dynamics.

Our proposed RNN (6) is inspired by two recent RNN architectures. The first one is *coRNN*, proposed recently in [41], where the underlying RNN architecture was also based on the use of a network of oscillators. As long as a constraint on the underlying weights was satisfied, coRNN was shown to mitigate the EVGP. In contrast to coRNN, our proposed RNN does not use a *damping* term. Moreover, each neuron, for any hidden layer, in UnICORN (6) is independent. This is very different from coRNN where all the neurons were coupled together. Finally, UnICORN is a multi-layer architecture whereas coRNN used a single hidden layer. These innovations allow us to admit a Hamiltonian structure for UnICORN and facilitate a fast and memory efficient implementation.

Our proposed architecture was also partly inspired by *IndRNN*, proposed in [32, 33], where the neurons in each hidden layers were independent of each other and interactions between neurons were mediated by stacking multiple RNN layers, with output of each hidden layer passed on to the next hidden layer, leading to a deep RNN. We clearly use this construction of independent neurons in each layer and stacking multiple layers in UnICORN (6). However in contrast to IndRNN, our proposed RNN is based on a discretized Hamiltonian system and we will not require any constraints on the weights to mitigate the EVGP.

Finally, we would like to point out that discrete Hamiltonian systems have already been used to design RNNs, for instance in [17] and also in [11], where a symplectic time-integrator for a Hamiltonian

system was proposed as the RNN architecture. However, these approaches are based on underlying time-independent Hamiltonians and are only relevant for mechanical systems as they cannot process time-dependent inputs, which arise in most realistic learning tasks. Moreover, as these methods enforce exact conservation of the Hamiltonian in time, they are not suitable for learning long-time dependencies, see [35] for a discussion and experiment on that issue. Although we use a Hamiltonian system as the basis of our proposed RNN (6), our underlying Hamiltonian (4) is time-dependent and the resulting RNN can readily process any time-dependent input signal.

### 2.3 On the Memory Efficiency of UnICORNN

As mentioned in the introduction, the standard BPTT training algorithm for RNNs requires one to store all the hidden states at every time step. To see this, we observe that for a standard multi-layer RNN with  $L$  layers and a mini-batch size of  $b$  (for any mini-batch stochastic gradient descent algorithm), the storage (in terms of floats) scales as  $\mathcal{O}(Nbd + LbmN)$ , with input and hidden sequences of length  $N$ . This memory requirement can be very high. Note that we have ignored the storage of trainable weights and biases for the RNN in the above calculation.

On the other hand, as argued before, our proposed RNN is a symplectic Euler discretization for a Hamiltonian system. Hence, it is invertible. In fact, one can explicitly write the *inverse* of UnICORNN (6) as,

$$\begin{aligned} \mathbf{y}_{n-1}^l &= \mathbf{y}_n^l - \Delta t \hat{\sigma}(\mathbf{c}^l) \odot \mathbf{z}_n^l, \\ \mathbf{z}_{n-1}^l &= \mathbf{z}_n^l + \Delta t \hat{\sigma}(\mathbf{c}^l) \odot [\sigma(\mathbf{w}^l \odot \mathbf{y}_{n-1}^l + \mathbf{V}^l \mathbf{y}_n^{\ell-1} + \mathbf{b}^l) + \alpha \mathbf{y}_{n-1}^l]. \end{aligned} \quad (7)$$

Thus, one can recover all the hidden states in a given hidden layer, only from the *stored* hidden state at the final time step, for that layer. Moreover, only the input signal needs to be stored as the other hidden states can be reconstructed from the formula (7). Hence, a straightforward calculation shows that the storage for UnICORNN scales as  $\mathcal{O}(Nbd + Lbm)$ . As  $L \ll N$ , we conclude that UnICORNN allows for a significant saving in terms of storage, when compared to a standard RNN.

## 3 Rigorous Analysis of UnICORNN

**Properties of the ODE (2).** In order to investigate the EVGP for the proposed RNN (6), we will first explore the dynamics of the gradients of hidden states  $\mathbf{y}, \mathbf{z}$  (solutions of the ODE (2)) with respect to the trainable parameters  $\mathbf{w}, \mathbf{V}$  and  $\mathbf{b}$ . Denote any scalar parameter as  $\theta$  and  $f_\theta = \frac{\partial f}{\partial \theta}$ , then differentiating the ODE (2) with respect to  $\theta$  results in the ODE,

$$\mathbf{y}'_\theta = \mathbf{z}_\theta, \quad \mathbf{z}'_\theta = -\sigma'(\mathbf{A}) \odot (\mathbf{w} \odot \mathbf{y}_\theta) - \alpha \mathbf{y}_\theta - \sigma'(\mathbf{A}) \odot \mathbf{C}(t), \quad (8)$$

where  $\mathbf{A} = \mathbf{w} \odot \mathbf{y} + \mathbf{V}\mathbf{u} + \mathbf{b}$  is the pre-activation and the coefficient  $\mathbf{C} \in \mathbb{R}^m$  is given by  $\mathbf{C}_i = \mathbf{y}_i$  if  $\theta = \mathbf{w}_i$ ,  $\mathbf{C}_i = \mathbf{u}_j$  if  $\theta = \mathbf{V}_{ij}$  and  $\mathbf{C}_i = 1$  if  $\theta = \mathbf{b}_i$ , with all other entries of the vector  $\mathbf{C}$  being zero.

It is easy to check that the ODE system (8) is a *Hamiltonian system* of form (3), with the following time-dependent Hamiltonian;

$$\mathbf{H}(\mathbf{y}_\theta, \mathbf{z}_\theta, t) := \frac{\alpha}{2} \|\mathbf{y}_\theta\|^2 + \frac{1}{2} \|\mathbf{z}_\theta\|^2 + \frac{1}{2} \sum_{i=1}^m \sigma'(\mathbf{A}_i) \mathbf{w}_i ((\mathbf{y}_\theta)_i)^2 + \sum_{i=1}^m \sigma'(\mathbf{A}_i) \mathbf{C}_i(t) (\mathbf{y}_\theta)_i. \quad (9)$$

Thus, by the well-known Liouville's theorem [42], we know that the phase space volume of (8) is preserved. Hence, this system cannot have any asymptotically stable fixed points. This implies that  $\{\mathbf{0}, \mathbf{0}\}$  cannot be a stable fixed point for the hidden state gradients  $(\mathbf{y}_\theta, \mathbf{z}_\theta)$ . Thus, we can expect that the hidden state gradients with respect to the system of oscillators (2) do not remain near zero.

On the other hand, as the Hamiltonian (9) for the hidden state gradient system (8) is time-dependent, we cannot directly infer that the gradients satisfy an upper bound. However, we prove (in Appendix C.2) the following bounds,

**Proposition 3.1.** Let  $\mathbf{y}_\theta(t), \mathbf{z}_\theta(t)$  be the solutions of the ODE system (8) for hidden state gradients, at any time  $t \in [0, 1]$ . Then for  $\alpha > 0$  and for all  $t \in [0, 1]$ , the hidden state gradients are bounded by,

$$\begin{aligned}\|\mathbf{y}_\theta(t)\|^2 &\leq \left( \|\mathbf{y}_\theta(0)\|^2 + \frac{1}{\alpha} \|\mathbf{z}_\theta(0)\|^2 \right) e^{\bar{\mathbf{w}}} + \frac{\bar{C}}{\alpha} e^{\bar{\mathbf{w}}}, \\ \|\mathbf{z}_\theta(t)\|^2 &\leq (\alpha \|\mathbf{y}_\theta(0)\|^2 + \|\mathbf{z}_\theta(0)\|^2) e^{\bar{\mathbf{w}}} + \bar{C} e^{\bar{\mathbf{w}}},\end{aligned}\tag{10}$$

with  $\bar{\mathbf{w}} = \max \left\{ \frac{\|\mathbf{w}\|_\infty^2}{\alpha}, 2 \right\}$ , and for some constant  $\bar{C}$  that only depends on the dimension  $m$  and on the constant  $U$  that bounds the input  $\mathbf{u}$ .

A similar upper bound for  $\alpha = 0$  is given in Appendix C.

**On the Exploding Gradient Problem for UnICORNN.** We train the RNN (6) to minimize the loss function,

$$\mathcal{E} := \frac{1}{N} \sum_{n=1}^N \mathcal{E}_n, \quad \mathcal{E}_n = \frac{1}{2} \|\mathbf{y}_n^L - \bar{\mathbf{y}}_n\|_2^2,\tag{11}$$

with  $\bar{\mathbf{y}}$  being the underlying ground truth (training data). Note that the loss function (11) only involves the output at the last hidden layer (we set the affine output layer to identity for the sake of simplicity). During training, we compute gradients of the loss function (11) with respect to the trainable weights and biases  $\Theta = [\mathbf{w}^\ell, \mathbf{V}^\ell, \mathbf{b}^\ell, \mathbf{c}^\ell]$ , for all  $1 \leq \ell \leq L$ , i.e.,

$$\frac{\partial \mathcal{E}}{\partial \theta} = \frac{1}{N} \sum_{n=1}^N \frac{\partial \mathcal{E}_n}{\partial \theta}, \quad \forall \theta \in \Theta.\tag{12}$$

Given the upper bounds (10) on the hidden state gradients for the underlying ODE (2), it is reasonable to expect that the gradient (12) is bounded. This is indeed the case and we have the following proposition,

**Proposition 3.2.** Let  $\alpha > 0, \Delta t < 1$  in the RNN (6) and let  $\mathbf{y}_n^\ell, \mathbf{z}_n^\ell$ , for  $1 \leq \ell \leq L$ , be the hidden states generated by the RNN (6). Then, the gradient of the loss function  $\mathcal{E}$  (11) with respect to any parameter  $\theta \in \Theta$  is bounded as,

$$\left| \frac{\partial \mathcal{E}}{\partial \theta} \right| \leq \frac{1 - (\Delta t)^L}{1 - \Delta t} e^{\bar{\mathbf{w}}^L} \bar{\mathbf{V}} (\bar{Y} + \mathbf{F}) (\Delta_1 + \Delta_2),\tag{13}$$

with  $\bar{Y} = \max_{1 \leq n \leq N} \|\bar{\mathbf{y}}_n\|_\infty$ , be a bound on the underlying training data and other quantities in (13) defined as,

$$\begin{aligned}\bar{\mathbf{w}}^L &= \max \{2, \|\mathbf{w}^L\|_\infty + \alpha\}, \quad \bar{\mathbf{V}} = \prod_{q=1}^L \max \{1, \|\mathbf{V}^q\|_\infty\}, \\ \mathbf{F} &= \sqrt{\frac{2m}{\alpha} (1 + 2\alpha \Delta t^3)} e^{\bar{\alpha}}, \quad \bar{\alpha} = 96^{\frac{1}{4}} \max \{1, \sqrt{\alpha}\}, \\ \Delta_1 &= 2 + \sqrt{2m (1 + 2\alpha \Delta t^3)} e^{\bar{\alpha}}, \\ \Delta_2 &= (2 + \alpha) \sqrt{\frac{2m}{\alpha} (1 + 2\alpha \Delta t^3)} e^{\bar{\alpha}}.\end{aligned}\tag{14}$$

This proposition, proved in Appendix D.2, demonstrates that as long as the weights  $\mathbf{w}^L, \mathbf{V}^q$  are bounded, there is a uniform bound on the hidden state gradients and the *exploding gradient problem* is mitigated for UnICORNN.

**On the Vanishing Gradient Problem for UnICORNN.** By applying the chain rule repeatedly to each term on the right-hand-side of (12), we obtain

$$\begin{aligned}\frac{\partial \mathcal{E}_n}{\partial \theta} &= \sum_{\ell=1}^L \sum_{k=1}^n \frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta}, \quad \frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta} := \frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^L} \frac{\partial \mathbf{X}_n^L}{\partial \mathbf{X}_k^\ell} \frac{\partial \mathbf{X}_k^\ell}{\partial \theta}, \\ \mathbf{X}_n^\ell &= [\mathbf{y}_n^{\ell,1}, \mathbf{z}_n^{\ell,1}, \dots, \mathbf{y}_n^{\ell,j}, \mathbf{z}_n^{\ell,j}, \dots, \mathbf{y}_n^{\ell,m}, \mathbf{z}_n^{\ell,m}].\end{aligned}\tag{15}$$

Here, the notation  $\frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta}$  refers to taking the partial derivative of  $\mathbf{X}_k^\ell$  with respect to the parameter  $\theta$ , while keeping the other arguments constant. The quantity  $\frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta}$  denotes the contribution from the  $k$ -recurrent step at the  $\ell$ -th hidden layer of the deep RNN (6) to the overall hidden state gradient at the step  $n$ . The vanishing gradient problem [38] arises if  $\left| \frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta} \right|$ , defined in (15),  $\rightarrow 0$  exponentially fast in  $k$ , for  $k \ll n$  (long-term dependencies). In that case, the RNN does not have long-term memory, as the contribution of the  $k$ -th hidden state at the  $\ell$ -th layer to error at time step  $t_n$  is infinitesimally small.

We have established that the hidden state gradients for the underlying continuous ODE (2) do not vanish. As we use a symplectic Euler discretization, the phase space volume for the discrete dynamical system (5) is also conserved [42, 19]. Hence, one can expect that the gradients of the multilayer RNN (6) do not vanish. However, these heuristic considerations need to be formalized. Observe that the vanishing gradient problem for RNNs focuses on the possible smallness of contributions of the gradient over a large number of recurrent steps. As this behavior of the gradient is independent of the number of layers, we focus on the vanishing gradient problem for a single hidden layer here, while presenting the multilayer results in Appendix D.4. Also, for the sake of definiteness, we set the scalar parameter  $\theta = \mathbf{w}^{1,p}$  for some  $1 \leq p \leq m$ . Similar results also hold for any other  $\theta \in \Theta$ . We have following representation formula (proved in Appendix D.3) for the hidden state gradients,

**Proposition 3.3.** *Let  $\mathbf{y}_n$  be the hidden states generated by the RNN (6). Then the gradient for long-term dependencies, i.e.  $k \ll n$ , satisfies the representation formula,*

$$\frac{\partial \mathcal{E}_{k,1}^{(n,1)}}{\partial \mathbf{w}^{1,p}} = -\Delta t \hat{\sigma}(\mathbf{c}^{1,p})^2 t_n \sigma'(\mathbf{A}_{k-1}^{1,p}) \mathbf{y}_{k-1}^{1,p} \left( \mathbf{y}_n^{1,p} - \bar{\mathbf{y}}_n^p \right) + \mathcal{O}(\Delta t^2). \quad (16)$$

It is clear from the representation formula (16) that there is no  $k$ -dependence for the gradient. In particular, as long as all the weights are of  $\mathcal{O}(1)$ , the leading-order term in (16) is  $\mathcal{O}(\Delta t)$ . Hence, the gradient can be small but is independent of the recurrent step  $k$ . Thus, we claim that the *vanishing gradient problem*, with respect to recurrent connections, is mitigated for UnICORNN (6).

## 4 Experiments

The details of the training procedure for each experiment can be found in Appendix A. Code to replicate the experiments can be found at <https://github.com/tk-rusch/unicorinn>.

**Implementation** The structure of UnICORNN (6) enables us to achieve a very fast implementation. First, the transformation of the input (i.e.  $\mathbf{V}^\ell \mathbf{y}_n^{\ell-1}$  for all  $l = 1, \dots, L$ ), which is the most computationally expensive part of UnICORNN, does not have a sequential structure and can thus be computed in parallel over time. Second, as the underlying ODEs of the UnICORNN are uncoupled for each neuron, the remaining recurrent part of UnICORNN is solved independently for each component. Hence, inspired by the implementation of Simple Recurrent Units (SRU) [31] and IndRNN, we present in Appendix B, the details of an efficient CUDA implementation, where the input transformation is computed in parallel and the dynamical system corresponding to each component of (6) is an independent CUDA thread.

We benchmark the training speed of UnICORNN with  $L = 2$  layers, against the fastest available RNN implementations, namely the cuDNN implementation [1] of LSTM (with 1 hidden layer), SRU and IndRNN (both with  $L = 2$  layers and with batch normalization). Fig. 1 shows the computational time (measured on a GeForce RTX 2080 Ti GPU) of the combined forward and backward pass for each network, averaged over 100 batches with each of size 128, for two different sequence lengths, i.e.  $N = 1000, 2000$ . We can see that while the cuDNN LSTM is relatively slow, the SRU, IndRNN and the UnICORNN perform similarly fast. Moreover, we also observe that UnICORNN is about 30 – 40 times faster per combined forward and backward pass, when compared to recently developed RNNs such as expRNN [7] and coRNN [41]. We thus conclude that the UnICORNN is among the fastest available RNN architectures.

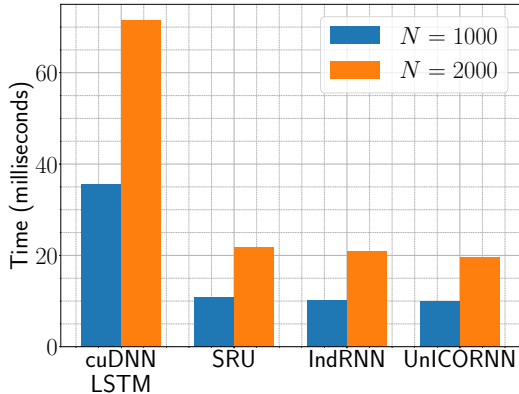


Figure 1: Measured computing time for the combined forward and backward pass for the UnICORNN as well as for three of the fastest available RNN implementations.




Figure 2: Test accuracy (mean and standard deviation) for the UnICORNN on EigenWorms for two types of sub-sampling approaches, i.e. using the last entries of the sequences as well as using a random subset of the entries. Both are shown for increasing number of entries used in each corresponding sub-sampling routine.

**Permuted sequential MNIST** A well-established benchmark for testing RNNs on input sequences with long time-dependencies is the permuted sequential MNIST (psMNIST) task [28]. Based on the classical MNIST data set [29], the flattened grey-scale matrices are randomly permuted (based on a fixed random permutation) and processed sequentially by the RNN. This makes the learning task more challenging than sequential MNIST, where one only flattens the MNIST matrices without permuting them. In order to make different methods comparable, we use the same fixed seed for the random permutation, as in [7, 6, 21]. Table 1 shows the results for UnICORNN with 3 layers, together with other recently proposed RNNs, which were explicitly designed to learn LTDs as well as two gated baselines. We see that UnICORNN clearly outperforms the other methods.

Table 1: Test accuracies on permuted sequential MNIST together with number of hidden units as well as total number of parameters for each network. All other results are taken from the corresponding original publication, cited in the main text, except that we are using the results of [9] for GRU and of [21] for LSTM.

Model	test accuracy	# units	# params
LSTM	92.9%	256	270k
GRU	94.1%	256	200k
expRNN	96.6%	512	127k
coRNN	97.3%	256	134k
IndRNN ( $L=6$ )	96.0%	128	86k
dense-IndRNN ( $L=6$ )	97.2%	128	257k
<b>UnICORNN (<math>L=3</math>)</b>	97.8%	128	35k
<b>UnICORNN (<math>L=3</math>)</b>	<b>98.4%</b>	256	135k

**Noise padded CIFAR-10** A more challenging test for the ability of RNNs to learn LTDs is provided by the recently proposed noise padded CIFAR-10 experiment [8]. In it, the CIFAR-10 data points [26] are fed to the RNN row-wise and flattened along the channels resulting in sequences of length 32. To test long term memory, entries of uniform random numbers are added such that the resulting sequences



have a length of 1000, i.e. the last 968 entries of each sequences are only noise to distract the RNNs. Table 2 shows the result of the UnICORNN with 3 layers together with the results of other recently proposed RNNs, namely for the LSTM, anti.sym. RNN and gated anti.sym. RNN [8], Lipschitz RNN [15], Incremental RNN [24], FastRNN [27] and coRNN [41]. We conclude that the proposed RNN readily outperforms all other methods on this experiment.

Table 2: Test accuracies on noise padded CIFAR-10 together with number of hidden units as well as total number of parameters for each network. All other results are taken from literature, specified in the main text.

Model	test accuracy	# units	# params
LSTM	11.6%	128	64k
Incremental RNN	54.5%	128	12k
Lipschitz RNN	55.8%	256	158k
FastRNN	45.8%	128	16k
anti.sym. RNN	48.3%	256	36k
gated anti.sym. RNN	54.7%	256	37k
coRNN	59.0%	128	46k
<b>UnICORNN (<math>L=3</math>)</b>	<b>62.4%</b>	128	47k

**EigenWorms** The EigenWorms data set [4] is a collecting of 259 very long sequences, i.e. length of 17984, describing the motion of a worm. The task is, based on the 6-dimensional motion sequences, to classify a worm as either wild-type or one of four mutant types. We use the same train/valid/test split as in [36], i.e. 70%/15%/15%. As the length of the input sequences is extremely long for this test case, we benchmark UnICORNN against three sub-sampling based baselines. These include the results of [36], which is based on signature sub-sampling routine for neural controlled differential equations. Additionally after a hyperparameter fine-tuning procedure, we perform a random sub-sampling as well as truncated back-propagation through time (BPTT) routine using LSTMs, where the random sub-sampling is based on 200 randomly selected time points of the sequences as well as the BPTT is truncated after the last 500 time points of the sequences. Furthermore, we compare UnICORNN with three leading RNN architectures for solving LTD tasks, namely expRNN, IndRNN and coRNN, which are all applied to the full-length sequences. The results, presented in Table 3, show that while sub-sampling approaches yield moderate test accuracies, expRNN as well as the IndRNN yield very poor accuracies. In contrast, coRNN performs very well. However, the best results are obtained for UnICORNN as it reaches a test accuracy of more than 90%, while at the same time yielding a relatively low standard deviation, further underlining the robustness of the proposed RNN.

Table 3: Test accuracies on EigenWorms using 5 re-trainings of each best performing network (based on the validation set) together with number of hidden units as well as total number of parameters for each network.

Model	test accuracy	# units	# params
t-BPTT LSTM	57.9%±7.0%	32	5.3k
sub-samp. LSTM	69.2%±8.3%	32	5.3k
sign.-NCDE	77.8%±5.9%	32	35k
expRNN	40.0%±10.1%	64	2.8k
IndRNN ( $L=2$ )	49.7%±4.8%	32	1.6k
coRNN	86.7%±3.0%	32	2.4k
<b>UnICORNN (<math>L=2</math>)</b>	<b>90.3%±3.0%</b>	32	1.5k

As this data set has only recently been proposed as a test for RNNs in learning LTDs, it is unclear if the input sequences truly exhibit very long time-dependencies. To investigate this further, we train

UnICORNN on a subset of the entries of the sequences. To this end, we consider using only the last entries as well as using a random subset of the entries. Fig. 2 shows the distributional results (10 re-trainings of the best performing UnICORNN) for the number of entries used in each sub-sampling routine, ranging from only using 1000 entries to using the full sequences for training. We can see that in order to reach a test accuracy of 80% when training on the last entries of the sequences, at least the last 10k entries are needed. Moreover, for both sub-sampling methods the test accuracy increases monotonically as the number of entries for training is increased. On the other hand, using a random subset of the entries increases the test accuracy significantly when compared to using only the last entries of the sequences. This indicates that the important entries of the sequences, i.e. information needed in order to classify them correctly, are uniformly distributed throughout the full sequence. We thus conclude that the EigenWorms data set indeed exhibits *very* long time-dependencies.

**Healthcare application: Vital signs prediction** We apply UnICORNN on two real-world data sets in health care, aiming to predict the vital signs of a patient, based on PPG and ECG signals. The data sets are part of the TSR archive [44] and are based on clinical data from the Beth Israel Deaconess Medical Center. The PPG and ECG signals are sampled with a frequency of 125Hz for 8 minutes each. The resulting two-dimensional sequences have a length of 4000. The goal is to predict a patient’s respiratory rate (RR) and heart rate (HR) based on these signals. We compare UnICORNN to 3 leading RNN architectures for solving LTDs, i.e. expRNN, IndRNN and coRNN. Additionally, we present two baselines using the LSTM as well as the recently proposed sub-sampling method of computing signatures for neural controlled differential equations (NCDE) [36]. Following [36], we split the 7949 sequences in a training set, validation set and testing set, using a 70%/15%/15% split. Table 4 shows the distributional results of all networks using 5 re-trainings of the best performing RNN. We observe that while the LSTM does not reach a low  $L^2$  testing error in both experiments, the other RNNs approximate the vital signs reasonably well. However, UnICORNN clearly outperforms all other methods on both benchmarks. We emphasize that UnICORNN significantly outperforms all other state-of-the-art methods on estimating the RR, which is of major importance in modern healthcare applications for monitoring hospital in-patients as well as for mobile health applications, as special invasive equipment (for instance capnometry or measurement of gas flow) is normally needed to do so [40].

Table 4:  $L^2$  test error on vital sign prediction using 5 re-trainings of each best performing network (based on the validation set), where the respiratory rate (RR) and heart rate (HR) is estimated based on PPG and ECG data.

Model	RR	HR
sign.-NCDE	1.51±0.08	2.97±0.45
LSTM	2.28±0.25	10.7±2.0
expRNN	1.57±0.16	1.87±0.19
IndRNN ( $L=3$ )	1.47±0.09	2.1±0.2
coRNN	1.45±0.23	1.71±0.1
<b>UnICORNN (<math>L=3</math>)</b>	<b>1.06±0.03</b>	<b>1.39±0.09</b>

**Sentiment analysis: IMDB** As a final experiment, we test the proposed UnICORNN on the widely used NLP benchmark data set IMDB [34], which consists of 50k online movie reviews with 25k reviews used for training and 25k reviews used for testing. This denotes a classical sentiment analysis task, where the model has to decide whether a movie review is positive or negative. We use 30% of the training set (i.e. 7.5k reviews) as the validation set and restrict the dictionary to 25k words. We choose an embedding size of 100 and initialize it with the pretrained 100d GloVe [39] vectors.

Table 5 shows the results for UnICORNN with 2 layers together with other recently proposed RNN architectures and gated baselines (which are known to perform very well on these tasks). The result of ReLU GRU is taken from [13], of coRNN from [41] and all other results are taken from [5]. We can see that UnICORNN outperforms the other methods while requiring significantly less parameters. We thus

Table 5: Test accuracies on IMDB together with number of hidden units as well as total number of parameters (without embedding) for each network. All other results are taken from literature, specified in the main text.

Model	test accuracy	# units	# params
LSTM	86.8%	128	220k
skip LSTM	86.6%	128	220k
GRU	85.2%	128	99k
ReLU GRU	84.8%	128	99k
skip GRU	86.6%	128	165k
coRNN	87.4 %	128	46k
<b>UnICORNN (<math>L=2</math>)</b>	<b>88.4%</b>	128	30k

conclude, that the UnICORNN can also be successfully applied to problems, which do not necessarily exhibit long term-dependencies.

## 5 Discussion

The design of RNNs that can accurately handle sequential inputs with long-time dependencies is very challenging. This is largely on account of the exploding and vanishing gradient problem (EVGP). Moreover, there is a significant increase in both computational time as well as memory requirements when LTD tasks have to be processed. Our main aim in this article was to present a novel RNN architecture which is fast, memory efficient, *invertible* and mitigates the EVGP. To this end, we proposed UnICORNN (6), an RNN based on the symplectic Euler discretization of a Hamiltonian system of second-order ODEs (2) modeling a network of independent, undamped, controlled and driven oscillators. In order to gain expressivity, we stack layers of RNNs and also endow this construction with a multi-scale feature by training the effective time step in (6).

Given the Hamiltonian structure of our continuous and discrete dynamical system, invertibility and volume preservation in phase space are guaranteed. Invertibility enables the proposed RNN to be memory efficient. The independence of neurons within each hidden layer allows us to build a highly efficient CUDA implementation of UnICORNN that is as fast as the fastest available RNN architectures. Motivated by the fact that the underlying ODE (1) as well as the ODE (8), governing the evolution of hidden state gradient possess a time-dependent Hamiltonian, we prove rigorous upper bounds (13) on the gradients and show that the exploding gradient problem is mitigated for UnICORNN. Moreover, we derive an explicit representation formula (16) for the gradients of (6), which shows that the vanishing gradient problem is also mitigated. Finally, we have tested UnICORNN on a suite of benchmarks that includes both synthetic as well as realistic learning tasks, designed to test the ability of an RNN to deal with long-time dependencies. In all the experiments, UnICORNN was able to show state of the art performance.

It is instructive to compare UnICORNN with two recently proposed RNN architectures, with which it shares some essential features. First, the use of coupled oscillators to design RNNs was already explored in the case of coRNN [41]. In contrast to coRNN, neurons in UnICORNN are independent (uncoupled) and as there is no damping, UnICORNN possesses a Hamiltonian structure. This paves the way for invertibility as well as for mitigating the EVGP without any assumptions on the weights whereas the mitigation of EVGP with coRNN was conditional on restrictions on weights. Finally, UnICORNN provides even better performance on benchmarks than coRNN, while being significantly faster. While we also use independent neurons in each hidden layer and stack RNN layers together as in IndRNN [32], our design principle is completely different as it is based on Hamiltonian ODEs. Consequently, we do not impose weight restrictions, which are necessary for IndRNN to mitigate the EVGP. Moreover, in contrast to IndRNNs, our architecture is invertible and hence, memory efficient.

This work can be extended in different directions. First, UnICORNN is a very flexible architecture in terms of stacking layers of RNNs together. We have used a fully connected stacking in (6) but other possibilities can be readily explored. See Appendix D.5 for a discussion on the use of residual connections

in stacking layers of UnICORNN. Second, the invertibility of UnICORNN can be leveraged in the context of normalizing flows [37], where the objective is to parametrize a flow such that the resulting Jacobian is readily computable. Finally, our focus in this article was on testing UnICORNN on learning tasks with long-time dependencies. Given that the underlying ODE (2) models oscillators, one can envisage that UnICORNN will be very competitive with respect to processing different time series data that arise in healthcare AI such as EEG and EMG data, as well as seismic time series from the geosciences.

## Acknowledgements.

The research of TKR and SM was partially supported by European Research Council Consolidator grant ERCCoG 770880: COMANFLO.

## References

- [1] J. Appleyard, T. Kocisky, and P. Blunsom. Optimizing performance of recurrent neural networks on gpus. *arXiv preprint arXiv:1604.01946*, 2016.
- [2] M. Arjovsky, A. Shah, and Y. Bengio. Unitary evolution recurrent neural networks. In *International Conference on Machine Learning*, pages 1120–1128, 2016.
- [3] V. I. Arnold. *Mathematical methods of classical mechanics*. Springer Verlag, New York, 1989.
- [4] A. Bagnall, H. A. Dau, J. Lines, M. Flynn, J. Large, A. Bostrom, P. Southam, and E. Keogh. The uea multivariate time series classification archive, 2018. *arXiv preprint arXiv:1811.00075*, 2018.
- [5] V. Campos, B. Jou, X. Giró-i-Nieto, J. Torres, and S. Chang. Skip RNN: learning to skip state updates in recurrent neural networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.
- [6] M. L. Casado. Trivializations for gradient-based optimization on manifolds. In *Advances in Neural Information Processing Systems*, pages 9154–9164, 2019.
- [7] M. L. Casado and D. Martínez-Rubio. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, volume 97 of *Proceedings of Machine Learning Research*, pages 3794–3803, 2019.
- [8] 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*, 2018.
- [9] S. Chang, Y. Zhang, W. Han, M. Yu, X. Guo, W. Tan, X. Cui, M. Witbrock, M. A. Hasegawa-Johnson, and T. S. Huang. Dilated recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 77–87, 2017.
- [10] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, pages 6571–6583, 2018.
- [11] Z. Chen, J. Zhang, M. Arjovsky, and L. Bottou. Symplectic recurrent neural networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- [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. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, 2014.
- [13] R. Dey and F. M. Salemt. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)*, pages 1597–1600. IEEE, 2017.

- [14] W. E. A proposal on machine learning via dynamical systems. *Commun. Math. Stat*, 5:1–11, 2017.
- [15] N. B. Erichson, O. Azencot, A. Queiruga, and M. W. Mahoney. Lipschitz recurrent neural networks. In *International Conference on Learning Representations*, 2021.
- [16] Y. Gal and Z. Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. *Advances in neural information processing systems*, 29:1019–1027, 2016.
- [17] S. Greydanus, M. Dzamba, and J. Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, pages 15379–15389, 2019.
- [18] J. Guckenheimer and P. Holmes. *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*. Springer Verlag, New York, 1990.
- [19] E. Hairer, C. Lubich, and G. Wanner. Geometric numerical integration illustrated by the störmer-verlet method. *Acta Numerica*, 14:399–450, 2003.
- [20] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [21] K. Helfrich, D. Willmott, and Q. Ye. Orthogonal recurrent neural networks with scaled cayley transform. In *International Conference on Machine Learning*, pages 1969–1978. PMLR, 2018.
- [22] M. Henaff, A. Szlam, and Y. LeCun. Recurrent orthogonal networks and long-memory tasks. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2034–2042, 2016.
- [23] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [24] A. Kag, Z. Zhang, and V. Saligrama. Rnns incrementally evolving on an equilibrium manifold: A panacea for vanishing and exploding gradients? In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- [25] G. Kerg, K. Goyette, M. P. 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*, pages 13591–13601, 2019.
- [26] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [27] A. Kusupati, M. Singh, K. Bhatia, A. Kumar, P. Jain, and M. Varma. Fastgrnn: A fast, accurate, stable and tiny kilobyte sized gated recurrent neural network. In *Advances in Neural Information Processing Systems*, pages 9017–9028, 2018.
- [28] Q. V. Le, N. Jaitly, and G. E. Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015.
- [29] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [30] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521:436–444, 2015.
- [31] T. Lei, Y. Zhang, S. I. Wang, H. Dai, and Y. Artzi. Simple recurrent units for highly parallelizable recurrence. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- [32] S. Li, W. Li, C. Cook, C. Zhu, and Y. Gao. Independently recurrent neural network (indrnn): Building a longer and deeper rnn. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5457–5466, 2018.

- [33] S. Li, W. Li, C. Cook, Y. Gao, and C. Zhu. Deep independently recurrent neural network (indrnn). *arXiv preprint arXiv:1910.06251*, 2019.
- [34] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pages 142–150. Association for Computational Linguistics, 2011.
- [35] M. MacKay, P. Vicol, J. Ba, and R. B. Grosse. Reversible recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 9029–9040, 2018.
- [36] J. Morrill, P. Kidger, C. Salvi, J. Foster, and T. Lyons. Neural cdes for long time series via the log-ode method. *arXiv preprint arXiv:2009.08295*, 2020.
- [37] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762v1*, 2019.
- [38] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *ICML'13*, page III–1310–III–1318. JMLR.org, 2013.
- [39] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [40] M. A. Pimentel, A. E. Johnson, P. H. Charlton, D. Birrenkott, P. J. Watkinson, L. Tarassenko, and D. A. Clifton. Toward a robust estimation of respiratory rate from pulse oximeters. *IEEE Transactions on Biomedical Engineering*, 64(8):1914–1923, 2016.
- [41] 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.
- [42] J. Sanz Serna and M. Calvo. *Numerical Hamiltonian problems*. Chapman and Hall, London, 1994.
- [43] S. Strogatz. *Nonlinear Dynamics and Chaos*. Westview, Boulder CO, 2015.
- [44] C. W. Tan, C. Bergmeir, F. Petitjean, and G. I. Webb. Monash university, uea, ucr time series regression archive. *arXiv preprint arXiv:2006.10996*, 2020.
- [45] P. J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.
- [46] S. Wisdom, T. Powers, J. Hershey, J. Le Roux, and L. Atlas. Full-capacity unitary recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 4880–4888, 2016.

**Supplementary Material for:**  
UnICORNN: A recurrent model for learning *very* long time dependencies

## A Training details

All experiments were run on GPU, namely NVIDIA GeForce GTX 1080 Ti and NVIDIA GeForce RTX 2080 Ti. The hidden weights  $\mathbf{w}$  of the UnICORNN are initialized according to  $\mathcal{U}(0,1)$ , while all biases are set to zero. The trained vector  $\mathbf{c}$  is initialized according to  $\mathcal{U}(-0.1,0.1)$ . The input weights  $\mathbf{V}$  are initialized according to the Kaiming uniform initialization [20] based on the input dimension mode and the negative slope of the rectifier set to  $a = 8$ .

The hyperparameters of the UnICORNN are selected using a random search algorithm based on a validation set. The hyperparameters of the best performing UnICORNN can be seen in Table 6. The value for  $\Delta t$  and  $\alpha$  is shared across all layers, except for the IMDB task and EigenWorms task, where we use a different  $\Delta t$  value for the first layer and the corresponding  $\Delta t$  value in Table 6 for all subsequent layers, i.e. we use  $\Delta t = 6.6 \times 10^{-3}$  for IMDB and  $\Delta t = 2.81 \times 10^{-5}$  for EigenWorms in the first layer. Additionally, the dropout column corresponds to variational dropout [16], which is applied after each consecutive layer. Note that for the IMDB task also an embedding dropout with  $p = 0.65$  is used.

We train the UnICORNN for a total of 50 epochs on the IMDB task and for a total of 250 epochs on the EigenWorms task. Moreover, we train UnICORNN for 650 epochs on psMNIST, after which we decrease the learning rate by a factor of 10 and proceed training for 3 times the amount of epochs used before reducing the learning rate. On all other tasks, UnICORNN is trained for 250 epochs, after which we decrease the learning rate by a factor of 10 and proceed training for additional 250 epochs. The resulting best performing networks are selected *based on a validation set*.

Table 6: Hyperparameters of the best performing UnICORNN architecture (based on a validation set) for each experiment.

Experiment	learning rate	dropout	batch size	$\Delta t$	$\alpha$
noise padded CIFAR-10	$3.14 \times 10^{-2}$	$1.0 \times 10^{-1}$	30	$1.26 \times 10^{-1}$	13.0
psMNIST (#units = 128)	$1.14 \times 10^{-3}$	$1.0 \times 10^{-1}$	64	$4.82 \times 10^{-1}$	12.53
psMNIST (#units = 256)	$2.51 \times 10^{-3}$	$1.0 \times 10^{-1}$	32	$1.9 \times 10^{-1}$	30.65
IMDB	$1.67 \times 10^{-4}$	$6.1 \times 10^{-1}$	32	$2.05 \times 10^{-1}$	0.0
EigenWorms	$8.59 \times 10^{-3}$	0.0	8	$3.43 \times 10^{-2}$	0.0
Healthcare: RR	$3.98 \times 10^{-3}$	$1.0 \times 10^{-1}$	32	$1.1 \times 10^{-2}$	9.0
Healthcare: HR	$2.88 \times 10^{-3}$	$1.0 \times 10^{-1}$	32	$4.6 \times 10^{-2}$	10.0

## B Implementation details

As already described in the implementation details of the main paper, we can speed up the computation of the forward and backward pass, by parallelizing the input transformation and computing the recurrent part for each independent dimension in an independent CUDA thread. While the forward/backward pass for the input transformation is simply that of an affine transformation, we discuss only the recurrent part. Since we compute the gradients of each dimension of the UnICORNN independently and add them up afterwards to get the full gradient, we will simplify to the following one-dimensional system:

$$\begin{aligned} z_n &= z_{n-1} - \Delta t \hat{\sigma}(c) [\sigma(wy_{n-1} + x_n) + \alpha y_{n-1}], \\ y_n &= y_{n-1} + \Delta t \hat{\sigma}(c) z_n, \end{aligned}$$

where  $x_n = (\mathbf{V}\mathbf{u}_n)_j$  is the transformed input corresponding to the respective dimension  $j = 1, \dots, m$ .

Since we wish to train the UnICORNN on some given objective

$$\mathcal{E} := \sum_{n=1}^N \tilde{\mathcal{E}}(y_n), \quad (17)$$

where  $\tilde{\mathcal{E}}$  is some loss function taking the hidden states  $y_n$  as inputs, for instance mean-square distance of (possibly transformed) hidden states  $y_n$  to some ground truth. During training, we compute gradients of the loss function (17) with respect to the following quantities  $\Theta = [w, \Delta t, x_n]$ , i.e.

$$\frac{\partial \mathcal{E}}{\partial \theta} = \sum_{n=1}^N \frac{\partial \tilde{\mathcal{E}}(y_n)}{\partial \theta}, \quad \forall \theta \in \Theta. \quad (18)$$

We can work out a recursion formula to compute the gradients in (18). We will exemplarily provide the formula for the gradient with respect to the hidden weight  $w$ . The computation of the gradients with respect to the other quantities follow similarly. Thus

$$\delta_k^z = \delta_{k-1}^z + \delta_{k-1}^y \Delta t \hat{\sigma}(c), \quad (19)$$

$$\delta_k^y = \delta_{k-1}^y - \delta_k^z \Delta t \hat{\sigma}(c) [\sigma'(w y_{N-k} + x_{N-k+1}) w + \alpha] + \frac{\partial \tilde{\mathcal{E}}}{\partial y_{N-k}}, \quad (20)$$

with initial values  $\delta_0^y = \frac{\partial \tilde{\mathcal{E}}}{\partial y_N}$  and  $\delta_0^z = 0$ . The gradient can then be computed as

$$\frac{\partial \mathcal{E}}{\partial w} = \sum_{k=1}^N a_k, \quad \text{with } a_k = -\delta_k^z \Delta t \hat{\sigma}(c) \sigma'(w y_{N-k} + x_{N-k+1}) y_{N-k}. \quad (21)$$

Note that this recursive formula is a direct formulation of the back-propagation through time algorithm [45] for the UnICORNN.

We can verify formula (19)-(21) by explicitly calculating the gradient in (18):

$$\begin{aligned} \frac{\partial \mathcal{E}}{\partial w} &= \sum_{n=1}^N \frac{\partial \tilde{\mathcal{E}}(y_n)}{\partial w} = \sum_{n=1}^{N-1} \frac{\partial \tilde{\mathcal{E}}(y_n)}{\partial w} + \frac{\partial \tilde{\mathcal{E}}}{\partial y_N} \left[ \frac{\partial y_{N-1}}{\partial w} + \Delta t \hat{\sigma}(c) \left( \frac{\partial z_{N-1}}{\partial w} - \Delta t \hat{\sigma}(c) (\sigma'(w y_{N-1} + x_N)) \right. \right. \\ &\quad \left. \left. (y_{N-1} + w \frac{\partial y_{N-1}}{\partial w}) + \alpha \frac{\partial y_{N-1}}{\partial w} \right) \right] = \sum_{n=1}^{N-2} \frac{\partial \tilde{\mathcal{E}}(y_n)}{\partial w} + a_1 + \delta_1^z \frac{\partial z_{N-1}}{\partial w} + \delta_1^y \frac{\partial y_{N-1}}{\partial w} \\ &= \sum_{n=1}^{N-2} \frac{\partial \tilde{\mathcal{E}}(y_n)}{\partial w} + a_1 + \delta_1^y \frac{\partial y_{N-2}}{\partial w} + (\delta_1^y \Delta t \hat{\sigma}(c) + \delta_1^z) \left( \frac{\partial z_{N-2}}{\partial w} - \Delta t \hat{\sigma}(c) (\sigma'(w y_{N-2} + x_{N-1})) \right. \\ &\quad \left. (y_{N-2} + w \frac{\partial y_{N-2}}{\partial w}) + \alpha \frac{\partial y_{N-2}}{\partial w} \right) = \sum_{n=1}^{N-3} \frac{\partial \tilde{\mathcal{E}}(y_n)}{\partial w} + \sum_{k=1}^2 a_k + \delta_2^z \frac{\partial z_{N-2}}{\partial w} + \delta_2^y \frac{\partial y_{N-2}}{\partial w}. \end{aligned}$$

Iterating the same reformulation yields the desired formula (19)-(21).

## C Bounds on the ODE (2) of the main text.

Rewriting the ODE (2) of the main text as,

$$\mathbf{z}' = -[\sigma(\mathbf{A}) + \alpha \mathbf{y}], \quad \mathbf{y}' = \mathbf{z}, \quad \mathbf{A} = \mathbf{w} \odot \mathbf{y} + \mathbf{V} \mathbf{u} + \mathbf{b}, \quad (22)$$

together with the initial data  $\mathbf{y}_0 = \mathbf{z}_0 = 0$ .



## C.1 Energy bounds.

We have the following energy bounds on the solutions of the ODE (22),

**Proposition C.1.** *Let  $\mathbf{y}(t), \mathbf{z}(t)$  be the solutions of the ODE system (22) at any time  $t \in [0, 1]$  and the initial data for (22) is given by,*

$$\mathbf{y}(0) = \mathbf{z}(0) \equiv 0.$$

Then for  $\alpha > 0$ , the solutions of (22) are bounded by,

$$\|\mathbf{y}(t)\|^2 \leq \frac{me}{\alpha}, \quad \|\mathbf{z}(t)\|^2 \leq me, \quad \forall t \in (0, 1], \quad (23)$$

and for  $\alpha = 0$ , the solutions of (22) are bounded by,

$$\|\mathbf{y}(t)\|^2 \leq me^2, \quad \|\mathbf{z}(t)\|^2 \leq me^2, \quad \forall t \in (0, T]. \quad (24)$$

Here  $\|\mathbf{x}\|^2 = \langle \mathbf{x}, \mathbf{x} \rangle$ , denotes the Euclidean norm of the vector  $\mathbf{x} \in \mathbb{R}^m$  and  $\langle \cdot, \cdot \rangle$ , the corresponding inner product.

*Proof.* To prove the bound (23), we take an inner product of the first equation in (22) with  $\alpha \mathbf{y}$  and an inner product of the second equation in (22) with  $\mathbf{z}$  and add the result to obtain,

$$\begin{aligned} \frac{d}{dt} \left( \frac{\alpha \|\mathbf{y}\|^2 + \|\mathbf{z}\|^2}{2} \right) &= -\langle \sigma(\mathbf{A}), \mathbf{z} \rangle \\ &\leq \frac{\|\sigma(\mathbf{A})\|^2}{2} + \frac{\|\mathbf{z}\|^2}{2} \quad (\text{Cauchy's inequality}) \\ &\leq \frac{m}{2} + \frac{\|\mathbf{z}\|^2}{2} + \frac{\alpha \|\mathbf{y}\|^2}{2} \quad (|\sigma| \leq 1). \end{aligned}$$

Applying the Grönwall's inequality to the above differential inequality directly leads to,

$$\alpha \|\mathbf{y}(t)\|^2 + \|\mathbf{z}(t)\|^2 \leq mte^t \leq me,$$

as  $t \leq 1$  and (23) follows as a direct consequence.

To proof (24), we set  $\alpha = 0$  in (22) and take an inner product of the first equation in (22) with  $\mathbf{y}$  and an inner product of the second equation in (22) with  $\mathbf{z}$  and add the result to obtain,

$$\begin{aligned} \frac{d}{dt} \left( \frac{\|\mathbf{y}\|^2 + \|\mathbf{z}\|^2}{2} \right) &= \langle \mathbf{y}, \mathbf{z} \rangle - \langle \sigma(\mathbf{A}), \mathbf{z} \rangle \\ &\leq \frac{\|\sigma(\mathbf{A})\|^2}{2} + \frac{\|\mathbf{y}\|^2}{2} + \|\mathbf{z}\|^2 \quad (\text{Cauchy's inequality}) \\ &\leq \frac{m}{2} + \|\mathbf{y}\|^2 + \|\mathbf{z}\|^2 \quad (|\sigma| \leq 1). \end{aligned}$$

The bound (24) follows as a direct consequence of the Grönwall's inequality.  $\square$

## C.2 Hidden state gradients and Proof of proposition 3.1 of main text.

We recall that the ODE (8) in the main text for the hidden state gradients with respect to any trainable parameter  $\theta \in \Theta$  is,

$$\mathbf{y}'_{\theta} = \mathbf{z}_{\theta}, \quad \mathbf{z}'_{\theta} = -\sigma'(\mathbf{A}) \odot (\mathbf{w} \odot \mathbf{y}_{\theta}) - \alpha \mathbf{y}_{\theta} - \sigma'(\mathbf{A}) \odot \mathbf{C}(t), \quad (25)$$

with the coefficient  $\mathbf{C} \in \mathbb{R}^m$  given by,

$$\mathbf{C} = \begin{cases} [0, 0, \dots, \mathbf{y}_i, \dots 0], & \text{if } \theta = \mathbf{w}_i, \\ [0, 0, \dots, \mathbf{u}_j, \dots 0], & \text{if } \theta = \mathbf{V}_{ij}, \\ [0, 0, \dots, 1, \dots 0], & \text{if } \theta = \mathbf{b}_i. \end{cases} \quad (26)$$

Note that only the  $i$ -th entry of  $\mathbf{C}$  is non-zero and all the remaining entries are zero. It is clear from the bounds (23) and (24) that as long as we consider bounded inputs i.e.,  $\|\mathbf{u}(t)\|^2 < c$ , for some constants  $c > 0$  and for all time  $t \in [0, 1]$ , the coefficient  $\mathbf{C}(t)$  is bounded in the following sense,

$$\|\mathbf{C}(t)\|^2 \leq \bar{C}, \quad \forall t \in [0, 1], \quad (27)$$

for some constant  $\bar{C}$  that only depends on the dimension  $m$  and on the constant  $c$  that bounds the input.

We restate Proposition 3.1 of the main text that provides an upper bound on the hidden state gradients,

**Proposition C.2.** *Let  $\mathbf{y}_\theta(t), \mathbf{z}_\theta(t)$  be the solutions of the ODE system (25) for hidden state gradients, at any time  $t \in [0, 1]$ . Then for  $\alpha > 0$  and for all  $t \in [0, 1]$ , the hidden state gradients are bounded by,*

$$\begin{aligned} \|\mathbf{y}_\theta(t)\|^2 &\leq \left( \|\mathbf{y}_\theta(0)\|^2 + \frac{1}{\alpha} \|\mathbf{z}_\theta(0)\|^2 \right) e^{\bar{\mathbf{w}}} + \frac{\bar{C}}{\alpha} e^{\bar{\mathbf{w}}}, \\ \|\mathbf{z}_\theta(t)\|^2 &\leq (\alpha \|\mathbf{y}_\theta(0)\|^2 + \|\mathbf{z}_\theta(0)\|^2) e^{\bar{\mathbf{w}}} + \bar{C} e^{\bar{\mathbf{w}}}, \end{aligned} \quad (28)$$

with  $\bar{\mathbf{w}} = \max \left\{ \frac{\|\mathbf{w}\|_\infty^2}{\alpha}, 2 \right\}$ .

For  $\alpha = 0$ , the hidden state gradients are bounded by,

$$\max (\|\mathbf{y}_\theta(t)\|^2, \|\mathbf{z}_\theta(t)\|^2) \leq (\|\mathbf{y}_\theta(0)\|^2 + \|\mathbf{z}_\theta(0)\|^2) e^{\underline{\mathbf{w}}} + \bar{C} e^{\underline{\mathbf{w}}}, \quad (29)$$

with  $\underline{\mathbf{w}} = \max \{1 + \|\mathbf{w}\|_\infty^2, 3\}$ .

*Proof.* To prove the bound (28), we take an inner product of the first equation in (25) with  $\alpha \mathbf{y}_\theta$  and an inner product of the second equation in (22) with  $\mathbf{z}_\theta$  and add the result to obtain,

$$\begin{aligned} \frac{d}{dt} \left( \frac{\alpha \|\mathbf{y}_\theta\|^2 + \|\mathbf{z}_\theta\|^2}{2} \right) &= -\langle \sigma'(\mathbf{A}) \odot (\mathbf{w} \odot \mathbf{y}_\theta), \mathbf{z}_\theta \rangle - \langle \sigma'(\mathbf{A}) \odot \mathbf{C}, \mathbf{z}_\theta \rangle \\ &\leq \frac{\|\mathbf{w}\|_\infty^2}{\alpha} \alpha \frac{\|\mathbf{y}_\theta\|^2}{2} + \|\mathbf{z}_\theta\|^2 + \frac{\|\mathbf{C}\|^2}{2} \quad (\text{Cauchy's inequality and } |\sigma'| \leq 1) \\ \Rightarrow \frac{d}{dt} (\alpha \|\mathbf{y}_\theta\|^2 + \|\mathbf{z}_\theta\|^2) &\leq \bar{\mathbf{w}} (\alpha \|\mathbf{y}_\theta\|^2 + \|\mathbf{z}_\theta\|^2) + \bar{C} \quad (\text{by definition of } \bar{\mathbf{w}} \text{ and (27)}). \end{aligned}$$

Applying the Grönwall's inequality to the above differential inequality directly leads to the bound,

$$(\alpha \|\mathbf{y}_\theta(t)\|^2 + \|\mathbf{z}_\theta(t)\|^2) \leq (\alpha \|\mathbf{y}_\theta(0)\|^2 + \|\mathbf{z}_\theta(0)\|^2) e^{\bar{\mathbf{w}}} + \bar{C} e^{\bar{\mathbf{w}}}. \quad (30)$$

This bound in turn, readily implies (28). The proof of the bound (29) is completely analogous.  $\square$

## D Rigorous bounds on UnICORNN

We rewrite UnICORNN (Eqn. (6) in the main text) in the following form: for all  $1 \leq \ell \leq L$  and for all  $1 \leq i \leq m$

$$\begin{aligned} \mathbf{y}_n^{\ell,i} &= \mathbf{y}_{n-1}^{\ell,i} + \Delta t \hat{\sigma}(\mathbf{c}^{\ell,i}) \mathbf{z}_n^{\ell,i}, \\ \mathbf{z}_n^{\ell,i} &= \mathbf{z}_{n-1}^{\ell,i} - \Delta t \hat{\sigma}(\mathbf{c}^{\ell,i}) \sigma(\mathbf{A}_{n-1}^{\ell,i}) - \alpha \Delta t \hat{\sigma}(\mathbf{c}^{\ell,i}) \mathbf{y}_{n-1}^{\ell,i}, \\ \mathbf{A}_{n-1}^{\ell,i} &= \mathbf{w}^{\ell,i} \mathbf{y}_{n-1}^{\ell,i} + \left( \mathbf{V}^\ell \mathbf{y}_{n-1}^{\ell-1} \right)^i + \mathbf{b}^{\ell,i}. \end{aligned} \quad (31)$$

Here, we have denoted the  $i$ -th component of a vector  $\mathbf{x}$  as  $\mathbf{x}^i$ .

We follow standard practice and set  $\mathbf{y}_0^\ell = \mathbf{z}_0^\ell \equiv 0$ , for all  $1 \leq \ell \leq L$ .

## D.1 Energy Bounds.

We have the following bounds on the discrete hidden states,

**Proposition D.1.** *Let  $\mathbf{y}_n, \mathbf{z}_n$  be the hidden states at the  $n$ -th time level  $t_n$  for UnICORN (31), then under the assumption that the time step  $\Delta t < 1$  and  $\alpha > 0$ , these hidden states are bounded as,*

$$\|\mathbf{y}_n^\ell\|^2 \leq \frac{2mT}{\alpha} (1 + 2\alpha\Delta t^3) e^{\bar{\alpha}T}, \quad \|\mathbf{z}_n^\ell\|^2 \leq 2mT (1 + 2\alpha\Delta t^3) e^{\bar{\alpha}T}, \quad \forall n, \forall 1 \leq \ell \leq L, \quad (32)$$

with the constant

$$\bar{\alpha} = 96^{\frac{1}{4}} \max\{1, \sqrt{\alpha}\}.$$

If  $\alpha = 0$  and  $\Delta t < 1/2$ , the hidden states satisfy the bound,

$$\|\mathbf{y}_n^\ell\|^2 \leq 2mT e^{2T}. \quad (33)$$

*Proof.* We fix  $\ell, n$  and multiply the first equation in (31) with  $\alpha \mathbf{y}_{n-1}^{\ell, i}$  and use the elementary identity

$$b(a - b) = \frac{a^2}{2} - \frac{b^2}{2} - \frac{1}{2}(a - b)^2,$$

to obtain

$$\begin{aligned} \frac{\alpha(\mathbf{y}_n^{\ell, i})^2}{2} &= \frac{\alpha(\mathbf{y}_{n-1}^{\ell, i})^2}{2} + \frac{\alpha}{2}(\mathbf{y}_n^{\ell, i} - \mathbf{y}_{n-1}^{\ell, i})^2 + \alpha\Delta t \hat{\sigma}(\mathbf{c}^{\ell, i}) \mathbf{y}_{n-1}^{\ell, i} \mathbf{z}_n^{\ell, i}, \\ &= \frac{\alpha(\mathbf{y}_{n-1}^{\ell, i})^2}{2} + \frac{\alpha\Delta t^2}{2} (\hat{\sigma}(\mathbf{c}^{\ell, i}))^2 (\mathbf{z}_n^{\ell, i})^2 + \alpha\Delta t \hat{\sigma}(\mathbf{c}^{\ell, i}) \mathbf{y}_{n-1}^{\ell, i} \mathbf{z}_n^{\ell, i}. \end{aligned} \quad (34)$$

Next, we multiply the second equation in (31) with  $\mathbf{z}_n^{\ell, i}$  and use the elementary identity

$$a(a - b) = \frac{a^2}{2} - \frac{b^2}{2} + \frac{1}{2}(a - b)^2,$$

to obtain

$$\begin{aligned} \frac{(\mathbf{z}_n^{\ell, i})^2}{2} &= \frac{(\mathbf{z}_{n-1}^{\ell, i})^2}{2} - \frac{1}{2}(\mathbf{z}_n^{\ell, i} - \mathbf{z}_{n-1}^{\ell, i})^2 - \Delta t \hat{\sigma}(\mathbf{c}^{\ell, i}) \sigma(\mathbf{A}_{n-1}^{\ell, i}) (\mathbf{z}_n^{\ell, i} - \mathbf{z}_{n-1}^{\ell, i}) \\ &\quad - \Delta t \hat{\sigma}(\mathbf{c}^{\ell, i}) \sigma(\mathbf{A}_{n-1}^{\ell, i}) \mathbf{z}_{n-1}^{\ell, i} - \alpha\Delta t \hat{\sigma}(\mathbf{c}^{\ell, i}) \mathbf{y}_{n-1}^{\ell, i} \mathbf{z}_n^{\ell, i}. \end{aligned} \quad (35)$$

Adding (34) and (35) and using Cauchy's inequality yields,

$$\begin{aligned} \frac{\alpha(\mathbf{y}_n^{\ell, i})^2}{2} + \frac{(\mathbf{z}_n^{\ell, i})^2}{2} &\leq \frac{\alpha(\mathbf{y}_{n-1}^{\ell, i})^2}{2} + \frac{(1 + \Delta t)(\mathbf{z}_{n-1}^{\ell, i})^2}{2} + \frac{\alpha\Delta t^2}{2} (\hat{\sigma}(\mathbf{c}^{\ell, i}))^2 (\mathbf{z}_n^{\ell, i})^2 \\ &\quad + (\hat{\sigma}(\mathbf{c}^{\ell, i}))^2 (\sigma(\mathbf{A}_{n-1}^{\ell, i}))^2 \Delta t + \frac{\Delta t - 1}{2} (\mathbf{z}_n^{\ell, i} - \mathbf{z}_{n-1}^{\ell, i})^2 \\ \Rightarrow \alpha(\mathbf{y}_n^{\ell, i})^2 + (\mathbf{z}_n^{\ell, i})^2 &\leq \alpha(\mathbf{y}_{n-1}^{\ell, i})^2 + (1 + \Delta t)(\mathbf{z}_{n-1}^{\ell, i})^2 + 2\Delta t + \alpha\Delta t^2 (\mathbf{z}_n^{\ell, i})^2, \end{aligned}$$

where the last inequality follows from the fact that  $|\sigma|, |\hat{\sigma}| \leq 1$  and  $\Delta t < 1$ . Using the elementary inequality,

$$(a + b + c)^2 \leq 4a^2 + 4b^2 + 2c^2,$$

and substituting for  $\mathbf{z}_n^{\ell, i}$  from the second equation of (31) in the last inequality leads to,

$$\alpha(\mathbf{y}_n^{\ell, i})^2 + (\mathbf{z}_n^{\ell, i})^2 \leq (1 + 4\alpha^2\Delta t^4)\alpha(\mathbf{y}_{n-1}^{\ell, i})^2 + (1 + \Delta t + 2\alpha\Delta t^2)(\mathbf{z}_{n-1}^{\ell, i})^2 + 2\Delta t + 4\alpha\Delta t^4.$$

Denoting  $H_n = \alpha(\mathbf{y}_n^{\ell, i})^2 + (\mathbf{z}_n^{\ell, i})^2$  and

$$G := 1 + \Delta t + 2\alpha\Delta t^2 + 4\alpha^2\Delta t^4,$$

yields the following inequality,

$$H_n \leq GH_{n-1} + 2\Delta t + 4\alpha\Delta t^4. \quad (36)$$

Using the Taylor expansion for the exponential and the definition of  $\bar{\alpha}$ , it is straightforward to check that,

$$G \leq e^{\bar{\alpha}\Delta t}.$$

Thus, we have from (36) that,

$$H_n \leq e^{\bar{\alpha}\Delta t} H_{n-1} + 2\Delta t + 4\alpha\Delta t^4. \quad (37)$$

Iterating (37)  $n$ -times and using the fact that the initial data is such that  $H_0 \equiv 0$  we obtain,

$$\begin{aligned} H_n &\leq (2\Delta t + 4\alpha\Delta t^4) \sum_{k=0}^{n-1} e^{\bar{\alpha}k\Delta t} \leq (2\Delta t + 4\alpha\Delta t^4) n e^{\bar{\alpha}n\Delta t} \\ &\leq t_n (2 + 4\alpha\Delta t^3) e^{\bar{\alpha}t_n} \\ &\leq 2T (1 + 2\alpha\Delta t^3) e^{\bar{\alpha}T}, \end{aligned} \quad (38)$$

as  $t_n \leq T$ . The bound (32) follows directly from the definition of  $H_n$ .

The proof of (33) follows analogously.  $\square$

## D.2 On the exploding gradient problem for UnICORN and Proof of proposition 3.2 of the main text.

We train the RNN (31) to minimize the loss function,

$$\mathcal{E} := \frac{1}{N} \sum_{n=1}^N \mathcal{E}_n, \quad \mathcal{E}_n = \frac{1}{2} \|\mathbf{y}_n^L - \bar{\mathbf{y}}_n\|_2^2, \quad (39)$$

with  $\bar{\mathbf{y}}$  being the underlying ground truth (training data). Note that the loss function (39) only involves the output at the last hidden layer (we set the affine output layer to identity for the sake of simplicity). During training, we compute gradients of the loss function (39) with respect to the trainable weights and biases  $\Theta = [\mathbf{w}^\ell, \mathbf{V}^\ell, \mathbf{b}^\ell, \mathbf{c}^\ell]$ , for all  $1 \leq \ell \leq L$  i.e.

$$\frac{\partial \mathcal{E}}{\partial \theta} = \frac{1}{N} \sum_{n=1}^N \frac{\partial \mathcal{E}_n}{\partial \theta}, \quad \forall \theta \in \Theta. \quad (40)$$

We have the following bound on the gradient (40),

**Proposition D.2.** *Let  $\alpha > 0, \Delta t < 1$  in the RNN (31) and let  $\mathbf{y}_n^\ell, \mathbf{z}_n^\ell$ , for  $1 \leq \ell \leq L$ , be the hidden states generated by the RNN (31). Then, the gradient of the loss function  $\mathcal{E}$  (39) with respect to any parameter  $\theta \in \Theta$  is bounded as,*

$$\left| \frac{\partial \mathcal{E}}{\partial \theta} \right| \leq \frac{1 - (\Delta t)^L}{1 - \Delta t} T e^{\bar{\mathbf{w}}^L T} \bar{\mathbf{V}} (\bar{\mathbf{Y}} + \mathbf{F}) \Delta, \quad (41)$$

with  $\bar{\mathbf{Y}} = \max_{1 \leq n \leq N} \|\bar{\mathbf{y}}_n\|_\infty$ , be a bound on the underlying training data and other quantities in (41) defined as,

$$\begin{aligned} \bar{\mathbf{w}}^L &= \max \{2, \|\mathbf{w}^L\|_\infty + \alpha\}, \\ \bar{\mathbf{V}} &= \prod_{q=1}^L \max \{1, \|\mathbf{V}^q\|_\infty\}, \\ \mathbf{F} &= \sqrt{\frac{2m\Gamma}{\alpha} (1 + 2\alpha\Delta t^3) e^{\bar{\alpha}T}}, \\ \Delta &= 2 + \sqrt{2m\Gamma (1 + 2\alpha\Delta t^3) e^{\bar{\alpha}T}} + (2 + \alpha) \sqrt{\frac{2m\Gamma}{\alpha} (1 + 2\alpha\Delta t^3) e^{\bar{\alpha}T}}. \end{aligned} \quad (42)$$

*Proof.* For any  $1 \leq n \leq N$  and  $1 \leq \ell \leq L$ , let  $\mathbf{X}_n^\ell \in \mathbb{R}^{2m}$  be the augmented hidden state vector defined by,

$$\mathbf{X}_n^\ell = [\mathbf{y}_n^{\ell,1}, \mathbf{z}_n^{\ell,1}, \dots, \mathbf{y}_n^{\ell,i}, \mathbf{z}_n^{\ell,i}, \dots, \mathbf{y}_n^{\ell,m}, \mathbf{z}_n^{\ell,m}]. \quad (43)$$

For any  $\theta \in \Theta$ , we can apply the chain rule repeatedly to obtain the following extension of the formula of [38] to a deep RNN,

$$\frac{\partial \mathcal{E}_n}{\partial \theta} = \sum_{\ell=1}^L \sum_{k=1}^n \underbrace{\frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^L} \frac{\partial \mathbf{X}_n^L}{\partial \mathbf{X}_k^\ell} \frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta}}_{\frac{\partial \varepsilon_{k,\ell}^{(n,L)}}{\partial \theta}}. \quad (44)$$

Here, the notation  $\frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta}$  refers to taking the partial derivative of  $\mathbf{X}_k^\ell$  with respect to the parameter  $\theta$ , while keeping the other arguments constant.

We remark that the quantity  $\frac{\partial \varepsilon_{k,\ell}^{(n,L)}}{\partial \theta}$  denotes the contribution from the  $k$ -recurrent step at the  $\ell$ -th hidden layer of the deep RNN (31) to the overall hidden state gradient at the step  $n$ .

It is straightforward to calculate that,

$$\frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^L} = [\mathbf{y}_n^{L,1} - \bar{\mathbf{y}}_n^1, 0, \dots, \mathbf{y}_n^{L,i} - \bar{\mathbf{y}}_n^i, 0, \dots, \mathbf{y}_n^{L,m} - \bar{\mathbf{y}}_n^m, 0]. \quad (45)$$

Repeated application of the chain and product rules yields,

$$\frac{\partial \mathbf{X}_n^L}{\partial \mathbf{X}_k^\ell} = \prod_{j=k+1}^n \frac{\partial \mathbf{X}_j^L}{\partial \mathbf{X}_{j-1}^L} \prod_{q=\ell+1}^L \frac{\partial \mathbf{X}_k^q}{\partial \mathbf{X}_k^{q-1}}. \quad (46)$$

For any  $j$ , a straightforward calculation using the form of the RNN (31) leads to the following representation formula for the matrix  $\frac{\partial \mathbf{X}_j^L}{\partial \mathbf{X}_{j-1}^L} \in \mathbb{R}^{2m} \times \mathbb{R}^{2m}$ :

$$\frac{\partial \mathbf{X}_j^L}{\partial \mathbf{X}_{j-1}^L} = \begin{bmatrix} \mathbf{B}_j^{L,1} & 0 & \dots & 0 \\ 0 & \mathbf{B}_j^{L,2} & \dots & 0 \\ \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & \mathbf{B}_j^{L,m} \end{bmatrix}, \quad (47)$$

with the block matrices  $\mathbf{B}_j^{L,i} \in \mathbb{R}^{2 \times 2}$  given by,

$$\mathbf{B}_j^{L,i} = \begin{bmatrix} 1 - (\hat{\sigma}(\mathbf{c}^{L,i}))^2 \Delta t^2 (\mathbf{w}^{L,i} \sigma'(\mathbf{A}_{j-1}^{L,i}) + \alpha) & \hat{\sigma}(\mathbf{c}^{L,i}) \Delta t \\ -\hat{\sigma}(\mathbf{c}^{L,i}) \Delta t (\mathbf{w}^{L,i} \sigma'(\mathbf{A}_{j-1}^{L,i}) + \alpha) & 1 \end{bmatrix}. \quad (48)$$

Similarly for any  $q$ , the matrix  $\frac{\partial \mathbf{X}_k^q}{\partial \mathbf{X}_k^{q-1}} \in \mathbb{R}^{2m \times 2m}$  can be readily computed as,

$$\frac{\partial \mathbf{X}_k^q}{\partial \mathbf{X}_k^{q-1}} = \begin{bmatrix} \mathbf{D}_{11}^{q,k} & 0 & \mathbf{D}_{12}^{q,k} & 0 & \dots & \dots & \mathbf{D}_{1m}^{q,k} & 0 \\ \mathbf{E}_{11}^{q,k} & 0 & \mathbf{E}_{12}^{q,k} & 0 & \dots & \dots & \mathbf{E}_{1m}^{q,k} & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \mathbf{D}_{m1}^{q,k} & 0 & \mathbf{D}_{m2}^{q,k} & 0 & \dots & \dots & \mathbf{D}_{mm}^{q,k} & 0 \\ \mathbf{E}_{m1}^{q,k} & 0 & \mathbf{E}_{m2}^{q,k} & 0 & \dots & \dots & \mathbf{E}_{mm}^{q,k} & 0 \end{bmatrix}, \quad (49)$$

with entries given by,

$$\mathbf{D}_{i,\bar{i}}^{q,k} = -\Delta t^2 (\hat{\sigma}(\mathbf{c}^{q,i}))^2 \sigma'(\mathbf{A}_{k-1}^{q,i}) \mathbf{V}_{i\bar{i}}^q, \quad \mathbf{E}_{i,\bar{i}}^{q,k} = -\Delta t \hat{\sigma}(\mathbf{c}^{q,i}) \sigma'(\mathbf{A}_{k-1}^{q,i}) \mathbf{V}_{i\bar{i}}^q. \quad (50)$$

A direct calculation with (48) leads to,

$$\begin{aligned} \|\mathbf{B}_j^{L,i}\|_\infty &\leq \max\left(1 + \Delta t + (|\mathbf{w}^{L,i}| + \alpha) \Delta t^2, 1 + (|\mathbf{w}^{L,i}| + \alpha) \Delta t\right) \\ &\leq 1 + \max\left(2, |\mathbf{w}^{L,i}| + \alpha\right) \Delta t + \left(\max\left(2, |\mathbf{w}^{L,i}| + \alpha\right)\right)^2 \frac{\Delta t^2}{2}. \end{aligned} \quad (51)$$

Using the definition of the  $L^\infty$  norm of a matrix, we use (51) to derive the following bound from (47),

$$\begin{aligned} \left\| \frac{\partial \mathbf{X}_j^L}{\partial \mathbf{X}_{j-1}^L} \right\|_\infty &\leq 1 + \max \left( 2, \|\mathbf{w}^L\|_\infty + \alpha \right) \Delta t + \left( \max \left( 2, \|\mathbf{w}^L\|_\infty + \alpha \right) \right)^2 \frac{\Delta t^2}{2} \\ &\leq e^{\bar{\mathbf{w}}^L \Delta t}, \end{aligned} \quad (52)$$

with  $\bar{\mathbf{w}}^L$  defined in (42).

As  $\Delta t < 1$ , it is easy to see that,

$$\left\| \frac{\partial \mathbf{X}_k^q}{\partial \mathbf{X}_k^{q-1}} \right\|_\infty \leq \|\mathbf{V}^q\|_\infty \Delta t. \quad (53)$$

Combining (52) and (53), we obtain from (46)

$$\begin{aligned} \left\| \frac{\partial \mathbf{X}_n^L}{\partial \mathbf{X}_k^L} \right\|_\infty &\leq \prod_{j=k+1}^n \left\| \frac{\partial \mathbf{X}_j^L}{\partial \mathbf{X}_{j-1}^L} \right\|_\infty \prod_{q=\ell+1}^L \left\| \frac{\partial \mathbf{X}_k^q}{\partial \mathbf{X}_k^{q-1}} \right\|_\infty \\ &\leq \Delta t^{L-\ell} \prod_{q=\ell+1}^L \|\mathbf{V}^q\|_\infty e^{(n-k)\bar{\mathbf{w}}^L \Delta t} \\ &\leq \bar{\mathbf{V}} \Delta t^{L-\ell} e^{\bar{\mathbf{w}}^L T}, \end{aligned} \quad (54)$$

where the last inequality follows from the fact that  $t_n = n\Delta t \leq T$  and the definition of  $\bar{\mathbf{V}}$  in (42).

Next, we observe that for any  $\theta \in \Theta$

$$\frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta} = \left[ \frac{\partial^+ \mathbf{y}_k^{\ell,1}}{\partial \theta}, \frac{\partial^+ \mathbf{z}_k^{\ell,1}}{\partial \theta}, \dots, \frac{\partial^+ \mathbf{y}_k^{\ell,i}}{\partial \theta}, \frac{\partial^+ \mathbf{z}_k^{\ell,i}}{\partial \theta}, \dots, \frac{\partial^+ \mathbf{y}_k^{\ell,m}}{\partial \theta}, \frac{\partial^+ \mathbf{z}_k^{\ell,m}}{\partial \theta} \right]^\top. \quad (55)$$

For any  $1 \leq i \leq m$ , a direct calculation with the RNN (31) yields,

$$\begin{aligned} \frac{\partial^+ \mathbf{y}_k^{\ell,i}}{\partial \theta} &= \Delta t \hat{\sigma}'(\mathbf{c}^{\ell,i}) \frac{\partial \mathbf{c}^{\ell,i}}{\partial \theta} \mathbf{z}_k^{\ell,i} + \Delta t \hat{\sigma}(\mathbf{c}^{\ell,i}) \frac{\partial^+ \mathbf{z}_k^{\ell,i}}{\partial \theta}, \\ \frac{\partial^+ \mathbf{z}_k^{\ell,i}}{\partial \theta} &= -\Delta t \hat{\sigma}'(\mathbf{c}^{\ell,i}) \frac{\partial \mathbf{c}^{\ell,i}}{\partial \theta} \sigma(\mathbf{A}_{k-1}^{\ell,i}) - \Delta t \hat{\sigma}(\mathbf{c}^{\ell,i}) \sigma'(\mathbf{A}_{k-1}^{\ell,i}) \frac{\partial \mathbf{A}_{k-1}^{\ell,i}}{\partial \theta} - \alpha \Delta t \hat{\sigma}'(\mathbf{c}^{\ell,i}) \frac{\partial \mathbf{c}^{\ell,i}}{\partial \theta} \mathbf{y}_{k-1}^{\ell,i}. \end{aligned} \quad (56)$$

Next, we have to compute explicitly  $\frac{\partial \mathbf{c}^{\ell,i}}{\partial \theta}$  and  $\frac{\partial \mathbf{A}_{k-1}^{\ell,i}}{\partial \theta}$  in order to complete the expressions (56). To this end, we need to consider explicit forms of the parameter  $\theta$  and obtain,

If  $\theta = \mathbf{w}^{q,p}$ , for some  $1 \leq q \leq L$  and  $1 \leq p \leq m$ , then,

$$\frac{\partial \mathbf{A}_{k-1}^{\ell,i}}{\partial \theta} = \begin{cases} \mathbf{y}_{k-1}^{\ell,i}, & \text{if } q = \ell, p = i, \\ 0, & \text{if otherwise.} \end{cases} \quad (57)$$

If  $\theta = \mathbf{b}^{q,p}$ , for some  $1 \leq q \leq L$  and  $1 \leq p \leq m$ , then,

$$\frac{\partial \mathbf{A}_{k-1}^{\ell,i}}{\partial \theta} = \begin{cases} 1, & \text{if } q = \ell, p = i, \\ 0, & \text{if otherwise.} \end{cases} \quad (58)$$

If  $\theta = \mathbf{V}_{p,\bar{p}}^q$ , for some  $1 \leq q \leq L$  and  $1 \leq p, \bar{p} \leq m$ , then,

$$\frac{\partial \mathbf{A}_{k-1}^{\ell,i}}{\partial \theta} = \begin{cases} \mathbf{y}_{k-1}^{\ell-1, \bar{p}}, & \text{if } q = \ell, p = i, \\ 0, & \text{if otherwise.} \end{cases} \quad (59)$$

If  $\theta = \mathbf{c}^{q,p}$  for any  $1 \leq q \leq L$  and  $1 \leq p \leq m$ , then,

$$\frac{\partial \mathbf{A}_{k-1}^{\ell,i}}{\partial \theta} \equiv 0. \quad (60)$$

Similarly, if  $\theta = \mathbf{w}^{q,p}$  or  $\theta = \mathbf{b}^{q,p}$ , for some  $1 \leq q \leq L$  and  $1 \leq p \leq m$ , or If  $\theta = \mathbf{V}_{p,\bar{p}}^q$ , for some  $1 \leq q \leq L$  and  $1 \leq p, \bar{p} \leq m$ , then

$$\frac{\partial \mathbf{c}^{\ell,i}}{\partial \theta} \equiv 0. \quad (61)$$

On the other hand, if  $\theta = \mathbf{c}^{q,p}$  for any  $1 \leq q \leq L$  and  $1 \leq p \leq m$ , then

$$\frac{\partial \mathbf{c}^{\ell,i}}{\partial \theta} = \begin{cases} 1, & \text{if } q = \ell, p = i, \\ 0, & \text{if otherwise.} \end{cases} \quad (62)$$

For any  $\theta \in \Theta$ , by substituting (57) to (62) into (56) and doing some simple algebra with norms, leads to the following inequalities,

$$\left| \frac{\partial^+ \mathbf{z}_k^{\ell,i}}{\partial \theta} \right| \leq \Delta t \left( 1 + \alpha |\mathbf{y}_{k-1}^{\ell,i}| + \max \left( |\mathbf{y}_{k-1}^{\ell,i}|, |\mathbf{y}_k^{\ell-1,\bar{p}}|, 1 \right) \right), \quad (63)$$

and,

$$\left| \frac{\partial^+ \mathbf{y}_k^{\ell,i}}{\partial \theta} \right| \leq \Delta t |\mathbf{z}_k^{\ell,i}| + \Delta t^2 \left( 1 + \alpha |\mathbf{y}_{k-1}^{\ell,i}| + \max \left( |\mathbf{y}_{k-1}^{\ell,i}|, |\mathbf{y}_k^{\ell-1,\bar{p}}|, 1 \right) \right), \quad (64)$$

for any  $1 \leq \bar{p} \leq m$ .

By the definition of  $L^\infty$  of a vector and some straightforward calculations with (64) yields,

$$\left\| \frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta} \right\|_\infty \leq \Delta t \left( 2 + \|\mathbf{z}_k^\ell\|_\infty + (1 + \alpha) \|\mathbf{y}_{k-1}^\ell\|_\infty + \|\mathbf{y}_k^{\ell-1}\|_\infty \right). \quad (65)$$

From the energy bounds (32), we can directly bound the above inequality further as,

$$\left\| \frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta} \right\|_\infty \leq \Delta t \left( 2 + \sqrt{2mT(1 + 2\alpha\Delta t^3)} e^{\bar{\alpha}T} + (2 + \alpha) \sqrt{\frac{2mT}{\alpha}} (1 + 2\alpha\Delta t^3) e^{\bar{\alpha}T} \right). \quad (66)$$

By (45) and the definition of  $\bar{Y}$  as well as the bound (32) on the hidden states, it is straightforward to obtain that,

$$\left\| \frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^L} \right\|_\infty \leq \bar{Y} + \sqrt{\frac{2mT}{\alpha}} (1 + 2\alpha\Delta t^3) e^{\bar{\alpha}T}. \quad (67)$$

From the definition in (44) and the identity (44), we have

$$\left| \frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta} \right| \leq \left\| \frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^L} \right\|_\infty \left\| \frac{\partial \mathbf{X}_n^L}{\partial \mathbf{X}_k^\ell} \right\|_\infty \left\| \frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta} \right\|_\infty. \quad (68)$$

Substituting (67), (66) and (52) into (68) yields,

$$\left| \frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta} \right| \leq \Delta t^{L-\ell+1} e^{\bar{w}^L T} \bar{\mathbf{V}} (\bar{Y} + \mathbf{F}) \mathbf{\Delta}, \quad (69)$$

with  $\mathbf{F}$  and  $\mathbf{\Delta}$  defined in (42).

Therefore, from the fact that  $\Delta t < 1, t_n = n\Delta t \leq T$  and (44), we obtain

$$\left| \frac{\partial \mathcal{E}_n}{\partial \theta} \right| \leq \frac{1 - (\Delta t)^L}{1 - \Delta t} T e^{\bar{w}^L T} \bar{\mathbf{V}} (\bar{Y} + \mathbf{F}) \mathbf{\Delta}. \quad (70)$$

By the definition of the loss function (39) and the fact that the right-hand-side of (70) is independent of  $n$  leads to the desired bound (41).  $\square$

### D.3 On the Vanishing gradient problem for UNICORNN and Proof of Proposition 3.3 of the main text.

By applying the chain rule repeatedly to the each term on the right-hand-side of (40), we obtain

$$\frac{\partial \mathcal{E}_n}{\partial \theta} = \sum_{\ell=1}^L \sum_{k=1}^n \frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta}, \quad \frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta} := \frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^L} \frac{\partial \mathbf{X}_n^L}{\partial \mathbf{X}_k^\ell} \frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta}. \quad (71)$$

Here, the notation  $\frac{\partial^+ \mathbf{X}_k^\ell}{\partial \theta}$  refers to taking the partial derivative of  $\mathbf{X}_k^\ell$  with respect to the parameter  $\theta$ , while keeping the other arguments constant. The quantity  $\frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta}$  denotes the contribution from the  $k$ -recurrent step at the  $\ell$ -th hidden layer of the deep RNN (31) to the overall hidden state gradient at the step  $n$ . The vanishing gradient problem [38] arises if  $\left| \frac{\partial \mathcal{E}_{k,\ell}^{(n,L)}}{\partial \theta} \right|$ , defined in (71),  $\rightarrow 0$  exponentially fast in  $k$ , for  $k \ll n$  (long-term dependencies). In that case, the RNN does not have long-term memory, as the contribution of the  $k$ -th hidden state at the  $\ell$ -th layer to error at time step  $t_n$  is infinitesimally small.

As argued in the main text, the vanishing gradient problem for RNNs focuses on the possible smallness of contributions of the gradient over a large number of recurrent steps. As this behavior of the gradient is independent of the number of layers, we start with a result on the vanishing gradient problem for a single hidden layer here. Also, for the sake of definiteness, we set the scalar parameter  $\theta = \mathbf{w}^{1,p}$  for some  $1 \leq p \leq m$ . Similar results also hold for any other  $\theta \in \Theta$ . Moreover, we introduce the following order-notation,

$$\begin{aligned} \beta &= O(\gamma), \text{ for } \gamma, \beta \in \mathbb{R}_+ \quad \text{if there exists constants } \bar{C}, \underline{C} \text{ such that } \underline{C}\gamma \leq \beta \leq \bar{C}\gamma. \\ \mathbf{M} &= O(\gamma), \text{ for } \mathbf{M} \in \mathbb{R}^{d_1 \times d_2}, \gamma \in \mathbb{R}_+ \quad \text{if there exists constant } \bar{C} \text{ such that } \|\mathbf{M}\| \leq \bar{C}\gamma. \end{aligned} \quad (72)$$

We restate Proposition 3.3 of the main text,

**Proposition D.3.** *Let  $\mathbf{y}_n$  be the hidden states generated by the RNN (31). Then the gradient for long-term dependencies, i.e.  $k \ll n$ , satisfies the representation formula,*

$$\frac{\partial \mathcal{E}_{k,1}^{(n,1)}}{\partial \mathbf{w}^{1,p}} = -\Delta t \hat{\sigma}(\mathbf{c}^{1,p})^2 t_n \sigma'(\mathbf{A}_{k-1}^{1,p}) \mathbf{y}_{k-1}^{1,p} \left( \mathbf{y}_n^{1,p} - \bar{\mathbf{y}}_n^p \right) + O(\Delta t^2). \quad (73)$$

*Proof.* Following the definition (71) and as  $L = 1$  and  $\theta = \mathbf{w}^{1,p}$ , we have,

$$\frac{\partial \mathcal{E}_{k,1}^{(n,1)}}{\partial \mathbf{w}^{1,p}} := \frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^1} \frac{\partial \mathbf{X}_n^1}{\partial \mathbf{X}_k^1} \frac{\partial^+ \mathbf{X}_k^1}{\partial \mathbf{w}^{1,p}}. \quad (74)$$

We will explicitly compute all three expressions on the right-hand-side of (74). To start with, using (55), (56) and (57), we obtain,

$$\begin{aligned} \frac{\partial^+ \mathbf{X}_k^1}{\partial \mathbf{w}^{1,p}} &= \left[ 0, 0, \dots, \dots, \frac{\partial^+ \mathbf{y}_k^{1,p}}{\partial \mathbf{w}^{1,p}}, \frac{\partial^+ \mathbf{z}_k^{1,p}}{\partial \mathbf{w}^{1,p}}, \dots, \dots, 0, 0 \right]^\top, \\ \left( \frac{\partial^+ \mathbf{X}_k^1}{\partial \mathbf{w}^{1,p}} \right)_{2p-1} &= \frac{\partial^+ \mathbf{y}_k^{1,p}}{\partial \mathbf{w}^{1,p}} = -\Delta t^2 (\hat{\sigma}(\mathbf{c}^{1,p}))^2 \sigma'(\mathbf{A}_{k-1}^{1,p}) \mathbf{y}_{k-1}^{1,p}, \\ \left( \frac{\partial^+ \mathbf{X}_k^1}{\partial \mathbf{w}^{1,p}} \right)_{2p} &= \frac{\partial^+ \mathbf{z}_k^{1,p}}{\partial \mathbf{w}^{1,p}} = -\Delta t \hat{\sigma}(\mathbf{c}^{1,p}) \sigma'(\mathbf{A}_{k-1}^{1,p}) \mathbf{y}_{k-1}^{1,p}. \end{aligned} \quad (75)$$

Using the product rule (46) we have,

$$\frac{\partial \mathbf{X}_n^1}{\partial \mathbf{X}_k^1} = \prod_{j=k+1}^n \frac{\partial \mathbf{X}_j^1}{\partial \mathbf{X}_{j-1}^1}. \quad (76)$$



Observing from the expressions (47) and (48) and using the *order*-notation (72), we obtain that,

$$\frac{\partial \mathbf{X}_j^1}{\partial \mathbf{X}_{j-1}^1} = \mathbf{I}_{2m \times 2m} + \Delta t \mathbf{C}_j^1 + \mathcal{O}(\Delta t^2), \quad (77)$$

with  $\mathbf{I}_{k \times k}$  is the  $k \times k$  Identity matrix and the matrix  $\mathbf{C}_j^1$  defined by,

$$\frac{\partial \mathbf{X}_j^L}{\partial \mathbf{X}_{j-1}^L} = \begin{bmatrix} \mathbf{C}_j^{1,1} & 0 & \dots & 0 \\ 0 & \mathbf{C}_j^{1,2} & \dots & 0 \\ \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & \mathbf{C}_j^{1,m} \end{bmatrix}, \quad (78)$$

with the block matrices  $\mathbf{C}_j^{1,i} \in \mathbb{R}^{2 \times 2}$  given by,

$$\mathbf{C}_j^{1,i} = \begin{bmatrix} 0 & \hat{\sigma}(\mathbf{c}^{1,i}) \\ -\hat{\sigma}(\mathbf{c}^{1,i}) \left( \mathbf{w}^{1,i} \sigma'(\mathbf{A}_{j-1}^{1,i}) + \alpha \right) & 0 \end{bmatrix}. \quad (79)$$

By a straightforward calculation and the use of induction, we claim that

$$\prod_{j=k+1}^n \frac{\partial \mathbf{X}_j^1}{\partial \mathbf{X}_{j-1}^1} = \mathbf{I}_{2m \times 2m} + \Delta t \mathbf{C}^1 + \mathcal{O}(\Delta t^2), \quad (80)$$

with

$$\mathbf{C}^1 = \begin{bmatrix} \mathbf{C}^{1,1} & 0 & \dots & 0 \\ 0 & \mathbf{C}^{1,2} & \dots & 0 \\ \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & \mathbf{C}^{1,m} \end{bmatrix}, \quad (81)$$

with the block matrices  $\mathbf{C}^{1,i} \in \mathbb{R}^{2 \times 2}$  given by,

$$\mathbf{C}^{1,i} = \begin{bmatrix} 0 & (n-k)\hat{\sigma}(\mathbf{c}^{1,i}) \\ -(n-k)\alpha\hat{\sigma}(\mathbf{c}^{1,i}) - \hat{\sigma}(\mathbf{c}^{1,i})\mathbf{w}^{1,i} \sum_{j=k+1}^n \sigma'(\mathbf{A}_{j-1}^{1,i}) & 0 \end{bmatrix}. \quad (82)$$

By the assumption that  $k \ll n$  and using the fact that  $t_n = n\Delta t$ , we have that,

$$\Delta t \mathbf{C}^{1,i} = \begin{bmatrix} 0 & t_n \hat{\sigma}(\mathbf{c}^{1,i}) \\ -t_n \alpha \hat{\sigma}(\mathbf{c}^{1,i}) - \hat{\sigma}(\mathbf{c}^{1,i}) \mathbf{w}^{1,i} \Delta t \sum_{j=k+1}^n \sigma'(\mathbf{A}_{j-1}^{1,i}) & 0 \end{bmatrix}. \quad (83)$$

Hence, the non-zero entries in the block matrices can be  $\mathcal{O}(1)$ . Therefore, the product formula (80) is modified to,

$$\prod_{j=k+1}^n \frac{\partial \mathbf{X}_j^1}{\partial \mathbf{X}_{j-1}^1} = \mathbf{C} + \mathcal{O}(\Delta t), \quad (84)$$

with the  $2m \times 2m$  matrix  $\mathbf{C}$  defined as,

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}^1 & 0 & \dots & 0 \\ 0 & \mathbf{C}^2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & \mathbf{C}^m \end{bmatrix}, \quad (85)$$

and,

$$\mathbf{C}^i = \begin{bmatrix} 1 & t_n \hat{\sigma}(\mathbf{c}^{1,i}) \\ -t_n \alpha \hat{\sigma}(\mathbf{c}^{1,i}) - \hat{\sigma}(\mathbf{c}^{1,i}) \mathbf{w}^{1,i} \Delta t & \sum_{j=k+1}^n \sigma'(\mathbf{A}_{j-1}^{1,i}) & 1 \end{bmatrix}. \quad (86)$$

Thus by taking the product of (84) with (75), we obtain that,

$$\prod_{j=k+1}^n \frac{\partial \mathbf{X}_j^1}{\partial \mathbf{X}_{j-1}^1} \frac{\partial^+ \mathbf{X}_k^1}{\partial \mathbf{w}^{1,p}} = \left[ 0, 0, \dots, \frac{\partial^+ \mathbf{y}_k^{1,p}}{\partial \mathbf{w}^{1,p}} + \mathbf{C}_{12}^p \frac{\partial^+ \mathbf{z}_k^{1,p}}{\partial \mathbf{w}^{1,p}}, \mathbf{C}_{21}^p \frac{\partial^+ \mathbf{y}_k^{1,p}}{\partial \mathbf{w}^{1,p}} + \frac{\partial^+ \mathbf{z}_k^{1,p}}{\partial \mathbf{w}^{1,p}} \dots, 0, 0 \right]^\top + \mathcal{O}(\Delta t^2), \quad (87)$$

with  $\mathbf{C}_{12}^p, \mathbf{C}_{21}^p$  are the off-diagonal entries of the corresponding block matrix, defined in (86). Note that the  $\mathcal{O}(\Delta t^2)$  remainder term arises from the  $\Delta t$ -dependence in (75).

From (45), we have that,

$$\frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^1} = [\mathbf{y}_n^{1,1} - \bar{\mathbf{y}}_n^1, 0, \dots, \mathbf{y}_n^{1,i} - \bar{\mathbf{y}}_n^i, 0, \dots, \mathbf{y}_n^{1,m} - \bar{\mathbf{y}}_n^m, 0]. \quad (88)$$

Therefore, taking the products of (88) and (87) and substituting the explicit expressions in (75), we obtain the desired identity (73).  $\square$

#### D.4 On the vanishing gradient problem for the multilayer version of UnICORNN.

The explicit representation formula (73) holds for 1 hidden layer in (31). What happens when additional hidden layers are stacked together as in UnICORNN (31)? To answer this question, we consider the concrete case of  $L = 3$  layers as this is the largest number of layers that we have used in the context of UnICORNN with fully connected stacked layers. As before, we set the scalar parameter  $\theta = \mathbf{w}^{1,p}$  for some  $1 \leq p \leq m$ . Similar results also hold for any other  $\theta \in \Theta$ . We have the following representation formula for the gradient in this case,

**Proposition D.4.** *Let  $\mathbf{y}_n$  be the hidden states generated by the RNN (31). the gradient for long-term dependencies satisfies the representation formula,*

$$\frac{\partial \mathcal{E}_{k,1}^{(n,3)}}{\partial \mathbf{w}^{1,p}} = \Delta t^4 \hat{\sigma}(\mathbf{c}^{1,p}) t_n \frac{\partial^+ \mathbf{z}_k^{1,p}}{\partial \mathbf{w}^{1,p}} \sum_{i=1}^m \bar{\mathbf{G}}_{2i-1,2p-1} (\mathbf{y}^{3,i} - \bar{\mathbf{y}}^i) + \mathcal{O}(\Delta t^6), \quad (89)$$

with the coefficients given by,

$$\begin{aligned} \frac{\partial^+ \mathbf{z}_k^{1,p}}{\partial \mathbf{w}^{1,p}} &= -\Delta t \hat{\sigma}(\mathbf{c}^{1,p}) \sigma'(\mathbf{A}_{k-1}^{1,p}) \mathbf{y}_{k-1}^{1,p}, \\ \bar{\mathbf{G}}_{2i-1,2p-1} &= \sum_{j=1}^m \mathbf{G}_{ij}^3 \mathbf{G}_{jp}^2, \quad \forall 1 \leq i \leq m, \quad \mathbf{G}_{r,s}^q = -(\hat{\sigma}(\mathbf{c}^{q,r}))^2 \sigma'(\mathbf{A}_{n-1}^{q,r}) \mathbf{V}_{rs}^q, \quad q = 2, 3. \end{aligned} \quad (90)$$

*Proof.* Following the definition (71) and as  $L = 3$  and  $\theta = \mathbf{w}^{1,p}$ , we have,

$$\frac{\partial \mathcal{E}_{k,1}^{(n,3)}}{\partial \mathbf{w}^{1,p}} := \frac{\partial \mathcal{E}_n}{\partial \mathbf{X}_n^3} \frac{\partial \mathbf{X}_n^3}{\partial \mathbf{X}_k^1} \frac{\partial^+ \mathbf{X}_k^1}{\partial \mathbf{w}^{1,p}}. \quad (91)$$

We will explicitly compute all three expressions on the right-hand-side of (91).

In (75), we have already explicitly computed the right most expression in the RHS of (91). Using the product rule (46) we have,

$$\frac{\partial \mathbf{X}_n^3}{\partial \mathbf{X}_k^1} = \frac{\partial \mathbf{X}_n^3}{\partial \mathbf{X}_n^2} \frac{\partial \mathbf{X}_n^2}{\partial \mathbf{X}_n^1} \prod_{j=k+1}^n \frac{\partial \mathbf{X}_j^1}{\partial \mathbf{X}_{j-1}^1}. \quad (92)$$

Note that we have already obtained an explicit representation formula for  $\prod_{j=k+1}^n \frac{\partial \mathbf{X}_j^1}{\partial \mathbf{X}_{j-1}^1}$  in (84).

Next we consider the matrices  $\frac{\partial \mathbf{X}_n^3}{\partial \mathbf{X}_n^2}$  and  $\frac{\partial \mathbf{X}_n^2}{\partial \mathbf{X}_n^1}$ . By the representation formula (49), we have the following decomposition for any  $1 \leq q \leq n$ ,

$$\frac{\partial \mathbf{X}_n^q}{\partial \mathbf{X}_n^{q-1}} = \Delta t^2 \mathbf{G}^{q,n} + \Delta t H^{q,n}, \quad (93)$$

with,

$$\mathbf{G}^{q,n} = \begin{bmatrix} \mathbf{G}_{11}^{q,n} & 0 & \mathbf{G}_{12}^{q,n} & 0 & \cdots & \cdots & \mathbf{G}_{1m}^{q,n} & 0 \\ 0 & 0 & 0 & 0 & \cdots & \cdots & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ \mathbf{G}_{m1}^{q,n} & 0 & \mathbf{G}_{m2}^{q,n} & 0 & \cdots & \cdots & \mathbf{G}_{mm}^{q,n} & 0 \\ 0 & 0 & 0 & 0 & \cdots & \cdots & 0 & 0 \end{bmatrix}, \quad \mathbf{G}_{i,\bar{i}}^{q,k} = -(\hat{\sigma}(\mathbf{c}^{q,i}))^2 \sigma'(\mathbf{A}_{n-1}^{q,i}) \mathbf{V}_{i\bar{i}}^q, \quad (94)$$

and

$$\mathbf{H}^{q,n} = \begin{bmatrix} 0 & 0 & 0 & 0 & \cdots & \cdots & 0 & 0 \\ \mathbf{H}_{11}^{q,n} & 0 & \mathbf{H}_{12}^{q,n} & 0 & \cdots & \cdots & \mathbf{H}_{1m}^{q,n} & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & 0 & \cdots & \cdots & 0 & 0 \\ \mathbf{H}_{m1}^{q,n} & 0 & \mathbf{H}_{m2}^{q,n} & 0 & \cdots & \cdots & \mathbf{H}_{mm}^{q,n} & 0 \end{bmatrix}, \quad \mathbf{H}_{i,\bar{i}}^{q,k} = -\hat{\sigma}(\mathbf{c}^{q,i}) \sigma'(\mathbf{A}_{n-1}^{q,i}) \mathbf{V}_{i\bar{i}}^q. \quad (95)$$

It is straightforward to see from (95) and (94) that,

$$\mathbf{H}^{3,n} \mathbf{H}^{2,n} \equiv \mathbf{0}_{2m \times 2m}, \quad \mathbf{G}^{3,n} \mathbf{H}^{2,n} \equiv \mathbf{0}_{2m \times 2m}, \quad (96)$$

and the entries of the  $2m \times 2m$  matrix  $\bar{\mathbf{G}} = \mathbf{G}^{3,n} \mathbf{G}^{2,n}$  are given by,

$$\bar{\mathbf{G}}_{2r-1,2s-1} = \sum_{j=1}^m \mathbf{G}_{r,j}^{3,n} \mathbf{G}_{j,s}^{2,n}, \quad \bar{\mathbf{G}}_{2r-1,2s} = \bar{\mathbf{G}}_{2r,2s-1} = \bar{\mathbf{G}}_{2r,2s} = 0, \quad \forall 1 \leq r, s \leq m, \quad (97)$$

while the entries of the  $2m \times 2m$  matrix  $\bar{\mathbf{H}} = \mathbf{H}^{3,n} \mathbf{G}^{2,n}$  are given by

$$\bar{\mathbf{H}}_{2r,2s-1} = \sum_{j=1}^m \mathbf{H}_{r,j}^{3,n} \mathbf{G}_{j,s}^{2,n}, \quad \bar{\mathbf{H}}_{2r-1,2s-1} = \bar{\mathbf{H}}_{2r-1,2s} = \bar{\mathbf{H}}_{2r,2s} = 0, \quad \forall 1 \leq r, s \leq m. \quad (98)$$

Hence we have,

$$\frac{\partial \mathbf{X}_n^3}{\partial \mathbf{X}_n^2} \frac{\partial \mathbf{X}_n^2}{\partial \mathbf{X}_n^1} = \Delta t^4 (\bar{\mathbf{G}} + \Delta t^{-1} \bar{\mathbf{H}}). \quad (99)$$

Taking the matrix-vector product of (99) with (87), we obtain

$$\begin{aligned} \frac{\partial \mathbf{X}_n^3}{\partial \mathbf{X}_k^1} \frac{\partial^+ \mathbf{X}_k^1}{\partial \mathbf{w}^{1,p}} &= \Delta t^4 \left( \frac{\partial^+ \mathbf{y}_k^{1,p}}{\partial \mathbf{w}^{1,p}} + \mathbf{C}_{12}^p \frac{\partial^+ \mathbf{z}_k^{1,p}}{\partial \mathbf{w}^{1,p}} \right) [\bar{\mathbf{G}}_{1,2p-1}, \Delta t^{-1} \bar{\mathbf{H}}_{2,2p-1}, \dots, \bar{\mathbf{G}}_{2m-1,2p-1}, \Delta t^{-1} \bar{\mathbf{H}}_{2m,2p-1}]^\top + \mathcal{O}(\Delta t^6) \\ &= \Delta t^4 \mathbf{C}_{12}^p \frac{\partial^+ \mathbf{z}_k^{1,p}}{\partial \mathbf{w}^{1,p}} [\bar{\mathbf{G}}_{1,2p-1}, \Delta t^{-1} \bar{\mathbf{H}}_{2,2p-1}, \dots, \bar{\mathbf{G}}_{2m-1,2p-1}, \Delta t^{-1} \bar{\mathbf{H}}_{2m,2p-1}]^\top + \mathcal{O}(\Delta t^6), \end{aligned} \quad (100)$$

where the last identify follows from the fact that  $\frac{\partial^+ \mathbf{y}_k^{1,p}}{\partial \mathbf{w}^{1,p}} = \mathcal{O}(\Delta t^2)$ .

Therefore, taking the products of (88) and (100), we obtain the desired identity (89).  $\square$

An inspection of the representation formula (89) shows that as long as the weights are  $\mathcal{O}(1)$  and from the energy bounds (28), we know that  $\mathbf{y} \sim \mathcal{O}(1)$ , the gradient

$$\frac{\partial \mathcal{E}_{k,1}^{(n,3)}}{\partial \mathbf{w}^{1,p}} \sim \mathcal{O}(\Delta t^5),$$

where the additional  $\Delta t$  stems from the  $\Delta t$ -term in (75). Thus the gradient does not depend on the recurrent step  $k$ . Hence, there is no vanishing gradient problem with respect to the number of recurrent connections, even in the multi-layer case.

However, it is clear from the representation formulas (73) and (89), as well as the proof of proposition D.4 that for  $L$ -hidden layers in UnICORNN (31), we have,

$$\frac{\partial \mathcal{E}_{k,1}^{(n,L)}}{\partial \mathbf{w}^{1,p}} \sim \mathcal{O}\left(\Delta t^{2L-1}\right). \quad (101)$$

Thus, the gradient can become very small if too many layers are stacked together. This is not at all surprising as such a behavior occurs even if there are no recurrent connections in UnICORNN (31). In that case, we simply have a fully connected deep neural network and it is well-known that the gradient can vanish as the number of layers increases, making it harder to train deep networks.

## D.5 Residual stacking of layers in UnICORNN.

Given the above considerations, it makes imminent sense to modify the fully-connected stacking of layers in UnICORNN (31) if a moderately large number of layers ( $L \geq 4$ ) are used. It is natural to modify the fully-connected stacking with a residual stacking, see [33]. We use the following form of residual stacking,

$$\mathbf{y}_n^\ell = \mathbf{y}_{n-1}^\ell + \Delta t \hat{\sigma}(\mathbf{c}^\ell) \odot \mathbf{z}_n^\ell, \quad (102)$$

$$\mathbf{z}_n^\ell = \mathbf{z}_{n-1}^\ell - \Delta t \hat{\sigma}(\mathbf{c}^\ell) \odot [\sigma(\mathbf{w}^\ell \odot \mathbf{y}_{n-1}^\ell + \mathbf{x}_n^\ell + \mathbf{b}^l) + \alpha \mathbf{y}_{n-1}^\ell], \quad (103)$$

where the input  $\mathbf{x}_n^\ell$  corresponds to a residual connection skipping  $S$  layers, i.e.

$$\mathbf{x}_n^\ell = \begin{cases} \mathbf{\Lambda}^\ell \mathbf{y}_n^{\ell-S-1} + \mathbf{V}^\ell \mathbf{y}_n^{\ell-1}, & \text{for } l > S \\ \mathbf{V}^\ell \mathbf{y}_n^{\ell-1}, & \text{for } l \leq S \end{cases}.$$

The number of skipped layers is  $2 \leq S$  and  $\mathbf{\Lambda}^\ell \in \mathbb{R}^{m \times m}$  is a trainable matrix.

The main advantages of using a residual staking such as (102) is to alleviate the vanishing gradient problem that arises from stacking multiple layers together and obtain a better scaling of the gradient than (101). To see this, we can readily follow the proof of proposition D.4, in particular the product,

$$\frac{\partial \mathbf{X}_n^L}{\partial \mathbf{X}_n^1} = \prod_{s=1}^v \frac{\partial \mathbf{X}_n^{L-(s-1)S}}{\partial \mathbf{X}_k^{L-sS}} \prod_{\ell=2}^{L-\nu S} \frac{\partial \mathbf{X}_n^\ell}{\partial \mathbf{X}_n^{\ell-1}} + \prod_{\ell=1}^{L-1} \frac{\partial \mathbf{X}_n^{\ell+1}}{\partial \mathbf{X}_k^\ell}, \quad (104)$$

with,

$$\nu = \begin{cases} \left\lfloor \frac{L}{S} \right\rfloor, & \text{if } L \bmod S \neq 0, \\ \left\lfloor \frac{L}{S} \right\rfloor - 1, & \text{if } L \bmod S = 0. \end{cases} \quad (105)$$

Here  $\lfloor x \rfloor \in \mathbb{N}$  is the largest natural number less than or equal to  $x \in \mathbb{R}$ .

Given the additive structure in the product of gradients and using induction over matrix products as in (96) and (97), we can compute that,

$$\frac{\partial \mathbf{X}_n^L}{\partial \mathbf{X}_n^1} = \mathcal{O}\left(\Delta t^{2(\nu+L-\nu S-1)}\right) + \mathcal{O}\left(\Delta t^{2(L-1)}\right). \quad (106)$$

By choosing  $S$  large enough, we clearly obtain that  $\nu + L - \nu S - 1 < L - 1$ . Hence by repeating the arguments of the proof of proposition D.4, we obtain that to leading order, the gradient of the residual stacked version of UnICORNN scales like,

$$\frac{\partial \mathcal{E}_{k,1}^{(n,L)}}{\partial \mathbf{w}^{1,p}} \sim \mathcal{O}\left(\Delta t^{2\nu+2L-2\nu S-1}\right). \quad (107)$$

Note that (107) is far more favorable scaling for the gradient than the scaling (101) for a fully connected stacking. As a concrete example, let us consider  $L = 7$  i.e., a network of 7 stacked layers of UniCORNN.

From (101), we see that the gradient scales like  $\mathcal{O}(\Delta t^{13})$  in this case. Even for a very moderate values of  $\Delta t < 1$ , this gradient will be very small and will ensure that the first layer will have very little, if any, influence on the loss function gradients. On the other hand, for the same number of layers  $L = 7$ , let us consider the residual stacking (102) with  $S = 3$  skipped connections. In this case  $\nu = 2$  and one directly concludes from (107) that the gradient scales like  $\mathcal{O}(\Delta t^5)$ , which is significantly larger than the gradient for the fully connected version of UnICORNN. In fact, it is exactly the same as the gradient scaling for fully connected UnICORNN (31) with 3 hidden layers (89). Thus, introducing skipped connections enabled the gradient to behave like a shallower fully-connected network, while possibly showing the expressivity of a deeper network.

## E Further experimental results

As we compare the results of the UnICORNN to the results of other recent RNN architecture, where only the best results of each RNN were published for the psMNIST, noise padded CIFAR-10 and IMDB task, we as well show the best (based on a validation set) obtained results for the UnICORNN in the main paper. However, distributional results, i.e. statistics of several re-trainings of the best performing UnICORNN based on different random initialization of the trainable parameters, provide additional insights into the performance. Table 7 shows the mean and standard deviation of 10 re-trainings of the best performing UnICORNN for the psMNIST, noise padded CIFAR-10 and IMDB task. We can see that in all experiments the standard deviation of the re-trainings are relatively low, which underlines the robustness of our presented results.

Table 7: Distributional information (mean and standard deviation) on the results for the classification experiment presented in the paper, where only the best results is shown, based on 10 re-trainings of the best performing UnICORNN using different random seeds.

Experiment	Mean	Standard deviation
psMNIST (128 units)	97.7%	0.09%
psMNIST (256 units)	98.2%	0.22%
Noise padded CIFAR-10	61.5%	0.52%
IMDB	88.1%	0.19%

As emphasized in the main paper and in the last section, naively stacking of many layers for the UnICORNN might result in a vanishing gradient for the deep multi-layer model, due to the vanishing gradient problem of stacking many (not necessarily recurrent) layers. Following section D.5, one can use skipped residual connections and we see that the estimate on the gradients scale preferably when using residual connections compared to a naively stacking, when using many layers. To test this also numerically, we train a standard UnICORNN (31) as well as a residual UnICORNN (res-UnICORNN) (102), with  $S = 2$  skipping layers, on the noise padded CIFAR-10 task. Fig. 3 shows the test accuracy (mean and standard deviation) of the best resulting model for different number of network layers  $L = 3, \dots, 6$ , for the standard UnICORNN and res-UnICORNN. We can see that while both models seem to perform comparably for using only few layers, i.e.  $L = 3, 4$ , the res-UnICORNN with  $S = 2$  skipping connections outperforms the standard UnICORNN when using more layers, i.e.  $L = 5, 6$ . In particular, we can see that the standard UnICORNN is not able to significantly improve the test accuracy when using more layers, while the res-UnICORNN seems to obtain higher test accuracies when using more layers.

Moreover, Fig. 3 also shows the test accuracy for a UnICORNN with an untrained time-step vector  $\mathbf{c}$ , resulting in a UnICORNN without the multi-scale property generated by the time-step. We can see that the UnICORNN without the multi-scale feature is inferior in performance, to the standard UnICORNN as well as its residual counterpart.

Finally, we recall that the estimate (41) on the gradients for UnICORNN (31) needs the weights to be bounded, see (42). One always initializes the training with bounded weights. However, it might happen that the weights explode during training. To check this issue, in Fig. 4, we plot the mean and standard deviation of the norms of the hidden weights  $\mathbf{w}^l$  for  $l = 1, 2, 3$  during training based on 10 re-trainings




Figure 3: Test accuracies (mean and standard deviation of 10 re-trainings of the best performing model) of the standard UnICORNN, res-UnICORNN and UnICORNN without multi-scale behavior on the noise padded CIFAR-10 experiment for different number of layers  $L$ .




Figure 4: Norms (mean and standard deviation of 10 re-trainings) of the hidden weights  $\|\mathbf{w}^l\|_\infty$ , for  $l = 1, 2, 3$ , of the UnICORNN during training.

of the best performing UnICORNN on the noise padded CIFAR-10 experiment. We can see that none of the norms of the weights explode during training. In fact the weight norms seem to saturate, mostly on account of reducing the learning rate after 250 epochs. Thus, the upper bound (41) can be explicitly computed and it is finite, even after training has concluded.