# Supplementary Material for "Noisy Recurrent Neural Networks" 

## A Notation and Background

We begin by introducing some notations that will be used in this SM.

- $\|\cdot\|_{F}$ denotes Frobenius norm, $\|\cdot\|_{p}$ denote $p$-norm $(p>0)$ of a vector/matrix (in particular, $\|v\|:=\|v\|_{2}$ denotes Euclidean norm of the vector $v$ and $\|A\|_{2}$ denotes the spectral norm of the matrix $A$ ).
- The $i$ th element of a vector $v$ is denoted as $v^{i}$ or $[v]^{i}$ and the $(i, j)$-entry of a matrix $A$ is denoted as $A^{i j}$ or $[A]^{i j}$.
- $I$ denotes identity matrix (the dimension should be clear from the context).
- $\operatorname{tr}$ denotes trace, the superscript $T$ denotes transposition, and $\mathbb{R}^{+}:=(0, \infty)$.
- For a function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ such that each of its first-order partial derivatives (with respect to $x$ ) exist on $\mathbb{R}^{n}, \frac{\partial f}{\partial x} \in \mathbb{R}^{m \times n}$ denotes the Jacobian matrix of $f$.
- For a scalar-valued function $g: \mathbb{R}^{n} \rightarrow \mathbb{R}, \nabla_{h} g$ denotes gradient of $g$ with respect to the variable $h \in \mathbb{R}^{n}$ and $H_{h} g$ denotes Hessian of $g$ with respect to $h$.
- The notation a.s. means $\mathbb{P}$-almost surely and $\mathbb{E}$ is expectation with respect to $\mathbb{P}$, where $\mathbb{P}$ is an underlying probability measure.
- For a matrix $M, M^{\text {sym }}=\left(M+M^{T}\right) / 2$ denote its symmetric part, $\lambda_{\min }(M)$ and $\lambda_{\max }(M)$ denote its minimum and maximum eigenvalue respectively, and $\sigma_{\min }(M)$ and $\sigma_{\max }(M)$ denote its minimum and maximum singular value respectively.
- For a vector $v=\left(v^{1}, \ldots, v^{d}\right), \operatorname{diag}(v)$ denotes the diagonal matrix with the $i$ th diagonal entry equal $v^{i}$.
- 1 denotes a vector with all entries equal to one.
- $e_{i j}$ denotes the Kronecker delta.
- $\mathcal{C}(I ; J)$ denotes the space of continuous $J$-valued functions defined on $I$.
- $\mathcal{C}^{2,1}(\mathcal{D} \times I ; J)$ denotes the space of all $J$-valued functions $V(x, t)$ defined on $\mathcal{D} \times I$ which are continuously twice differentiable in $x \in \mathcal{D}$ and once differentiable in $t \in I$.

Next, we recall the RNN models considered in the main paper.
Continuous-Time NRNNs. For a terminal time $T>0$ and an input signal $x=\left(x_{t}\right)_{t \in[0, T]} \in$ $\mathcal{C}\left([0, T] ; \mathbb{R}^{d_{x}}\right)$, the output $y_{t} \in \mathbb{R}^{d_{y}}$, for $t \in[0, T]$, is a linear map of the hidden states $h_{t} \in \mathbb{R}^{d_{h}}$ satisfying the Itô stochastic differential equation (SDE):

$$
\begin{equation*}
\mathrm{d} h_{t}=f\left(h_{t}, x_{t}\right) \mathrm{d} t+\sigma\left(h_{t}, x_{t}\right) \mathrm{d} B_{t}, \quad y_{t}=V h_{t} \tag{23}
\end{equation*}
$$

where $V \in \mathbb{R}^{d_{y} \times d_{h}}, f: \mathbb{R}^{d_{h}} \times \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{h}}, \sigma: \mathbb{R}^{d_{h}} \times \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{h} \times r}$ and $\left(B_{t}\right)_{t \geq 0}$ is an $r$-dimensional Wiener process.
In particular, as an example and for empirical experiments, we focus on the choice of drift function:

$$
\begin{equation*}
f(h, x)=A h+a(W h+U x+b), \tag{24}
\end{equation*}
$$

where $a: \mathbb{R} \rightarrow \mathbb{R}$ is a Lipschitz continuous scalar activation function (such as tanh) extended to act on vectors pointwise, $A, W \in \mathbb{R}^{d_{h} \times d_{h}}, U \in \mathbb{R}^{d_{h} \times d_{x}}$ and $b \in \mathbb{R}^{d_{h}}$, and the choice of diffusion coefficient:

$$
\begin{equation*}
\sigma(h, x)=\epsilon\left(\sigma_{1} I+\sigma_{2} \operatorname{diag}(f(h, x))\right), \tag{25}
\end{equation*}
$$

where the noise level $\epsilon>0$ is small, and $\sigma_{1} \geq 0$ and $\sigma_{2} \geq 0$ are tunable parameters describing the relative strength of additive noise and a multiplicative noise respectively.

We consider the following NRNN models by discretizing the SDE (23), as discussed in detail in the main paper.

Discrete-Time NRNNs. Let $0:=t_{0}<t_{1}<\cdots<t_{M}:=T$ be a partition of the interval $[0, T]$. Denote $\delta_{m}:=t_{m+1}-t_{m}$ for each $m=0,1, \ldots, M-1$, and $\delta:=\left(\delta_{m}\right)$. The Euler-Mayurama (E-M) scheme provides a family (parametrized by $\delta$ ) of approximations to the solution of the SDE in (23):

$$
\begin{equation*}
h_{m+1}^{\delta}=h_{m}^{\delta}+f\left(h_{m}^{\delta}, \hat{x}_{m}\right) \delta_{m}+\sigma\left(h_{m}^{\delta}, \hat{x}_{m}\right) \sqrt{\delta_{m}} \xi_{m} \tag{26}
\end{equation*}
$$

for $m=0,1, \ldots, M-1$, where $\left(\hat{x}_{m}\right)_{m=0, \ldots, M-1}$ is a given sequential data, the $\xi_{m} \sim \mathcal{N}(0, I)$ are independent $r$-dimensional standard normal random vectors, and $h_{0}^{\delta}=h_{0}$. Eq. 26) describes the update equation of our NRNN models, an example of which is when $f$ and $\sigma$ are taken to be (24) and (25) respectively (see also the experiments in the main paper). In the special case when $\epsilon:=0$ in this example, we recover the Lipschitz RNN of [6].

It is worth mentioning that while higher-order integrators are also possible to consider, the presence of Itô white noise poses a significant challenge over the standard ODE case. Generally speaking, implementations of higher-order schemes require additional computational effort which may outweigh the benefit of using them. For instance, in implicit E-M schemes the zero of a nonlinear equation has to be determined in each time step [12]. In Milstein and stochastic Runge-Kutta schemes, there is an extra computational cost in simulating the Lévy area [21]. Similar challenges arise for other multistep schemes and higher-order schemes.

Organizational Details. This SM is organized as follows.

- In Section B we provide results that guarantee existence and uniqueness of solutions to the SDE defining our continuous-time NRNNs.
- In Section C, we provide results that guarantee stability and convergence of our discrete-time NRNNs. These results are in fact very general and may be of independent interest.
- In Section D, we provide results on implicit regularization due to noise injection in both continuous-time and discrete-time NRNNs, in particular the proof of Theorem 1 in the main paper.
- In Section E, we provide some background and results to study classification margin and generalization bound of the corresponding discrete-time deterministic RNNs, in particular the proof of Theorem 2 in the main paper.
- In Section F, we discuss stability of continuous-time NRNNs and the noise-induced stabilization phenomenon, and provide conditions that guarantee almost sure exponential stability of the NRNNs, in particular the proof of Theorem 3 in the main paper.
- In Section G, we provide details on the empirical results in the main paper and additional results.


## B Existence and Uniqueness of Solutions

Essential to any discussion concerning SDEs is the existence and uniqueness of solutions - in this case, we are interested in strong solutions [14].

In the following, we fix a complete filtered probability space $\left(\Omega, \mathcal{F},\left(\mathcal{F}_{t}\right)_{t \geq 0}, \mathbb{P}\right)$ which satisfies the usual conditions [14] and on which there is defined an $r$-dimensional Wiener process $\left(B_{t}\right)_{t \geq 0}$. We also fix a $T>0$ and denote $f\left(h_{t}, t\right):=f\left(h_{t}, x_{t}\right), \sigma\left(h_{t}, t\right):=\sigma\left(h_{t}, x_{t}\right)$ to emphasize the explicit dependence of the functions on time $t$ through the input $x_{t}$.

We start with the following assumptions on the SDE (23).

Assumption B. (a) (Global Lipschitz condition) The coefficients $f$ and $\sigma$ are $L$-Lipschitz, i.e., there exists a constant $L>0$ such that

$$
\begin{equation*}
\left\|f(h, t)-f\left(h^{\prime}, t\right)\right\|+\left\|\sigma(h, t)-\sigma\left(h^{\prime}, t\right)\right\|_{F} \leq L\left\|h-h^{\prime}\right\| \tag{27}
\end{equation*}
$$

for all $h, h^{\prime} \in \mathbb{R}^{d_{h}}$ and $t \in[0, T]$.
(b) (Linear growth condition) $f$ and $\sigma$ satisfy the following linear growth condition, i.e., there exists a constant $K>0$ such that

$$
\begin{equation*}
\|f(h, t)\|^{2}+\|\sigma(h, t)\|_{F}^{2} \leq K\left(1+\|h\|^{2}\right) \tag{28}
\end{equation*}
$$

for all $h \in \mathbb{R}^{d_{h}}$ and $t \in[0, T]$.
Under Assumption B, it is a standard result from stochastic analysis that the SDE (23) has a unique solution (which is a continuous and adapted process $\left(h_{t}\right)_{t \in[0, T]}$ satisfying the integral equation $\left.h_{t}=h_{0}+\int_{0}^{t} f\left(h_{s}, s\right) d s+\int_{0}^{t} \sigma\left(h_{s}, s\right) d B_{s}\right)$ for every initial value $h_{0} \in \mathbb{R}^{d_{h}}$, for $t \in[0, T]$ (see, for instance, Theorem 3.1 in Section 2.3 of [23]). The uniqueness is in the sense that for any other solution $h_{t}^{\prime}$ satisfying the SDE ,

$$
\begin{equation*}
\mathbb{P}\left[h_{t}=h_{t}^{\prime} \text { for all } t \in[0, T]\right]=1 \tag{29}
\end{equation*}
$$

For our purpose, the following conditions suffice to satisfy Assumption B
Assumption C. The function $a: \mathbb{R} \rightarrow \mathbb{R}$ is an activation function (i.e., a non-constant and Lipschitz continuous function), and $\sigma$ is $L_{\sigma}$-Lipschitz for some $L_{\sigma}>0$.
Lemma 1. Consider the SDE (23) defining our CT-NRNN. Then, under Assumption C Assumption B is satisfied.

Proof. Note that $f(h, t)=A h+a\left(W h+U x_{t}+b\right)$, where $a$ is an activation function. For any $t \in[0, T]$,

$$
\begin{align*}
\left\|f(h, t)-f\left(h^{\prime}, t\right)\right\| & \leq\left\|A\left(h-h^{\prime}\right)\right\|+\left\|a\left(W h+U x_{t}+b\right)-a\left(W h^{\prime}+U x_{t}+b\right)\right\|  \tag{30}\\
& \leq\|A\|\left\|h-h^{\prime}\right\|+L_{a}\left\|W\left(h-h^{\prime}\right)\right\|  \tag{31}\\
& \leq\left(\|A\|+L_{a}\|W\|\right)\left\|h-h^{\prime}\right\| \tag{32}
\end{align*}
$$

for all $h, h^{\prime} \in \mathbb{R}^{d_{h}}$, where $L_{a}>0$ is the Lipschitz constant of the (non-constant) activation function $a$. Therefore, the condition (a) in Assumption B is satisfied since by our assumption $\sigma$ is $L_{\sigma}$-Lipschitz for some constant $L_{\sigma}>0$. In this case one can take $L=\max \left(\|A\|+L_{a}\|W\|, L_{\sigma}\right)$ in Eq. 27).
Since $f$ and $\sigma$ are $L$-Lipschitz, they satisfy the linear growth condition (b) in Assumption B Indeed, if $f$ is $L$-Lipschitz, then for $t \in[0, T]$,

$$
\begin{equation*}
\|f(h, t)\|=\|f(h, t)-f(0, t)+f(0, t)\| \leq L\|h\|+\|f(0, t)\| \leq L\|h\|+C_{f} \tag{33}
\end{equation*}
$$

for some constant $C_{f} \in(0, \infty)$, where we have used the fact that $f(0, t)=a\left(U x_{t}+b\right)$ is bounded for $t \in[0, T]$ (since continuous functions on compact sets are bounded). So,

$$
\begin{equation*}
\|f(h, t)\|^{2} \leq\left(L\|h\|+C_{f}\right)^{2} \leq L^{2}\|h\|^{2}+2 L C_{f}\|h\|+C_{f}^{2} \tag{34}
\end{equation*}
$$

For $\|h\| \geq 1$, we have:

$$
\begin{equation*}
\|f(h, t)\|^{2} \leq\left(L^{2}+2 L C_{f}\right)\|h\|^{2}+C_{f}^{2} \leq\left(L^{2}+2 L C_{f}+C_{f}^{2}\right)\left(1+\|h\|^{2}\right) \tag{35}
\end{equation*}
$$

For $\|h\|<1$, we have:

$$
\begin{equation*}
\|f(h, t)\|^{2} \leq\left(L^{2}+2 L C_{f}\right)\|h\|+C_{f}^{2} \leq\left(L^{2}+2 L C_{f}\right)+C_{f}^{2} \leq\left(L^{2}+2 L C_{f}+C_{f}^{2}\right)\left(1+\|h\|^{2}\right) \tag{36}
\end{equation*}
$$

Choosing $K=L^{2}+2 L C_{f}+C_{f}^{2}$ gives us the linear growth condition for $f$.
Similarly, one can show that $\sigma$ satisfies the linear growth condition. The proof is done.
Throughout the paper, we work with SDEs satisfying Assumption C. The following additional assumption on the SDEs will be needed and invoked.

Assumption D. For $t \in[0, T]$, the partial derivatives of the coefficients $f^{i}(h, t), \sigma^{i j}(h, t)$ with respect to $h$ up to order three (inclusive) exist. Moreover, the coefficients $f^{i}(h, t), \sigma^{i j}(h, t)$ and all these partial derivatives are:
(i) bounded and Borel measurable in $t$, for fixed $h \in \mathbb{R}^{d_{h}}$;
(ii) Lipschitz continuous in $h$, for fixed $t \in[0, T]$.

In particular, Assumption Dimplies that these partial derivatives (with respect to $h$ ) of $f$ and $\sigma$ satisfy (a)-(b) in Assumption B Assumption D holds for SDEs with commonly used activation functions such as hyperbolic tangent. We remark that Assumption C.D may be weakened in various directions (for instance, to locally Lipschitz coefficients) but for the purpose of this paper we need not go beyond these assumptions.

## C Stability and Convergence of the Euler-Maruyama Schemes

We provide stability and strong convergence results for the explicit Euler-Mayurama (E-M) approximations of the SDE (23), which is time-inhomogeneous due to the dependence of the drift and possibly diffusion coefficient on a time-varying input, here. Intuitively, strong convergence results ensure that the approximated path follows the continuous path accurately, in contrast to weak convergence results which can only guarantee this at the level of probability distribution. The latest version of strong convergence results for time-homogeneous SDEs can be found in [7, 8]. The results for our time-inhomogeneous SDEs can be obtained by adapting the proof in [7] without much difficulty. Since we cannot find them in the literature, we provide them in this section.

First, we recall the discretization scheme. Let $0:=t_{0}<t_{1}<\cdots<t_{M}:=T$ and $t_{m+1}=t_{m}+\delta_{m}$, for $m=0,1, \ldots, M-1$ and some time step $\delta_{m}>0$. Note that we work at full generality here since the step sizes $\delta_{m}$ are not necessarily uniform and may even depend on the numerical solution, i.e., $\delta_{m}=\delta\left(h_{m}^{\delta}\right)$ (see Example 1). The general results will be of independent interest, in particular for further explorations in designing other variants of NRNNs.

For $m=0,1, \ldots, M-1$, consider

$$
\begin{equation*}
h_{m+1}^{\delta}=h_{m}^{\delta}+f\left(h_{m}^{\delta}, \hat{x}_{m}\right) \delta_{m}+\sigma\left(h_{m}^{\delta}, \hat{x}_{m}\right) \Delta B_{m}, \tag{37}
\end{equation*}
$$

where $\Delta B_{m}:=B_{t_{m+1}}-B_{t_{m}},\left(\hat{x}_{m}\right)_{m=0,1, \ldots, M-1}$ is a given input sequential data, and $h_{0}^{\delta}=h_{0}$.
Let $\underline{t}=\max \left\{t_{m}: t_{m} \leq t\right\}, m_{t}=\max \left\{m: t_{m} \leq t\right\}$ for the nearest time point before time $t$, and its index. Denote the piecewise constant interpolant process $\bar{h}_{t}=h_{\underline{t}}^{\delta}$. It is convenient to use continuous-time approximations, so we consider the continuous interpolant that satisfies:

$$
\begin{equation*}
h_{t}^{\delta}=h_{\underline{t}}^{\delta}+f\left(h_{\underline{t}}^{\delta}, x_{\underline{t}}\right)(t-\underline{t})+\sigma\left(h_{\underline{t}}^{\delta}, x_{\underline{t}}\right)\left(B_{t}-B_{\underline{t}}\right) \tag{38}
\end{equation*}
$$

so that $h_{t}^{\delta}$ is the solution of the SDE:

$$
\begin{equation*}
d h_{t}^{\delta}=f\left(h_{\underline{t}}^{\delta}, x_{\underline{t}}\right) d t+\sigma\left(h_{\underline{t}}^{\delta}, x_{\underline{t}}\right) d B_{t}=f\left(\bar{h}_{t}, x_{t}\right) d t+\sigma\left(\bar{h}_{t}, x_{t}\right) d B_{t} . \tag{39}
\end{equation*}
$$

We make the following assumptions about the time step.
Assumption E. The (possibly adaptive) time step function $\delta: \mathbb{R}^{d_{h}} \rightarrow \mathbb{R}^{+}$is continuous and strictly positive, and there exist constants $\alpha, \beta>0$ such that for all $h \in \mathbb{R}^{d_{h}}, \delta$ satisfies

$$
\begin{equation*}
\langle h, f(h, t)\rangle+\frac{1}{2} \delta(h)\|f(h, t)\|^{2} \leq \alpha\|h\|^{2}+\beta \tag{40}
\end{equation*}
$$

for every $t \in[0, T]$.
Note that if another time step function $\delta^{\epsilon}(h)$ is smaller than $\delta(h)$, then $\delta^{\epsilon}(h)$ also satisfies Assumption E.

A simple adaptation of the proof of Theorem 2.1.1 in [7] to our case of time-inhomogeneous SDE gives the following result.

Proposition 1 (Finite-time stability). Under Assumption $B$ and Assumption $E T$ is a.s. attainable (i.e., for $\omega \in \Omega, \mathbb{P}\left[\exists N(\omega)<\infty\right.$ s.t. $\left.t_{N(\omega)} \geq T\right]=1$ ) and for all $p>0$ there exists a constant $C>0$ (depending on only $p$ and $T$ ) such that

$$
\begin{equation*}
\mathbb{E}\left[\sup _{t \in[0, T]}\left\|h_{t}^{\delta}\right\|^{p}\right] \leq C \tag{41}
\end{equation*}
$$

This is the discrete-time analogue of the result that $\mathbb{E}\left[\sup _{t \in[0, T]}\left\|h_{t}\right\|^{p}\right]<\infty$ for all $p>0$, which can be proven by simply adapting the proof of Lemma 2.1.1. in [7].
In the case where the time step is adaptive, we take the following lower bound on the time step to bound the expected number of time steps (how quickly $\delta(h) \rightarrow 0$ as $\|h\| \rightarrow 0$ ).
Assumption F. There exist constants $a, b, q>0$ such that the adaptive time step function satisfies:

$$
\begin{equation*}
\delta(h) \geq \frac{1}{a\|h\|^{q}+b} \tag{42}
\end{equation*}
$$

Next, we provide strong convergence result for the numerical approximation with the time step $\delta$. When the time step $\delta$ is adaptive, one needs to rescale the time step function by a small scalar-valued magnitude $\epsilon>0$ and then consider the limit as $\epsilon \rightarrow 0$. Following [7], we make the following assumption.
Assumption G. The rescaled time step function $\delta^{\epsilon}$ satisfies

$$
\begin{equation*}
\epsilon \min (T, \delta(h)) \leq \delta^{\epsilon}(h) \leq \min (\epsilon T, \delta(h)) \tag{43}
\end{equation*}
$$

where $\delta$ satisfies Assumption E|F
Under this additional assumption, we have the following convergence result, which can be proven by adapting the proof of Theorem 2.1.2 in [7] to our time-inhomogeneous SDE case. The proof is based on the argument used for the uniform time step analysis (see Theorem 2.2 in [11]), taking into account the adaptive nature of the time step appropriately.
Theorem 4 (Strong convergence). Let the $S D E \sqrt{23}$ ) satisfy Assumption $B$ and the time step function satisfy Assumption $G$ Then, for all $p>0$,

$$
\begin{equation*}
\lim _{\epsilon \rightarrow 0} \mathbb{E}\left[\sup _{t \in[0, T]}\left\|h_{t}^{\delta^{\epsilon}}-h_{t}\right\|^{p}\right]=0 \tag{44}
\end{equation*}
$$

where $h_{t}^{\delta}$ is the continuous interpolant satisfying (38) and $h_{t}$ satisfies the SDE (23).
In particular, the non-adaptive time stepping scheme satisfies the above assumptions. Therefore, stability and strong convergence of the schemes are guaranteed by the above results.
Under stronger assumptions on the drift $f$ we can obtain the order of strong convergence for the numerical schemes; see Theorem 2.1.3 in [7] for the case of time-homogeneous SDEs. This result can be adapted to our case to obtain order- $\frac{1}{2}$ strong convergence, which is also obtained in the special case when the step sizes are uniform (see Theorem 10.2.2 in [16]).
Theorem 5 (Strong convergence rate). Assume that $f$ satisfies the following one-sided Lipschitz condition, i.e., there exists a constant $\alpha>0$ such that for all $h, h^{\prime} \in \mathbb{R}^{d_{h}}$,

$$
\begin{equation*}
\left\langle h-h^{\prime}, f(h, t)-f\left(h^{\prime}, t\right)\right\rangle \leq \alpha\left\|h-h^{\prime}\right\|^{2} \tag{45}
\end{equation*}
$$

for all $t \in[0, T]$, and the following locally polynomial growth Lipschitz condition, i.e., there exists $\gamma, \mu, q>0$ such that for all $h, h^{\prime} \in \mathbb{R}^{d_{h}}$,

$$
\begin{equation*}
\left\|f(h, t)-f\left(h^{\prime}, t\right)\right\| \leq\left(\gamma\left(\|h\|^{q}+\left\|h^{\prime}\right\|^{q}\right)+\mu\right)\left\|h-h^{\prime}\right\| \tag{46}
\end{equation*}
$$

for all $t \in[0, T]$. Moreover, assume that $\sigma$ is globally Lipschitz and the time step function satisfies Assumption $G$ Then, for all $p>0$, there exists a constant $C>0$ such that

$$
\begin{equation*}
\mathbb{E}\left[\sup _{t \in[0, T]}\left\|h_{t}^{\delta^{\epsilon}}-h_{t}\right\|^{p}\right] \leq C \epsilon^{p / 2} \tag{47}
\end{equation*}
$$

Lastly, it is worth mentioning the following adaptive scheme, which may be a useful option when designing NRNNs.
Example 1 (Adaptive E-M). Under the same setup as the classical E-M setting, we may also introduce an adaptive step size scheme through a sequence of random vectors $d_{m}$. In this case,

$$
\begin{equation*}
h_{m+1}^{\delta}=h_{m}^{\delta}+d_{m} \odot f\left(h_{m}^{\delta}, \hat{x}_{m}\right) \Delta t_{m}+\sigma\left(h_{m}^{\delta}, \hat{x}_{m}\right)\left(\Delta t_{m}\right)^{1 / 2} \xi_{m}, \tag{48}
\end{equation*}
$$

where $\odot$ denotes the pointwise (Hadamard) product, and each $d_{n}$ may be dependent on $h_{i}^{\delta}$ for $i \leq m$. Provided that $d_{m} \rightarrow(1, \ldots, 1)$ uniformly almost surely as $\Delta t \rightarrow 0$, one could also obtain the same convergence as the classical E-M case. The adaptive setting allows for potentially better approximations by shrinking step sizes in places where the solution changes rapidly. An intuitive explanation for the instability of the standard E-M approximation of SDEs is that there is always a very small probability of a large Brownian increment which causes the approximation to produce a solution with undesirable growth. Using an adaptive time step eliminates this problem. Moreover, this scheme includes, in appropriate sense, the stochastic depth in [20] (see page 9 there) and the dropout in [19] as special cases upon choosing an appropriate $d_{m}$ and $\sigma$.
In particular, one can consider the following drift-tamed E-M scheme, where all components, $d_{m}^{i}$, of the elements of the sequence are generated as a function of $h_{m}^{\delta}$, i.e.,

$$
\begin{equation*}
d_{m}=\frac{1}{\max \left\{1, c_{1}\left\|h_{m}^{\delta}\right\|+c_{2}\right\}} \mathbf{1} \tag{49}
\end{equation*}
$$

for some $c_{1}, c_{2}>0$. In this way, the drift term is "tamed" by a solution-dependent multiplicative factor no larger than one, which prevents the hidden state in the next time step from becoming too large. This adaptive scheme is related to the one introduced in [13] to provide an explicit numerical method that would display strong convergence in circumstances where the standard E-M method does not. Under certain conditions strong convergence of this scheme can be proven (even for SDEs with superlinearly growing drift coefficients). Other adaptive schemes include the increment-tamed scheme of [12] and many others.

## D Implicit Regularization in NRNNs

As discussed in the main paper, although the learning is carried out in discrete time, it is worth studying the continuous-time setting. The results for the continuous-time case may provide alternative perspectives and, more importantly, will be useful as a reference for exploring other discretization schemes for the CT-NRNNs. In Subsection D.1, we study implicit regularization for the continuoustime NRNNs. In Subsection D.2, we study implicit regularization for discrete-time NRNNs and comment on the difference between the continuous-time and discrete-time case.

We remark that the approach presented here is standard in showing implicit regularization. The essence of the approach is to view NRNN as a training scheme for the deterministic RNN. Also, note that had one attempted to conduct an analysis based on NRNN directly, the resulting bound would be stochastic due to the presence of the diffusion term and it is not clear how this bound helps explaining implicit regularization.

## D. 1 Continuous-Time Setting

Main Result and Discussions. For the sake of brevity, we denote $f_{t}(\cdot):=f\left(\cdot, x_{t}\right)$ and $\sigma_{t}(\cdot):=$ $\sigma\left(\cdot, x_{t}\right)$ for $t \in[0, T]$ in the following.
To begin, consider the process $\left(\bar{h}_{t}\right)_{t \in[0, T]}$ satisfying the following initial value problem (IVP):

$$
\begin{equation*}
d \bar{h}_{t}=f_{t}\left(\bar{h}_{t}\right) d t, \quad \bar{h}_{0}=h_{0} \tag{50}
\end{equation*}
$$

Let $\Psi$ denote the unique fundamental matrix satisfying the following properties: for $0 \leq s \leq u \leq$ $t \leq T$,
(a)

$$
\begin{equation*}
\frac{\partial \Psi(t, s)}{\partial t}=\frac{\partial f_{t}}{\partial \bar{h}}\left(\bar{h}_{t}\right) \Psi(t, s) ; \quad \frac{\partial \Psi(t, s)}{\partial s}=-\Psi(t, s) \frac{\partial f_{t}}{\partial \bar{h}}\left(\bar{h}_{t}\right) ; \tag{51}
\end{equation*}
$$

(b) $\Psi(t, s)=\Psi(t, u) \Psi(u, s)$;
(c) $\Psi(t, s)=\Psi^{-1}(s, t)$;
(d) $\Psi(s, s)=I$.

Also, let $\Sigma(t, s):=\Psi(t, s) \sigma_{s}\left(\bar{h}_{s}\right)$ for $0 \leq s \leq t \leq T$.
The following result links the expected loss function used for training CT-NRNNs to that for training deterministic CT-RNNs when the noise amplitude is small.
Theorem 6 (Explicit regularization induced by noise injection for CT-NRNNs). Under Assumption A in the main paper,

$$
\begin{equation*}
\mathbb{E} \ell\left(h_{T}\right)=\ell\left(\bar{h}_{T}\right)+\frac{\epsilon^{2}}{2}[Q(\bar{h})+R(\bar{h})]+\mathcal{O}\left(\epsilon^{3}\right) \tag{52}
\end{equation*}
$$

as $\epsilon \rightarrow 0$, where $Q$ and $R$ are given by

$$
\begin{align*}
& Q(\bar{h})=\left(\nabla l\left(\bar{h}_{T}\right)\right)^{T} \int_{0}^{T} \mathrm{~d} s \Psi(T, s) \int_{0}^{s} \mathrm{~d} u v(u)+\left(\nabla l\left(\bar{h}_{T}\right)\right)^{T} \int_{0}^{T} \mathrm{~d} s w(s)  \tag{53}\\
& R(\bar{h})=\int_{0}^{T} \mathrm{~d} s \operatorname{tr}\left(\Sigma(T, s) \Sigma(T, s)^{\top} \nabla^{2} \ell\left(\bar{h}_{T}\right)\right) \tag{54}
\end{align*}
$$

with $v(u)$ a vector with the pth component $\left(p=1,2, \ldots, d_{h}\right)$ :

$$
\begin{equation*}
v^{p}(u)=\operatorname{tr}\left(\Sigma(s, u) \Sigma^{T}(s, u) \nabla^{2}\left[f_{s}\right]^{p}\left(\bar{h}_{s}\right)\right), \tag{55}
\end{equation*}
$$

and $w(s)$ a vector with the qth component $\left(q=1,2, \ldots, d_{h}\right)$ :

$$
\begin{equation*}
w^{q}(s)=\sum_{k=1}^{r} \sum_{j, l=1}^{d_{h}} \Psi_{k, l}^{q j}(T, s) \partial_{l} \sigma_{s}^{j k}\left(\bar{h}_{s}\right) \sigma_{s}^{l k}\left(\bar{h}_{s}\right) . \tag{56}
\end{equation*}
$$

Therefore, to study the difference between the CT-NRNNs and their deterministic version, it remains to investigate the role of $Q$ and $R$ in Theorem 6 . If the Hessian is positive semi-definite, then $R(\bar{h})$ is also positive semi-definite and thus a viable regularizer. On the other hand, $Q(\bar{h})$ need not be non-negative. However, by assuming that $\nabla^{2} f$ and $\nabla \sigma^{i j}$ are small (that is, $f$ is approximately linear and $\sigma$ relatively independent of $\bar{h}$ ), then $Q$ can be perceived negligible and we may focus predominantly on $R$. An argument of this kind was used in [4] in the context of Gauss-Newton Hessian approximations. In particular, $Q=0$ for linear NRNNs with additive noise. Therefore, Theorem 6 essentially tells us that injecting noise to deterministic RNN is approximately equivalent to considering a regularized objective functional. Moreover, the explicit regularizer is solely determined by the flow generated by the Jacobian $\frac{\partial f_{t}}{\partial h}\left(\bar{h}_{t}\right)$, the diffusion coefficient $\sigma_{t}$ and the Hessian of the loss function, all evaluated along the dynamics of the deterministic RNN.
Under these assumptions, ignoring higher-order terms and bounding the Frobenius inner product in (54), we can interpret training with CT-NRNN as an approximation of the following optimal control problem [29] with the running cost $C(t):=\frac{1}{2} \operatorname{tr}\left(\sigma_{t}\left(\bar{h}_{t}\right)^{T} \Psi(T, t)^{T} \nabla^{2} \ell\left(\bar{h}_{t}\right) \Psi(T, t) \sigma_{t}\left(\bar{h}_{t}\right)\right)$ :

$$
\begin{align*}
& \min \mathbb{E}_{(\boldsymbol{x}, y) \sim \mu}\left[\ell\left(\bar{h}_{T}\right)+\epsilon^{2} \int_{0}^{T} C(t) d t\right]  \tag{57}\\
& \text { s.t. } \mathrm{d} \bar{h}_{t}=f_{t}\left(\bar{h}_{t}\right) \mathrm{d} t, t \in[0, T], \bar{h}_{0}=h_{0} \tag{58}
\end{align*}
$$

where $\left(\boldsymbol{x}:=\left(x_{t}\right)_{t \in[0, T]}, y\right)$ denotes a training example drawn from the distribution $\mu$ and the minimization is with respect to the parameters (controls) in the corresponding deterministic RNN. On the other hand, we can interpret training with the deterministic RNN as the above optimal control problem with zero running cost or regularization. Note that if the Hessian matrix is symmetric positive semi-definite, then $C(t)$ is a quadratic form with the associated metric tensor $M_{t}^{T} M_{t}:=\nabla^{2} \ell\left(\bar{h}_{t}\right)$ and

$$
\begin{equation*}
C(t)=\frac{1}{2}\left\langle\Psi(T, t) \sigma_{t}, \Psi(T, t) \sigma_{t}\right\rangle_{M_{t}}=\frac{1}{2}\left\|M_{t} \Psi(T, t) \sigma_{t}\right\|_{F}^{2} \leq \frac{1}{2}\left\|\sigma_{t}\right\|_{F}^{2}\left\|M_{t}\right\|_{F}^{2}\|\Psi(T, t)\|_{F}^{2} . \tag{59}
\end{equation*}
$$

Overall, we can see that the use of NRNNs as a regularization mechanism reduces the fundamental matrices $\Psi(T, s)$ according to the magnitude of the elements of $\sigma_{t}$.

Proof of Theorem 6. Next, we prove Theorem6 We will need some auxiliary results before doing so.

For a small perturbation parameter $\epsilon>0$, the hidden states now satisfy the SDE

$$
\mathrm{d} h_{t}=f_{t}\left(h_{t}\right) \mathrm{d} t+\epsilon \sigma_{t}\left(h_{t}\right) \mathrm{d} B_{t}
$$

where we have used the shorthand $f_{t}(\cdot)=f\left(\cdot, x_{t}\right)$ and $\sigma_{t}(\cdot)=\sigma\left(\cdot, x_{t}\right)$. To investigate the effect of the perturbation, consider the following hierarchy of differential equations:

$$
\begin{align*}
\mathrm{d} h_{t}^{(0)} & =f_{t}\left(h_{t}^{(0)}\right) \mathrm{d} t  \tag{60}\\
\mathrm{~d} h_{t}^{(1)} & =\frac{\partial f_{t}}{\partial h}\left(h_{t}^{(0)}\right) h_{t}^{(1)} \mathrm{d} t+\sigma_{t}\left(h_{t}^{(0)}\right) \mathrm{d} B_{t}  \tag{61}\\
\mathrm{~d} h_{t}^{(2)} & =\frac{\partial f_{t}}{\partial h}\left(h_{t}^{(0)}\right) h_{t}^{(2)} \mathrm{d} t+\Phi_{t}^{(1)}\left(h_{t}^{(0)}, h_{t}^{(1)}\right) \mathrm{d} t+\Phi_{t}^{(2)}\left(h_{t}^{(0)}, h_{t}^{(1)}\right) \mathrm{d} B_{t} \tag{62}
\end{align*}
$$

with $h_{0}^{(0)}=h_{0}, h_{0}^{(1)}=0$, and $h_{0}^{(2)}=0$, and where

$$
\begin{align*}
\Phi_{t}^{(1)}\left(h_{0}, h_{1}\right) & =\frac{1}{2} \sum_{i, j} \frac{\partial^{2} f_{t}}{\partial h^{i} \partial h^{j}}\left(h_{0}\right) h_{1}^{i} h_{1}^{j}  \tag{63}\\
\Phi_{t}^{(2)}\left(h_{0}, h_{1}\right) & =\sum_{i} \frac{\partial \sigma_{t}}{\partial h^{i}}\left(h_{0}\right) h_{1}^{i} . \tag{64}
\end{align*}
$$

In the sequel, we will suppose Assumption A in the main paper (which is equivalent to Assumption B and Assumption D) holds. Under this assumption, each of these initial value problems have a unique solution for $t \in[0, T]$. The processes $h_{t}^{(0)}, h_{t}^{(1)}$ and $h_{t}^{(2)}$ denote the zeroth-, first-, and second-order terms in an expansion of $h_{t}$ about $\epsilon=0$. This can be easily seen using Kunita's theory of stochastic flows. In particular, by Theorem 3.1 in [17], letting $h_{\epsilon, t}^{(1)}=\frac{\partial h_{t}}{\partial \epsilon}$, we find that

$$
\mathrm{d} h_{\epsilon, t}^{(1)}=\frac{\partial f_{t}}{\partial h}\left(h_{t}\right) h_{\epsilon, t}^{(1)} \mathrm{d} t+\left(\sigma_{t}\left(h_{t}\right)+\epsilon \Phi_{t}^{(2)}\left(h_{t}, h_{\epsilon, t}^{(1)}\right)\right) \mathrm{d} B_{t},
$$

and so we find that $h_{0, t}^{(1)}=h_{t}^{(1)}$. Similarly, $h_{\epsilon, t}^{(2)}=\frac{\partial^{2} h_{t}}{\partial \epsilon^{2}}=\frac{\partial h_{\epsilon, t}^{(1)}}{\partial \epsilon}$ can be shown to satisfy

$$
\begin{aligned}
\mathrm{d} h_{\epsilon, t}^{(2)}= & \frac{\partial f_{t}}{\partial h}\left(h_{t}\right) h_{\epsilon, t}^{(2)} \mathrm{d} t+2 \Phi_{t}^{(1)}\left(h_{t}, h_{\epsilon, t}^{(1)}\right) \mathrm{d} t \\
& +\left(2 \Phi_{t}^{(2)}\left(h_{t}, h_{\epsilon, t}^{(1)}\right)+\epsilon \sum_{k}\left[h_{\epsilon, t}^{(1)}\right]^{k} \frac{\partial}{\partial h^{k}} \Phi_{t}^{(2)}\left(h_{t}, h_{\epsilon, t}^{(1)}\right)\right) \mathrm{d} B_{t} .
\end{aligned}
$$

This equation is obtained by applying Theorem 3.1 in [17] to find the first derivative of the system $\left(h_{t}, h_{\epsilon, t}^{(1)}\right)$ with respect to $\epsilon$ and projecting to the second coordinate. Taking $\epsilon=0$, we find that $h_{0, t}^{(2)}=2 h_{t}^{(2)}$. Therefore, informally, a pathwise second-order Taylor expansion about $\epsilon=0$ reveals that $h_{t}=h_{t}^{(0)}+\epsilon h_{t}^{(1)}+\epsilon^{2} h_{t}^{(2)}+\mathcal{O}\left(\epsilon^{3}\right)$. To formalize this statement, we will later bound the third-order error term in Lemma 3
While the equation for $h_{t}^{(0)}$ is not explicitly solvable, both $h_{t}^{(1)}$ and $h_{t}^{(2)}$ are. In particular, for $t \in[0, T]$ (see Eq. (4.28) in [25]):

$$
\begin{align*}
& h_{t}^{(1)}=\int_{0}^{t} \Psi(t, s) \sigma_{s}\left(h_{s}^{(0)}\right) \mathrm{d} B_{s}=\int_{0}^{t} \Sigma(t, s) \mathrm{d} B_{s},  \tag{65}\\
& h_{t}^{(2)}=\int_{0}^{t} \Psi(t, s) \Phi_{s}^{(1)}\left(h_{s}^{(0)}, h_{s}^{(1)}\right) \mathrm{d} s+\int_{0}^{t} \Psi(t, s) \Phi_{s}^{(2)}\left(h_{s}^{(0)}, h_{s}^{(1)}\right) \mathrm{d} B_{s} . \tag{66}
\end{align*}
$$

The key result needed to prove Theorem6 is contained in the following theorem. In the sequel, big $\mathcal{O}$ notation is to be understood in the almost sure sense.
Theorem 7. For a scalar-valued loss function $\ell \in \mathcal{C}^{2}\left(\mathbb{R}^{d_{h}}\right)$, for $t \in[0, T]$,
$\ell\left(h_{t}\right)=\ell\left(h_{t}^{(0)}\right)+\epsilon \nabla \ell\left(h_{t}^{(0)}\right) \cdot h_{t}^{(1)}+\epsilon^{2}\left(\nabla \ell\left(h_{t}^{(0)}\right) \cdot h_{t}^{(2)}+\frac{1}{2}\left(h_{t}^{(1)}\right)^{\top} \nabla^{2} \ell\left(h_{t}^{(0)}\right)\left(h_{t}^{(1)}\right)\right)+\mathcal{O}\left(\epsilon^{3}\right)$,
as $\epsilon \rightarrow 0$.

We now prove Theorem 7 The proof relies on two lemmas. The first bounds the solutions $h^{(i)}$ over $[0, T]$.
Lemma 2. For any $p>0, \sup _{s \in[0, T]}\left\|h_{s}^{(0)}\right\|^{p}<\infty$ and $\mathbb{E} \sup _{s \in[0, T]}\left\|h_{s}^{(i)}\right\|^{p}<\infty$ for $i=1,2$.
Proof. For $s \in[0, T], h_{s}^{(0)}=h_{0}+\int_{0}^{s} f_{u}\left(h_{u}^{(0)}\right) \mathrm{d} u$, so recalling $(x+y)^{2} \leq 2 x^{2}+2 y^{2}$ and $\left\|f_{t}(h)\right\|^{2} \leq$ $K\left(1+\|h\|^{2}\right)$,

$$
\begin{aligned}
\left\|h_{s}^{(0)}\right\|^{2} & \leq 2\left\|h_{0}\right\|^{2}+2 \int_{0}^{s}\left\|f_{u}\left(h_{u}^{(0)}\right)\right\|^{2} \mathrm{~d} u \\
& \leq 2\left(\left\|h_{0}\right\|^{2}+K^{2} s+K^{2} \int_{0}^{s}\left\|h_{u}^{(0)}\right\|^{2} \mathrm{~d} u\right)
\end{aligned}
$$

Therefore, by Gronwall's inequality,

$$
\begin{aligned}
\left\|h_{s}^{(0)}\right\|^{2} & \leq 2\left(\left\|h_{0}\right\|^{2}+K^{2} s\right) e^{2 K^{2} s} \\
& \leq 2\left(\left\|h_{0}\right\|^{2}+K^{2} T\right) e^{2 K^{2} T}<+\infty
\end{aligned}
$$

and so $\sup _{s \in[0, T]}\left\|h_{s}^{(0)}\right\|<\infty$. Similarly, for $s \in[0, T]$,

$$
h_{s}^{(1)}=\int_{0}^{s} \frac{\partial f_{u}}{\partial h}\left(h_{u}^{(0)}\right) h_{u}^{(1)} \mathrm{d} u+\int_{0}^{s} \sigma_{u}\left(h_{u}^{(0)}\right) \mathrm{d} B_{u}
$$

Therefore, for $p \geq 2$ (since $(x+y)^{p} \leq 2^{p-1}\left(x^{p}+y^{p}\right)$ by Jensen's inequality):

$$
\left\|h_{s}^{(1)}\right\|^{p} \leq 2^{p-1} \int_{0}^{s}\left\|\frac{\partial f_{u}}{\partial h}\left(h_{u}^{(0)}\right)\right\|^{p}\left\|h_{u}^{(1)}\right\|^{p} \mathrm{~d} u+2^{p-1}\left\|\int_{0}^{s} \sigma_{u}\left(h_{u}^{(0)}\right) \mathrm{d} B_{u}\right\|^{p}
$$

Because the Itô integral is a continuous martingale, the Burkholder-Davis-Gundy inequality (see Theorem 3.28 in [14]) implies that for positive constants $C_{p}$ depending only on $p$ (but not necessarily the same in each appearance),

$$
\mathbb{E} \sup _{s \in[0, T]}\left\|h_{s}^{(1)}\right\|^{p} \leq C_{p} \int_{0}^{s} \mathbb{E} \sup _{s \in[0, u]}\left\|h_{s}^{(1)}\right\|^{p} \mathrm{~d} u+C_{p}\left(\int_{0}^{t}\left\|\sigma_{u}\left(h_{u}^{(0)}\right)\right\|^{2} \mathrm{~d} u\right)^{p / 2}
$$

An application of Gronwall's inequality yields

$$
\mathbb{E} \sup _{s \in[0, T]}\left\|h_{s}^{(1)}\right\|^{p} \leq C_{p}\left(\int_{0}^{T}\left\|\sigma_{u}\left(h_{u}^{(0)}\right)\right\|^{2} \mathrm{~d} u\right)^{p / 2} e^{C_{p} T}
$$

Therefore, $\mathbb{E} \sup _{s \in[0, T]}\left\|h_{s}^{(1)}\right\|^{p}<\infty$ for all $p \geq 2$. The $p \in(0,2)$ case follows from Hölder's inequality. Repeating this same approach for $h_{s}^{(2)}$ completes the proof.

The second of our two critical lemmas provides a pathwise expansion of $h_{t}$ about $\epsilon$ in the vein of [3]. Doing so characterizes the response of the NRNN hidden states to small noise perturbations at the sample path level. It can be seen as a strengthening of Theorem 2.2 in [ 9 ] for our time-inhomogeneous SDEs.
Lemma 3. For a fixed $\epsilon_{0}>0$, and any $0<\epsilon \leq \epsilon_{0}$, with probability one,

$$
h_{t}=h_{t}^{(0)}+\epsilon h_{t}^{(1)}+\epsilon^{2} h_{t}^{(2)}+\epsilon^{3} R_{3}^{\epsilon}(t)
$$

where for any $p>0$,

$$
\begin{equation*}
\sup _{\epsilon \in\left(0, \epsilon_{0}\right)} \mathbb{E} \sup _{t \in[0, T]}\left\|R_{3}^{\epsilon}(t)\right\|^{p}<\infty \tag{67}
\end{equation*}
$$

Proof. It suffices to show that $\sup _{\epsilon \in\left(0, \epsilon_{0}\right)} \mathbb{E} \sup _{t \in[0, T]}\left\|R_{3}^{\epsilon}(t)\right\|^{p}<\infty$ for $p \geq 2-$ the $p \in(0,2)$ case follows from Hölder's inequality. In the sequel, we shall let $K$ denote a finite number (not necessarily the same in each appearance) depending only on $f, \sigma, T, \epsilon_{0}$, and $p$, and therefore independent of $t, \epsilon$.

For $\epsilon>0$, let $h_{t}^{\epsilon}=h_{t}^{(0)}+\epsilon h_{t}^{(1)}+\epsilon^{2} h_{t}^{(2)}$ and $R_{3}(t)=\epsilon^{-3}\left(h_{t}-h_{t}^{\epsilon}\right)$, where $h_{t}, h_{t}^{(1)}, h_{t}^{(2)}$ are coupled together through the same Brownian motion. Then

$$
\begin{aligned}
\epsilon^{3} R_{3}(t)= & \int_{0}^{t}\left(f_{s}\left(h_{s}\right)-f_{s}\left(h_{s}^{(0)}\right)-\epsilon \frac{\partial f_{s}}{\partial h}\left(h_{s}^{(0)}\right) h_{s}^{(1)}-\epsilon^{2} \frac{\partial f_{s}}{\partial h}\left(h_{s}^{(0)}\right) h_{s}^{(2)}-\epsilon^{2} \Phi_{s}^{(1)}\left(h_{s}^{(0)}, h_{s}^{(1)}\right)\right) \mathrm{d} s \\
& +\epsilon \int_{0}^{t} \sigma_{s}\left(h_{s}\right)-\sigma_{s}\left(h_{s}^{(0)}\right)-\epsilon \Phi_{s}^{(2)}\left(h_{s}^{(0)}, h_{s}^{(1)}\right) \mathrm{d} B_{s}
\end{aligned}
$$

To simplify, we decompose $\epsilon^{3} R_{3}(t)$ into the sum of four random variables $\theta_{i}(t), i=1, \ldots, 4$, given by

$$
\begin{aligned}
& \theta_{1}(t)=\int_{0}^{t}\left[f_{s}\left(h_{s}\right)-f_{s}\left(h_{s}^{\epsilon}\right)\right] \mathrm{d} s \\
& \theta_{2}(t)=\int_{0}^{t}\left[f_{s}\left(h_{s}^{\epsilon}\right)-f_{s}\left(h_{s}^{(0)}\right)-\epsilon \frac{\partial f_{s}}{\partial h}\left(h_{s}^{(0)}\right) h_{s}^{(1)}-\epsilon^{2} \frac{\partial f_{s}}{\partial h}\left(h_{s}^{(0)}\right) h_{s}^{(2)}-\epsilon^{2} \Phi_{s}^{(1)}\left(h_{s}^{(0)}, h_{s}^{(1)}\right)\right] \mathrm{d} s \\
& \theta_{3}(t)=\epsilon \int_{0}^{t}\left[\sigma_{s}\left(h_{s}\right)-\sigma_{s}\left(h_{s}^{(0)}+\epsilon h_{s}^{(1)}\right)\right] \mathrm{d} B_{s} \\
& \theta_{4}(t)=\epsilon \int_{0}^{t}\left[\sigma_{s}\left(h_{s}^{(0)}+\epsilon h_{s}^{(1)}\right)-\sigma_{s}\left(h_{s}^{(0)}\right)-\epsilon \Phi_{s}^{(2)}\left(h_{s}^{(0)}, h_{s}^{(1)}\right)\right] \mathrm{d} B_{s} .
\end{aligned}
$$

Beginning with the more straightforward terms $\theta_{1}(t), \theta_{3}(t)$, by Lipschitz continuity of $f$,

$$
\left\|f_{s}\left(h_{s}\right)-f_{s}\left(h_{s}^{\epsilon}\right)\right\| \leq L_{f} \epsilon^{3}\left\|R_{3}(s)\right\|
$$

and so

$$
\mathbb{E} \sup _{s \in[0, t]}\left\|\theta_{1}(s)\right\|^{p} \leq K \epsilon^{3 p} \int_{0}^{t} \mathbb{E} \sup _{s \in[0, u]}\left\|R_{3}(s)\right\| \mathrm{d} u .
$$

In the same way, $\left\|\sigma_{s}\left(h_{s}\right)-\sigma_{s}\left(h_{s}^{(0)}+\epsilon h_{s}^{(1)}\right)\right\| \leq L_{\sigma} \epsilon^{2}\left\|h_{s}^{(2)}+\epsilon R_{3}(s)\right\|$. Recall that $\left(\int_{0}^{t} g(s) \mathrm{d} s\right)^{p} \leq$ $t^{p-1} \int_{0}^{t} g(s)^{p} \mathrm{~d} s$ by Jensen's inequality. Now, $\theta_{3}(s)$ is a continuous martingale, and hence, the Burkholder-Davis-Gundy inequality (see Theorem 3.28 in [14]) implies that for some constant $C_{p}>0$ depending only on $p$,

$$
\begin{aligned}
\mathbb{E} \sup _{s \in[0, t]}\left\|\theta_{3}(s)\right\|^{p} & \leq \epsilon^{p} C_{p}\left(\int_{0}^{t} \mathbb{E}\left\|\sigma_{s}\left(h_{s}\right)-\sigma_{s}\left(h_{s}^{(0)}+\epsilon h_{s}^{(1)}\right)\right\|^{2} \mathrm{~d} s\right)^{p / 2} \\
& \leq C_{p} L_{\sigma}^{p} \epsilon^{3 p}\left(\int_{0}^{t} \mathbb{E}\left\|h_{s}^{(2)}+\epsilon R_{3}(s)\right\|^{2} \mathrm{~d} s\right)^{p / 2} \\
& \leq C_{p} L_{\sigma}^{p} \epsilon^{3 p} T^{p / 2-1} \int_{0}^{t} \mathbb{E}\left\|h_{s}^{(2)}+\epsilon R_{3}(s)\right\|^{p} \mathrm{~d} s \\
& \leq C_{p} L_{\sigma}^{p} \epsilon^{3 p} 2^{p-1} T^{p / 2-1}\left(\int_{0}^{t} \mathbb{E}\left\|h_{s}^{(2)}\right\|^{p} \mathrm{~d} s+\epsilon^{p} \int_{0}^{t} \mathbb{E}\left\|R_{3}(s)\right\|^{p} \mathrm{~d} s\right) .
\end{aligned}
$$

From Lemma2, it follows that

$$
\mathbb{E} \sup _{s \in[0, t]}\left\|\theta_{3}(s)\right\|^{p} \leq K \epsilon^{3 p}\left(1+\epsilon^{p} \int_{0}^{t} \mathbb{E} \sup _{s \in[0, u]}\left\|R_{3}(s)\right\|^{p} \mathrm{~d} u\right)
$$

Treating the $\theta_{2}$ term next, for each $s \in[0, t]$, by Taylor's theorem, there exists some $\epsilon_{s} \in(0, \epsilon)$ such that

$$
f_{s}\left(h_{s}^{\epsilon}\right)-f_{s}\left(h_{0}\right)-\epsilon \frac{\partial f_{s}}{\partial h}\left(h_{0}\right) h_{1}=\epsilon^{2} \frac{\partial f_{s}}{\partial h}\left(h_{s}^{\epsilon_{s}}\right) h_{2}+\epsilon^{2} \Phi_{s}^{(1)}\left(h_{s}^{\epsilon_{s}}, h_{1}\right)
$$

Therefore, by Lipschitz continuity of the derivatives of $f$,

$$
\begin{aligned}
\theta_{2}(t) & =\epsilon^{2} \int_{0}^{t}\left(\frac{\partial f_{s}}{\partial h}\left(h_{s}^{\epsilon_{s}}\right) h_{2}+\Phi_{s}^{(1)}\left(h_{s}^{\epsilon_{s}}, h_{1}\right)-\frac{\partial f_{s}}{\partial h}\left(h_{0}\right) h_{2}-\Phi_{s}^{(1)}\left(h_{0}, h_{1}\right)\right) \mathrm{d} s \\
& \leq K \epsilon^{2} \int_{0}^{t}\left\|h_{s}^{\epsilon_{s}}-h_{0}\right\| \mathrm{d} s \\
& \leq K \epsilon^{3} \int_{0}^{t}\left\|h_{s}^{(1)}\right\|+\epsilon\left\|h_{s}^{(2)}\right\| \mathrm{d} s
\end{aligned}
$$

From Lemma 2, it follows that

$$
\mathbb{E} \sup _{s \in[0, T]}\left\|\theta_{2}(s)\right\|^{p} \leq K \epsilon^{3 p}
$$

Similarly, by Taylor's theorem, there exists $\epsilon_{s} \in(0, \epsilon)$ such that

$$
\sigma_{s}\left(h_{s}^{(0)}+\epsilon h_{s}^{(1)}\right)-\sigma_{s}\left(h_{s}^{(0)}\right)=\epsilon \Phi_{s}^{(2)}\left(h_{s}^{(0)}+\epsilon_{s} h_{s}^{(1)}, h_{s}^{(1)}\right),
$$

and so for $p \geq 2$, by the Burkholder-Davis-Gundy inequality and Lipschitz continuity of the derivatives of $\sigma$,

$$
\begin{aligned}
\mathbb{E} \sup _{s \in[0, t]}\left\|\theta_{4}(s)\right\|^{p} & \leq C_{p} \epsilon^{2 p}\left(\int_{0}^{t} \mathbb{E}\left\|\Phi_{s}^{(2)}\left(h_{s}^{(0)}+\epsilon_{s} h_{s}^{(1)}, h_{s}^{(1)}\right)-\Phi_{s}^{(2)}\left(h_{s}^{(0)}, h_{s}^{(1)}\right)\right\|^{2} \mathrm{~d} s\right)^{p / 2} \\
& \leq K C_{p} \epsilon^{3 p}\left(\int_{0}^{t} \mathbb{E}\left\|h_{s}^{(1)}\right\|^{2} \mathrm{~d} s\right)^{p / 2} \\
& \leq K C_{p} T^{p / 2} \epsilon^{3 p} \mathbb{E} \sup _{s \in[0, T]}\left\|h_{s}^{(1)}\right\|^{p} \leq K \epsilon^{3 p} .
\end{aligned}
$$

Combining estimates for $\theta_{1}, \theta_{2}, \theta_{3}, \theta_{4}$,

$$
\begin{aligned}
& \mathbb{E} \sup _{s \in[0, t]}\left\|R_{3}(s)\right\|^{p} \\
& =4^{p-1} \epsilon^{-3 p}\left(\mathbb{E} \sup _{s \in[0, t]}\left\|\theta_{1}(s)\right\|^{p}+\mathbb{E} \sup _{s \in[0, t]}\left\|\theta_{2}(s)\right\|^{p}+\mathbb{E} \sup _{s \in[0, t]}\left\|\theta_{3}(s)\right\|^{p}+\mathbb{E} \sup _{s \in[0, T]}\left\|\theta_{4}(s)\right\|^{p}\right) \\
& \leq K\left(1+\int_{0}^{t} \mathbb{E} \sup _{s \in[0, u]}\left\|R_{3}(s)\right\|^{p} \mathrm{~d} u\right),
\end{aligned}
$$

and so by Gronwall's inequality, $\mathbb{E} \sup _{s \in[0, t]}\left\|R_{3}(s)\right\|^{p} \leq K e^{K t}$. Since $K$ is independent of $t \leq T$ and $\epsilon \leq \epsilon_{0}$, it follows that

$$
\sup _{\epsilon \in\left(0, \epsilon_{0}\right)} \mathbb{E} \sup _{s \in[0, t]}\left\|R_{3}(s)\right\|^{p} \leq K e^{K t}<+\infty
$$

and the result follows.

We remark that perturbative techniques such as the one used to obtain Theorem 3 are standard in the theory of stochastic flows.
Theorem 7 now follows in a straightforward fashion from Lemma 3 by taking a second-order Taylor expansion of $\ell\left(h_{t}^{(0)}+\epsilon h_{t}^{(1)}+\epsilon^{2} h_{t}^{(2)}+\mathcal{O}\left(\epsilon^{3}\right)\right)$ about $\epsilon=0$.
We are now in a position to prove Theorem 6 using Theorem 7 .
Proof of Theorem 6. From Theorem 7, we have, upon taking expectation:
$\mathbb{E} \ell\left(h_{t}\right)=\ell\left(h_{t}^{(0)}\right)+\epsilon\left(\nabla_{h^{(0)}} \ell\right)^{T} \mathbb{E} h_{t}^{(1)}+\epsilon^{2}\left(\left(\nabla_{h^{(0)}} \ell\right)^{T} \mathbb{E} h_{t}^{(2)}+\frac{1}{2} \mathbb{E}\left(h_{t}^{(1)}\right)^{T}\left(H_{h^{(0)}} \ell\right) h_{t}^{(1)}\right)+\mathcal{O}\left(\epsilon^{3}\right)$,
for $t \in[0, T]$, as $\epsilon \rightarrow 0$, where $H_{h^{(0)}}$ denotes Hessian operator and the $h_{t}^{(i)}$ satisfy Eq. (60)-(62).
Since $\nabla f_{t}$ and its derivative are bounded and are thus Lipschitz continuous, by Picard's theorem the IVP has a unique solution. Moreover, it follows from our assumptions that the solution to the IVP is square-integrable (i.e., $\int_{0}^{t}\|\Psi(t, s)\|_{F}^{2} d s<\infty$ for any $t \in[0, T]$ ). Therefore, the solution $h_{t}^{(1)}$ to Eq. (61) can be uniquely represented as the following Itô integral:

$$
\begin{equation*}
h_{t}^{(1)}=\int_{0}^{t} \Psi(t, s) \sigma\left(h_{s}^{(0)}, s\right) d B_{s} \tag{69}
\end{equation*}
$$

where $\Psi(t, s)$ is the (deterministic) fundamental matrix solving the IVP (51). We have $\mathbb{E} h_{t}^{(1)}=0$ and

$$
\begin{equation*}
\mathbb{E}\left\|h_{t}^{(1)}\right\|^{2}=\int_{0}^{t}\left\|\Psi(t, s) \sigma\left(h_{s}^{(0)}, s\right)\right\|_{F}^{2} d s<\infty \tag{70}
\end{equation*}
$$

Similar argument together with Assumption D shows that the solution $h_{t}^{(2)}$ to Eq. (62) admits the following unique integral representation, with the $i$ th component:

$$
\begin{equation*}
h_{t}^{(2) i}=\frac{1}{2} \int_{0}^{t} \Psi^{i j}(t, s)\left[h_{s}^{(1)}\right]^{l} \frac{\partial^{2} b^{j}}{\partial\left[h_{s}^{(0)}\right]^{l} \partial\left[h_{s}^{(0)}\right]^{k}}\left[h_{s}^{(1)}\right]^{k} d s+\int_{0}^{t} \Psi^{i j}(t, s) \frac{\partial \sigma^{j k}}{\partial\left[h_{s}^{(0)}\right]^{l}}\left[h_{s}^{(1)}\right]^{l} d B_{s}^{k} \tag{71}
\end{equation*}
$$

where the last integral above is a uniquely defined Itô integral.
Plugging Eq. 69) into the above expression and then taking expectation, we have:

$$
\begin{align*}
\mathbb{E} h_{t}^{(2) i}= & \frac{1}{2} \mathbb{E} \int_{0}^{t} d s \Psi^{i j}(t, s) \frac{\partial^{2} b^{j}}{\partial\left[h_{s}^{(0)}\right]^{l} \partial\left[h_{s}^{(0)}\right]^{k}} \int_{0}^{s} d B_{u_{1}}^{l_{2}} \int_{0}^{s} d B_{u_{2}}^{k_{2}} \Psi^{l l_{1}}\left(s, u_{1}\right) \sigma^{l_{1} l_{2}} \sigma^{k_{1} k_{2}} \Psi^{k k_{1}}\left(s, u_{2}\right) \\
& +\frac{1}{2} \mathbb{E} \int_{0}^{t} d B_{s}^{k} \Psi^{i j}(t, s) \frac{\partial \sigma^{j k}}{\partial\left[h_{s}^{(0)}\right]^{l}} \int_{0}^{t} d B_{u}^{l_{2}} \Psi^{l l_{1}}(s, u) \sigma^{l_{1} l_{2}}\left(h_{u}^{(0)}, u\right) \tag{72}
\end{align*}
$$

where we have performed change of variable to arrive at the last double integral above.
Using the semigroup property of $\Psi$, we have $\Psi(t, s)=\Psi(t, 0) \Psi^{-1}(s, 0)$ for any $s \leq t$ (and so $\Psi^{i j}(t, s)=\Psi^{i j_{1}}(t, 0)\left(\Psi^{-1}\right)^{j_{1} j}(s, 0)$ and $\Psi^{l l_{1}}(s, u)=\Psi^{l l_{2}}(s, 0)\left(\Psi^{-1}\right)^{l_{2} l_{1}}(u, 0)$ etc. $)$. Using this property in (72) and then evaluating the resulting expression using properties of moments of stochastic integrals (applying Eq. (5.7) and Proposition 4.16 in [10] - note that Itô isometry follows from Eq. (5.7) there), we obtain $\left(\nabla_{h^{(0)}} \ell\right)^{T} \mathbb{E} h_{t}^{(2)}=Q\left(h^{(0)}\right)$, where $Q$ satisfies Eq. 53).

Similarly, plugging Eq. 69) into $\mathbb{E} h_{t}^{(1) T}\left(H_{h^{(0)}} \ell\right) h_{t}^{(1)}$, and then proceeding as above and applying the cyclic property of trace, give $\frac{1}{2} \mathbb{E}\left(h_{t}^{(1)}\right)^{T}\left(H_{h^{(0)}} \ell\right) h_{t}^{(1)}=R\left(h^{(0)}\right)$, where $R$ satisfies Eq. 54. The proof is done.

## D. 2 Discrete-Time Setting: Proof of Theorem 1 in the Main Paper

The goal in this subsection is to prove Theorem 1 in the main paper, the discrete-time analogue of Theorem6 We recall the theorem in the following.
Theorem 8 (Explicit regularization induced by noise injection for discrete-time NRNNs - Theorem 1 in the main paper). Under Assumption A in the main paper,

$$
\begin{equation*}
\mathbb{E} \ell\left(h_{M}^{\delta}\right)=\ell\left(\bar{h}_{M}^{\delta}\right)+\frac{\epsilon^{2}}{2}\left[\hat{Q}\left(\bar{h}^{\delta}\right)+\hat{R}\left(\bar{h}^{\delta}\right)\right]+\mathcal{O}\left(\epsilon^{3}\right) \tag{73}
\end{equation*}
$$

as $\epsilon \rightarrow 0$, where the terms $\hat{Q}$ and $\hat{R}$ are given by

$$
\begin{align*}
& \hat{Q}\left(\bar{h}^{\delta}\right)=\left(\nabla l\left(\bar{h}_{M}^{\delta}\right)\right)^{T} \sum_{k=1}^{M} \delta_{k-1} \hat{\Phi}_{M-1, k} \sum_{m=1}^{M-1} \delta_{m-1} \boldsymbol{v}_{m},  \tag{74}\\
& \hat{R}\left(\bar{h}^{\delta}\right)=\sum_{m=1}^{M} \delta_{m-1} \operatorname{tr}\left(\sigma_{m-1}^{T} \hat{\Phi}_{M-1, m}^{T} H_{\bar{h}^{\delta}} l \hat{\Phi}_{M-1, m} \sigma_{m-1}\right), \tag{75}
\end{align*}
$$

with $\boldsymbol{v}_{m}$ a vector with the pth component $\left(p=1, \ldots, d_{h}\right)$ :

$$
\left[v_{m}\right]^{p}=\operatorname{tr}\left(\sigma_{m-1}^{T} \hat{\Phi}_{M-2, m}^{T} H_{\bar{h}^{\delta}}\left[f_{M}\right]^{p} \hat{\Phi}_{M-2, m} \sigma_{m-1}\right)
$$

Moreover,

$$
\begin{equation*}
\left|\hat{Q}\left(\bar{h}^{\delta}\right)\right| \leq C_{Q} \Delta^{2}, \quad\left|\hat{R}\left(\bar{h}^{\delta}\right)\right| \leq C_{R} \Delta \tag{76}
\end{equation*}
$$

for $C_{Q}, C_{R}>0$ independent of $\Delta$.
To prove Theorem 8 , the key idea is to first obtain a discretized version of the loss function in Theorem 7 by either discretizing the results in Theorem 7 or by proving directly from the discretized equations (26). It then remains to compute the expectation of this loss as functional of the discrete-time process. The first part is straightforward while the second part involves some tedious recursive computations.
Let $0:=t_{0}<t_{1}<\cdots<t_{M}:=T$ be a partition of the interval $[0, T]$ and let $\delta_{m}=t_{m+1}-t_{m}$ for each $m=0,1, \ldots, M-1$. For small parameter $\epsilon>0$, the E-M scheme is given by:

$$
\begin{equation*}
h_{m+1}^{\delta}=h_{m}^{\delta}+f\left(h_{m}^{\delta}, \hat{x}_{m}\right) \delta_{m}+\epsilon \sigma\left(h_{m}^{\delta}, \hat{x}_{m}\right) \sqrt{\delta_{m}} \xi_{m} \tag{77}
\end{equation*}
$$

where $\left(\hat{x}_{m}\right)_{m=0, \ldots, M-1}$ is a given sequential data, each $\xi_{m} \sim \mathcal{N}(0, I)$ is an independent $r$ dimensional standard normal random vector, and $h_{0}^{\delta}=h_{0}$.
Consider the following hierarchy of recursive equations. For the sake of notation cleanliness, we replace the superscript $\delta$ by hat when denoting the $\delta$-dependent approximating solutions in the following.
For $m=0,1, \ldots, M-1$ :

$$
\begin{align*}
& \hat{h}_{m+1}^{(0)}=\hat{h}_{m}^{(0)}+\delta_{m} f\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right), \quad \hat{h}_{0}^{(0)}=h_{0}  \tag{78}\\
& \hat{h}_{m+1}^{(1)}=\hat{J}_{m} \hat{h}_{m}^{(1)}+\sqrt{\delta_{m}} \sigma\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right) \xi_{m}, \hat{h}_{0}^{(1)}=0  \tag{79}\\
& \hat{h}_{m+1}^{(2)}=\hat{J}_{m} \hat{h}_{m}^{(2)}+\sqrt{\delta_{m}}+\delta_{m} \Psi_{1}\left(\hat{h}_{m}^{(0)}, \hat{h}_{m}^{(1)}\right)+\delta_{m} \Psi_{2}\left(\hat{h}_{m}^{(0)}, \hat{h}_{m}^{(1)}\right) \xi_{m}, \quad \hat{h}_{0}^{(2)}=0, \tag{80}
\end{align*}
$$

where the

$$
\begin{equation*}
\hat{J}_{m}=I+\delta_{m} f^{\prime}\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right) \tag{81}
\end{equation*}
$$

are the state-to-state Jacobians and

$$
\begin{align*}
\Psi_{1}\left(h_{0}, h_{1}\right) & =\frac{1}{2} \sum_{i, j} \frac{\partial^{2} f_{m}}{\partial h^{i} \partial h^{j}}\left(h_{0}\right) h_{1}^{i} h_{1}^{j}  \tag{82}\\
\Psi_{2}\left(h_{0}, h_{1}\right) & =\sum_{i} \frac{\partial \sigma_{m}}{\partial h^{i}}\left(h_{0}\right) h_{1}^{i} . \tag{83}
\end{align*}
$$

Note that the above equations can also be obtained by E-M discretization of Eq. 60)-61).
The following theorem is a discrete-time analogue of Theorem 7 . Recall that the big $\mathcal{O}$ notation is to be understood in the almost sure sense.
Theorem 9. Under the same assumption as before, for a scalar-valued loss function $\ell \in \mathcal{C}^{2}\left(\mathbb{R}^{d_{h}}\right)$, for $m=0,1, \ldots, M-1$, we have

$$
\begin{align*}
\ell\left(\hat{h}_{m+1}\right)= & \ell\left(\hat{h}_{m}^{(0)}\right)+\epsilon \nabla \ell\left(\hat{h}_{m}^{(0)}\right) \cdot \hat{h}_{m}^{(1)}+\epsilon^{2}\left(\nabla \ell\left(\hat{h}_{m}^{(0)}\right) \cdot \hat{h}_{m}^{(2)}+\frac{1}{2}\left(\hat{h}_{m}^{(1)}\right)^{\top} \nabla^{2} \ell\left(\hat{h}_{m}^{(0)}\right)\left(\hat{h}_{m}^{(1)}\right)\right) \\
& +\mathcal{O}\left(\epsilon^{3}\right) \tag{84}
\end{align*}
$$

as $\epsilon \rightarrow 0$, where the $\hat{h}_{m}^{(i)}, i=0,1,2$, satisfy Eq. 78)-(80).

Proof. The proof is analogous to the one for continuous-time case, working with the discrete-time process (77) instead of continuous-time process.

We begin by recalling a remark from the main text.
Remark 1. Interestingly, Theorem 8 looks like discrete-time analogue of Theorem 6 for CT-RNN, except that, unlike the term $Q$ there, the term $\hat{Q}$ for the discrete-time case has no explicit dependence on the derivative (with respect to $h$ ) of the noise coefficient $\sigma$. Therefore, a direct discretization of the result in Theorem6 would not give us the correct explicit regularizer for discrete-time NRNNs. This remark highlights the difference between learning in the practical discrete-time setting versus learning in the idealized continuous-time setting with NRNNs. This also means that we need to work out an independently crafted proof for the discrete-time case.

The proof of Theorem 8 involves some tedious, albeit technically straightforward, computations. The key ingredients are the recursive relations (78) and the property of standard Gaussian random vectors that

$$
\begin{equation*}
\mathbb{E} \xi_{p}^{l} \xi_{q}^{j}=e_{p q} e_{l j} \tag{85}
\end{equation*}
$$

where the $e_{p q}$ denote the Kronecker delta.
To organize our proof, we begin by introducing some notation and proving a lemma.

Notation. For $m=1, \ldots, M-1$, let us denote $f_{m}^{\prime}:=f^{\prime}\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right), \sigma_{m}:=\sigma\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right)$,

$$
\begin{align*}
H_{l j} f_{m}^{i} & :=\frac{\partial^{2} f^{i}\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right)}{\partial\left[\left[_{m}^{(0)}\right]^{l} \partial\left[\hat{h}_{m}^{(0)}\right]^{j}\right.}  \tag{86}\\
D_{l} \sigma_{m}^{i j} & :=\frac{\partial \sigma^{i j}\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right)}{\partial\left[\hat{h}_{m}^{(0)}\right]^{l}} \tag{87}
\end{align*}
$$

and

$$
\begin{equation*}
\hat{\Phi}_{m, k}:=J_{m} J_{m-1} \cdots J_{k}, \quad \hat{\Phi}_{k, k+1}=I \tag{88}
\end{equation*}
$$

for $k=1, \ldots, m$. For computational convenience, we are using Einstein's summation notation for repeated indices in the following.
Lemma 4. For $m=0,1, \ldots, M, \mathbb{E} \hat{h}_{m}^{(1)}=0$ and

$$
\begin{equation*}
\mathbb{E}\left[\hat{h}_{m}^{(1)}\right]^{l}\left[\hat{h}_{m}^{(1)}\right]^{j}=\delta_{m-1} \sigma_{m-1}^{l l_{1}} \sigma_{m-1}^{j l_{1}}+\sum_{k=1}^{m-1} \delta_{k-1} \hat{\Phi}_{m-1, k}^{l l_{2}} \hat{\Phi}_{m-1, k}^{j j_{2}} \sigma_{k-1}^{l_{2} l_{3}} \sigma_{k-1}^{j_{2} l_{3}} . \tag{89}
\end{equation*}
$$

Proof. From Eq. (79), we have $\hat{h}_{0}^{(1)}=0, \hat{h}_{1}^{(1)}=\sqrt{\delta_{0}} \sigma_{t_{0}} \xi_{0}$ and, upon iterating, for $m=1, \ldots, M-$ 1,

$$
\begin{equation*}
\hat{h}_{m+1}^{(1)}=\sqrt{\delta_{m}} \sigma_{m} \xi_{m}+\sum_{k=1}^{m} \sqrt{\delta_{k-1}} \hat{\Phi}_{m, k} \sigma_{k-1} \xi_{k-1} \tag{90}
\end{equation*}
$$

The first equality in the lemma follows from taking expectation of Eq. 90) and using the fact that the $\xi_{k}$ are (mean zero) standard Gaussian random variables. The second equality in the lemma follows from taking expectation of a product of components of the $\hat{h}_{m+1}^{(1)}$ in Eq. 90) and applying the property 85].

Proof of Theorem 8. Iterating Eq. (80), we obtain $\hat{h}_{0}^{(2)}=0, \hat{h}_{1}^{(2)}=\delta_{0} \Psi_{1}\left(h_{0}, 0\right)+\sqrt{\delta_{0}} \Psi_{2}\left(h_{0}, 0\right) \xi_{0}$ and, for $m=1, \ldots, M-1$,

$$
\begin{align*}
\hat{h}_{m+1}^{(2)}= & \delta_{m} \Psi_{1}\left(\hat{h}_{m}^{(0)}, \hat{h}_{m}^{(1)}\right)+\sqrt{\delta_{m}} \Psi_{2}\left(\hat{h}_{m}^{(0)}, \hat{h}_{m}^{(1)}\right)+\sum_{k=1}^{m} \delta_{k-1} \hat{\Phi}_{m, k} \Psi_{1}\left(\hat{h}_{k-1}^{(0)}, \hat{h}_{k-1}^{(1)}\right) \\
& +\sum_{k=1}^{m} \sqrt{\delta_{k-1}} \hat{\Phi}_{m, k} \Psi_{1}\left(\hat{h}_{k-1}^{(0)}, \hat{h}_{k-1}^{(1)}\right) \xi_{k-1} \tag{91}
\end{align*}
$$

Substituting in the formulae (82)-(83) in the right hand side above and then using Eq. (90):

$$
\begin{align*}
{\left[\hat{h}_{m+1}^{(2)}\right]^{i}=} & \frac{\delta_{m}}{2}\left[\hat{h}_{m}^{(1)}\right]^{l} H_{l j} f_{m}^{i}\left[\hat{h}_{m}^{(1)}\right]^{j}+\sum_{k=1}^{m} \frac{\delta_{k-1}}{2} \hat{\Phi}_{m, k}^{i p}\left[\hat{h}_{k-1}^{(1)}\right]^{l} H_{l j} f_{k-1}^{p}\left[\hat{h}_{k-1}^{(1)}\right]^{j} \\
& +\sqrt{\delta_{m}} D_{l} \sigma_{m}^{i j}\left[\hat{h}_{m}^{(1)}\right]^{l} \xi_{m}^{j}+\sum_{k=1}^{m} \sqrt{\delta_{k-1}} \hat{\Phi}_{m, k}^{i q} D_{l} \sigma_{k-1}^{q r}\left[\hat{h}_{k-1}^{(1)}\right]^{l} \xi_{k-1}^{r}  \tag{92}\\
= & \frac{\delta_{m}}{2}\left[\hat{h}_{m}^{(1)}\right]^{l} H_{l j} f_{m}^{i}\left[\hat{h}_{m}^{(1)}\right]^{j}+\sum_{k=1}^{m} \frac{\delta_{k-1}}{2} \hat{\Phi}_{m, k}^{i p}\left[\hat{h}_{k-1}^{(1)}\right]^{l} H_{l j} f_{k-1}^{p}\left[\hat{h}_{k-1}^{(1)}\right]^{j} \\
& +\sqrt{\delta_{m}} D_{l} \sigma_{m}^{i j} \xi_{m}^{j}\left(\sqrt{\delta_{m-1}} \sigma_{m-1}^{l l_{1}} \xi_{m-1}^{l_{1}}+\sum_{k=1}^{m-1} \sqrt{\delta_{k-1}} \hat{\Phi}_{m-1, k}^{l l_{1}} \sigma_{k-1}^{l_{1} l_{2}} \xi_{k-1}^{l_{2}}\right) \\
& +\sqrt{\delta_{1}} \hat{\Phi}_{m, 2}^{i q} D_{l} \sigma_{1}^{q r} \xi_{1}^{r}\left(\sqrt{\delta_{0}} \sigma_{0}^{l l_{1}} \xi_{0}^{l_{1}}\right)  \tag{93}\\
& +\sum_{k=3}^{m} \sqrt{\delta_{k-1}} \hat{\Phi}_{m, k}^{i q} D_{l} \sigma_{k-1}^{q r} \xi_{k-1}^{r}\left(\sqrt{\delta_{k-2}} \sigma_{k-2}^{l p_{1}} \xi_{k-2}^{p_{1}}+\sum_{k^{\prime}=1}^{k-2} \sqrt{\delta_{k^{\prime}-1}} \hat{\Phi}_{k-2, k^{\prime}}^{l p_{1}} \sigma_{k^{\prime}-1}^{p_{1} p_{2}} \xi_{k^{\prime}-1}^{p_{2}}\right)
\end{align*}
$$

where we have made use of the fact that $\hat{h}_{0}^{(1)}=0$ and $\hat{h}_{1}^{(1)}=\sqrt{\delta_{0}} \sigma_{0} \xi_{0}$ in the last two lines above to rewrite the summation (so that the summation over $k$ in the last line above starts at $k=3$ ).

Therefore, using the above result, Lemma 4 and Eq. 85], we compute the expectation of $\left[\hat{h}_{m+1}^{(2)}\right]^{i}$ :

$$
\begin{equation*}
\mathbb{E}\left[\hat{h}_{m+1}^{(2)}\right]^{i}=\frac{1}{2} \sum_{k=1}^{m+1} \delta_{k-1} \hat{\Phi}_{m, k}^{i p} H_{l j} f_{m}^{p} \sum_{k=1}^{m} \delta_{k-1} \hat{\Phi}_{m-1, k}^{l l_{2}} \sigma_{k-1}^{l_{2} l_{3}} \sigma_{k-1}^{j_{2} l_{3}} \hat{\Phi}_{m-1, k}^{j j_{2}} \tag{94}
\end{equation*}
$$

Moreover, using Lemma 4 , we obtain, for $m=1,2, \ldots, M-1$,

$$
\begin{equation*}
\mathbb{E}\left[\hat{h}_{m+1}^{(1)}\right]^{l}\left[H_{\hat{h}^{(0)}} l\right]^{l j}\left[\hat{h}_{m+1}^{(1)}\right]^{j}=\sum_{k=1}^{m+1} \delta_{k-1} \sigma_{k-1}^{l_{2} l_{3}} \hat{\Phi}_{m, k}^{l l_{2}}\left[H_{\hat{h}^{(0)}} l\right]^{l j} \hat{\Phi}_{m, k}^{j j_{2}} \sigma_{k-1}^{j_{2} l_{3}} . \tag{95}
\end{equation*}
$$

The first statement of the theorem then follows from Theorem 9 and Eq. 94)-95) (with $m:=M-1$ ):

$$
\begin{align*}
& \hat{Q}\left(\bar{h}^{\delta}\right)=\partial_{i} l\left(\bar{h}_{M}^{\delta}\right) \sum_{k=1}^{M} \delta_{k-1} \hat{\Phi}_{M-1, k}^{i p} \sum_{m=1}^{M-1} \delta_{m-1} \partial_{l j}\left[f_{M}\right]^{p} \hat{\Phi}_{M-2, m}^{l l_{2}} \sigma_{m-1}^{l_{2} l_{3}} \sigma_{m-1}^{j_{2} l_{3}} \hat{\Phi}_{M-2, m}^{j j_{2}},  \tag{96}\\
& \hat{R}\left(\bar{h}^{\delta}\right)=\sum_{m=1}^{M} \delta_{m-1} \sigma_{m-1}^{l_{2} l_{3}} \hat{\Phi}_{M-1, m}^{l l_{2}}\left[H_{\bar{h}^{\delta}} l\right]^{l j} \hat{\Phi}_{M-1, m}^{j j_{2}} \sigma_{m-1}^{j_{2} l_{3}} . \tag{97}
\end{align*}
$$

The last statement of the theorem follows from taking straightforward bounds.
Remark 2. We remark that the computed $\hat{h}_{m}^{(2)}$ (a key step in the above proof), like that for $h_{t}^{(2)}$ in the continuous-time case, has explicit dependence on the noise coefficient. It is only upon taking the expectation (see Eq. 94 ) that the dependence on the noise coefficient vanishes (whereas $\mathbb{E} h_{t}^{(2)} \neq 0$ retains its dependence on the noise coefficient). This fully reconciles with Remark 1
Remark 3. Moreover, one can compute the variance of $l\left(\hat{h}_{M}\right)$ to be $\epsilon^{2}\left(\nabla l\left(\hat{h}_{M}^{(0)}\right)\right)^{T} C \nabla l\left(\hat{h}_{M}^{(0)}\right)+$ $\mathcal{O}\left(\epsilon^{3}\right)$, as $\epsilon \rightarrow 0$, where $C$ is a PSD matrix whose $(l, j)$-entry is given by Eq. 89) with $m:=M$. So we see that the spread of $l\left(\hat{h}_{M}\right)$ about its average is $\mathcal{O}\left(\epsilon^{2}\right)$ as $\epsilon \rightarrow 0$.

## E Bound on Classification Margin and a Generalization Bound for Deterministic RNNs: Proof of Theorem 2 in the Main Paper

We recall the setting considered in the main paper before providing proof to the results presented there.

Let $\mathcal{S}_{N}$ denote a set of training samples $s_{n}:=\left(\boldsymbol{x}_{n}, y_{n}\right)$ for $n=1, \ldots, N$, where each input sequence $\boldsymbol{x}_{n}=\left(x_{n, 0}, x_{n, 1}, \ldots, x_{n, M-1}\right) \in \mathcal{X} \subset \mathbb{R}^{d_{x} M}$ has a corresponding class label $y_{n} \in$ $\mathcal{Y}=\left\{1, \ldots, d_{y}\right\}$. Following the statistical learning framework, these samples are assumed to be independently drawn from an underlying probability distribution $\mu$ on the sample space $\mathcal{S}=\mathcal{X} \times \mathcal{Y}$. An RNN-based classifier $g^{\delta}(\boldsymbol{x})$ is constructed in the usual way by taking

$$
\begin{equation*}
g^{\delta}(\boldsymbol{x})=\operatorname{argmax}_{i=1, \ldots, d_{y}} p^{i}\left(V \bar{h}_{M}^{\delta}[\boldsymbol{x}]\right) \tag{98}
\end{equation*}
$$

where $p^{i}(x)=e^{x^{i}} / \sum_{j} e^{x^{j}}$ is the softmax function. Letting $\ell$ denoting the cross-entropy loss, such a classifier is trained from $\mathcal{S}_{N}$ by minimizing the empirical risk (training error)

$$
\mathcal{R}_{N}\left(g^{\delta}\right):=\frac{1}{N} \sum_{n=1}^{N} \ell\left(g^{\delta}\left(\boldsymbol{x}_{n}\right), y_{n}\right)
$$

as a proxy for the true (population) risk (testing error) $\mathcal{R}\left(g^{\delta}\right)=\mathbb{E}_{(\boldsymbol{x}, y) \sim \mu} \ell\left(g^{\delta}(\boldsymbol{x}), y\right)$ with $(\boldsymbol{x}, y) \in \mathcal{S}$. The measure used to quantify the prediction quality is the generalization error (or estimation error), which is the difference between the empirical risk of the classifier on the training set and the true risk:

$$
\begin{equation*}
G E\left(g^{\delta}\right):=\left|\mathcal{R}\left(g^{\delta}\right)-\mathcal{R}_{N}\left(g^{\delta}\right)\right| \tag{99}
\end{equation*}
$$

The classifier is a function of the output of the deterministic RNN, which is an Euler discretization of the ODE (1) in the main paper with step sizes $\delta=\left(\delta_{m}\right)$. In particular, for the Lipschitz RNN,

$$
\begin{equation*}
\hat{\Phi}_{m, k}=\hat{J}_{m} \hat{J}_{m-1} \cdots \hat{J}_{k} \tag{100}
\end{equation*}
$$

where $\hat{J}_{l}=I+\delta_{l}\left(A+D_{l} W\right)$, with $D_{l}^{i j}=a^{\prime}\left(\left[W \bar{h}_{l}^{\delta}+U \hat{x}_{l}+b\right]^{i}\right) e_{i j}$.
In the following, we let $\operatorname{conv}(\mathcal{X})$ denote the convex hull of $\mathcal{X}$. We denote $\hat{\boldsymbol{x}}_{0: m}:=\left(\hat{x}_{0}, \ldots, \hat{x}_{m}\right)$ so that $\hat{\boldsymbol{x}}=\hat{\boldsymbol{x}}_{0: M-1}$, and use the notation $f[\boldsymbol{x}]$ to indicate the dependence of the function $f$ on the vector $\boldsymbol{x}$. Moreover, we will need the following two definitions to characterize a training sample $s_{i}=\left(\boldsymbol{x}_{i}, y_{i}\right)$
Working in the above setting, we now recall and prove the second main result in the main paper, providing bounds for classification margin for the deterministic RNN classifiers $g^{\delta}$.
Theorem 10 (Classification margin bound for the deterministic RNN - Theorem 2 in the main paper). Suppose that Assumption A in the main paper holds. Assume that the $o\left(s_{i}\right)>0$ and

$$
\begin{equation*}
\gamma\left(s_{i}\right):=\frac{o\left(s_{i}\right)}{C \sum_{m=0}^{M-1} \delta_{m} \sup _{\hat{\boldsymbol{x}} \in \operatorname{conv}(\mathcal{X})}\left\|\hat{\Phi}_{M, m+1}[\hat{\boldsymbol{x}}]\right\|_{2}}>0 \tag{101}
\end{equation*}
$$

where

$$
C=\|V\|_{2}\left(\max _{m=0,1, \ldots, M-1}\left\|\frac{\partial f\left(\bar{h}_{m}^{\delta}, \hat{x}_{m}\right)}{\partial \hat{x}_{m}}\right\|_{2}\right)>0
$$

is a constant (in particular, $C=\|V\|_{2}\left(\max _{m=0, \ldots, M-1}\left\|D_{m} U\right\|_{2}\right)$ for Lipschitz $\left.R N N s\right)$, the $\hat{\Phi}_{m, k}$ are defined in 100) and the $\delta_{m}$ are the step sizes. Then, we have the following upper bound on the classification margin for the training sample $s_{i}$ :

$$
\begin{equation*}
\gamma^{d}\left(s_{i}\right) \geq \gamma\left(s_{i}\right) \tag{102}
\end{equation*}
$$

Moreover, under additional assumptions one can obtain the following generalization bound, which follows from Theorem 10
Theorem 11 (A generalization bound for the deterministic RNN). Under the same setting as Theorem 10 if we further assume that $\mathcal{X}$ is a (subset of) $k$-dimensional manifold with $k \leq d_{x} M, \gamma:=$ $\min _{s_{i} \in \mathcal{S}_{N}} \gamma\left(s_{i}\right)>0$, and $\ell\left(g^{\delta}(\boldsymbol{x}), y\right) \leq L_{g}$ for all $s \in \mathcal{S}$, then for any $\delta^{\prime}>0$, with probability at least $1-\delta^{\prime}$,

$$
\begin{equation*}
G E\left(g^{\delta}\right) \leq L_{g}\left(\frac{1}{\gamma^{k / 2}} \sqrt{\frac{d_{y} C_{M}^{k} 2^{k+1} \log 2}{N}}+\sqrt{\frac{2 \log \left(1 / \delta^{\prime}\right)}{N}}\right) \tag{103}
\end{equation*}
$$

where $C_{M}>0$ is a constant that measures complexity of $\mathcal{X}, N$ is the number of training examples and $d_{y}$ is the number of label classes.
Remark 4. Generalization bounds involving classification margins (for RNNs in particular) are a separate topic with a significant presence in the literature. We emphasize that the generalization bound above is one of the many bounds that one can derive for RNNs. There exist much tighter bounds (for various variants of RNNs under various assumptions and settings) which may be equally applicable and lead to the same claimed conclusion, but are much more difficult to state (see, for instance, Theorem E. 1 in [28]). There are also other types of generalization bounds that are not obtained in terms of classification margin in the literature. Although they are interesting in their own, our focus here is on bounds that can be expressed in terms of classification margin. Therefore, meaningful comparisons between these generalization bounds are not straightforward.

In order to prove Theorem 10 and Theorem 11 , we place ourselves in the algorithmic robustness framework of [30]. This framework provides bounds for the generalization error based on the robustness of a learning algorithm that learns a classifier $g$ by exploiting the structure of the training set $\mathcal{S}_{N}$. Robustness is, roughly speaking, the desirable property for a learning algorithm that if a testing sample is "similar" to a training sample, then the testing error is close to the training error (i.e., the algorithm is insensitive to small perturbations in the training data).

To ensure that our exposition is self-contained, we recall important definitions and results from [30, 27] to formalize the previous statement in the context of our deterministic RNNs in the following.

Definition 4. Let $\mathcal{S}_{N}$ be a training set and $\mathcal{S}$ the sample space. A learning algorithm is $\left(K, \epsilon\left(\mathcal{S}_{N}\right)\right)$ robust if $\mathcal{S}$ can be partitioned into $K$ disjoint sets denoted by $\mathcal{K}_{k}, k=1, \ldots, K$ :

$$
\begin{align*}
& \mathcal{K}_{k} \subset \mathcal{S}, k=1, \ldots, K  \tag{104}\\
& \mathcal{S}=\cup_{k=1}^{K} \mathcal{K}_{k}, \text { and } \mathcal{K}_{k} \cap \mathcal{K}_{k^{\prime}}=\emptyset, \forall k \neq k^{\prime} \tag{105}
\end{align*}
$$

such that for all $s_{i} \in \mathcal{S}_{N}$ and all $s \in \mathcal{S}$,

$$
\begin{equation*}
s_{i}=\left(\boldsymbol{x}_{i}, y_{i}\right) \in \mathcal{K}_{k} \wedge s=(\boldsymbol{x}, y) \in \mathcal{K}_{k} \Longrightarrow\left|\ell\left(g\left(\boldsymbol{x}_{i}\right), y_{i}\right)-\ell(g(\boldsymbol{x}), y)\right| \leq \epsilon\left(\mathcal{S}_{N}\right) \tag{106}
\end{equation*}
$$

The above definition says that a robust learning algorithm selects a classifier $g$ for which the losses of any $s$ and $s_{i}$ in the same partition $\mathcal{K}_{k}$ are close.
The following result from Theorem 1 in [30] will be critical to the proof of Theorem 10. It provides a generalization bound for robust algorithms.
Theorem 12. If a learning algorithm is $\left(K, \epsilon\left(\mathcal{S}_{N}\right)\right)$-robust and $\ell(g(\boldsymbol{x}), y) \leq M$ for all $s=(\boldsymbol{x}, y) \in$ $\mathcal{S}$, for some constant $M>0$, then for any $\delta>0$, with probability at least $1-\delta$,

$$
\begin{equation*}
G E(g) \leq \epsilon\left(\mathcal{S}_{N}\right)+M \sqrt{\frac{2 K \log (2)+2 \log (1 / \delta)}{m}} \tag{107}
\end{equation*}
$$

Note that the above generalization bound is data-dependent, in contrast to bounds obtained via approaches based on complexity or stability arguments that give bounds in terms of data agnostic measures such as the Rademacher complexity or the VC dimension, which are found not sufficient for explaining the good generalization properties of deep neural networks.

The number of partition $K$ in the above can be bounded in terms of the covering number of the sample space $\mathcal{S}$, which gives a way to measure the complexity of sets. We recall the definition of covering number in the following.
Definition 5 (Covering). Let $\mathcal{A}$ be a set. We say that $\mathcal{A}$ is $\rho$-covered by a set $\mathcal{A}^{\prime}$, with respect to the (pseudo-)metric $d$, if for all $a \in \mathcal{A}$, there exists $a^{\prime} \in \mathcal{A}^{\prime}$ with $d\left(a, a^{\prime}\right) \leq \rho$. We call the cardinality of the smallest $\mathcal{A}^{\prime}$ that $\rho$-covers $\mathcal{A}$ covering number, denoted by $\mathcal{N}(\mathcal{S} ; d, \rho)$.

The covering number is the smallest number of (pseudo-)metric balls of radius $\rho$ needed to cover $\mathcal{S}$ and we denote it by $\mathcal{N}(\mathcal{S} ; d, \rho)$, where $d$ denotes the (pseudo-)metric. The choice of metric $d$ determines how efficiently one may cover $\mathcal{X}$. For example, the Euclidean metric $d\left(x, x^{\prime}\right)=\left\|x-x^{\prime}\right\|_{2}$ for $x, x^{\prime} \in \mathcal{X}$. The covering number of many structured low-dimensional data models can be bounded in terms of their intrinsic properties. Since in our case the space $\mathcal{S}=\mathcal{X} \times \mathcal{Y}$, we write $\mathcal{N}(\mathcal{S} ; d, \rho) \leq d_{y} \cdot \mathcal{N}(\mathcal{X} ; d, \rho)$, where $d_{y}$ is the number of label classes. We take $d$ to be the Euclidean metric: $d\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right)=\left\|\boldsymbol{x}-\boldsymbol{x}^{\prime}\right\|_{2}$ for $\boldsymbol{x}, \boldsymbol{x}^{\prime} \in \mathcal{X}$, unless stated otherwise.
Lemma 5 (Example 27.1 from [26]). Assume that $\mathcal{X} \subset \mathbb{R}^{m}$ lies in a $k$-dimensional subspace of $\mathbb{R}^{m}$. Let $c=\max _{x \in \mathcal{X}}\|x\|$ and take $d$ to be the Euclidean metric. Then $\mathcal{N}(\mathcal{X} ; d, \rho) \leq(2 c \sqrt{k} / \rho)^{k}$.

In other words, a subset, $\mathcal{X}$, of a $k$-dimensional manifold has the covering number $\left(C_{M} / \rho\right)^{k}$, where $C_{M}>0$ is a constant. We remark that other complexity measures such as Rademacher complexity can be bounded based on the covering number (see [26] for details).

The class of robust learning algorithms that is of interest to us is the large margin classifiers. We define classification margin in the following.
Definition 6 (Classification margin). The classification margin of a training sample $s_{i}=\left(\boldsymbol{x}_{i}, y_{i}\right)$ measured by a metric $d$ is defined as the radius of the largest $d$-metric ball in $\mathcal{X}$ centered at $\boldsymbol{x}_{i}$ that is contained in the decision region associated with the class label $y_{i}$, i.e., it is:

$$
\begin{equation*}
\gamma^{d}\left(s_{i}\right)=\sup \left\{a: d\left(\boldsymbol{x}_{i}, \boldsymbol{x}\right) \leq a \Longrightarrow g(\boldsymbol{x})=y_{i} \forall \boldsymbol{x}\right\} \tag{108}
\end{equation*}
$$

Intuitively, a larger classification margin allows a classifier to associate a larger region centered on a point $\boldsymbol{x}_{i}$ in the input space to the same class. This makes the classifier less sensitive to input perturbations and a noisy perturbation of $\boldsymbol{x}_{i}$ is still likely to fall within this region, keeping the classifier prediction. In this sense, the classifier becomes more robust.
The following result follows from Example 9 in [30].

Proposition 2. If there exists a $\gamma>0$ such that $\gamma^{d}\left(s_{i}\right)>\gamma$ for all $s_{i} \in \mathcal{S}_{N}$, then the classifier $g$ is $\left(d_{y} \cdot \mathcal{N}(\mathcal{X} ; d, \gamma / 2), 0\right)$-robust.

In our case the networks are trained by a loss (cross-entropy) that promotes separation of different classes at the network output. The training aims at maximizing a certain notion of score of each training sample.
Definition 7 (Score). For a training sample $s_{i}=\left(\boldsymbol{x}_{i}, y_{i}\right)$, we define its score as

$$
\begin{equation*}
o\left(s_{i}\right)=\min _{j \neq y_{i}} \sqrt{2}\left(e_{y_{i}}-e_{j}\right)^{T} S^{\delta}\left[\boldsymbol{x}_{i}\right] \geq 0 \tag{109}
\end{equation*}
$$

where $e_{i} \in \mathbb{R}^{d_{y}}$ is the Kronecker delta vector with $e_{i}^{i}=1$ and $e_{i}^{j}=0$ for $i \neq j, S^{\delta}\left[\boldsymbol{x}_{i}\right]:=$ $p\left(V \bar{h}_{M}^{\delta}\left[\boldsymbol{x}_{i}\right]\right)$ with $\bar{h}_{M}^{\delta}\left[\boldsymbol{x}_{i}\right]$ denoting the hidden state of the RNN, driven by the input sequence $\boldsymbol{x}_{i}$, at terminal index $M$.

The RNN classifier $g^{\delta}$ is defined as

$$
\begin{equation*}
g^{\delta}(\boldsymbol{x})=\arg \max _{i \in\left\{1, \ldots, d_{y}\right\}} S^{i}[\boldsymbol{x}] \tag{110}
\end{equation*}
$$

and the decision boundary between class $i$ and class $j$ in the output space is given by the hyperplane $\left\{z=p\left(V \bar{h}_{M}^{\delta}\right): z^{i}=z^{j}\right\}$. A positive score implies that at the network output, classes are separated by a margin that corresponds to the score. However, a large score may not imply a large classification margin - recall that the classification margin is a function of the decision boundary in the input space, whereas the training algorithm aims at optimizing the decision boundary at the network output in the output space.
We need the following lemma relating a pair of vectors in the input space and the output space.
Lemma 6. For any $\boldsymbol{x}, \boldsymbol{x}^{\prime} \in \mathcal{X} \subset \mathbb{R}^{d_{x} M}$, and a given $R N N$ output functional $\mathcal{F}[\cdot]$,

$$
\begin{equation*}
\left\|\mathcal{F}[\boldsymbol{x}]-\mathcal{F}\left[\boldsymbol{x}^{\prime}\right]\right\|_{2} \leq \sup _{\overline{\boldsymbol{x}} \in \operatorname{conv}(\mathcal{X})}\|J[\overline{\boldsymbol{x}}]\|_{2} \cdot\left\|\boldsymbol{x}-\boldsymbol{x}^{\prime}\right\|_{2}, \tag{111}
\end{equation*}
$$

where $\boldsymbol{J}[\boldsymbol{x}]=d \mathcal{F}[\boldsymbol{x}] / d \boldsymbol{x}$ is the input-output Jacobian of the RNN output functional.
Proof. Let $t \in[0,1]$ and define the function $F(t)=\mathcal{F}\left[\boldsymbol{x}+t\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right)\right]$. Note that

$$
\begin{equation*}
\frac{d F(t)}{d t}=\boldsymbol{J}\left[\boldsymbol{x}+t\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right)\right]\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right) . \tag{112}
\end{equation*}
$$

Therefore,

$$
\begin{equation*}
\mathcal{F}\left[\boldsymbol{x}^{\prime}\right]-\mathcal{F}[\boldsymbol{x}]=F(1)-F(0)=\int_{0}^{1} \frac{d F(t)}{d t} d t=\left(\int_{0}^{1} \boldsymbol{J}\left[\boldsymbol{x}+t\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right)\right] d t\right)\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right), \tag{113}
\end{equation*}
$$

where we have used the fundamental theorem of calculus.
Now,

$$
\begin{align*}
\left\|\mathcal{F}[\boldsymbol{x}]-\mathcal{F}\left[\boldsymbol{x}^{\prime}\right]\right\|_{2} & \leq\left\|\int_{0}^{1} \boldsymbol{J}\left[\boldsymbol{x}+t\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right)\right] d t\right\|_{2} \cdot\left\|\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right)\right\|_{2}  \tag{114}\\
& \leq \sup _{\boldsymbol{x}, \boldsymbol{x}^{\prime} \in \mathcal{X}, t \in[0,1]}\left\|\boldsymbol{J}\left[\boldsymbol{x}+t\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right)\right]\right\|_{2} \cdot\left\|\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right)\right\|_{2}  \tag{115}\\
& \leq \sup _{\overline{\boldsymbol{x}} \in \operatorname{conv}(\mathcal{X})}\|J[\overline{\boldsymbol{x}}]\|_{2} \cdot\left\|\boldsymbol{x}-\boldsymbol{x}^{\prime}\right\|_{2} \tag{116}
\end{align*}
$$

where we have used the fact that $\boldsymbol{x}+t\left(\boldsymbol{x}^{\prime}-\boldsymbol{x}\right) \in \operatorname{conv}(\mathcal{X})$ for all $t \in[0,1]$ to arrive at the last line. The proof is done.

The classification margin depends on the score and the network's expansion and contraction of distances around the training points. These can be quantified by studying the network's input-output Jacobian matrix. The following proposition provides classification margin bounds in terms of the score and input-output Jacobian associated to the RNN classifier.

Proposition 3. Assume that a RNN classifier $g^{\delta}(\boldsymbol{x})$, defined in (110), classifies a training sample $\boldsymbol{x}_{i}$ with the score $o\left(s_{i}\right)>0$. Then we have the following lower bound for the classification margin:

$$
\begin{equation*}
\gamma^{d}\left(s_{i}\right) \geq \frac{o\left(s_{i}\right)}{\sup _{\boldsymbol{x} \in \operatorname{conv}(\mathcal{X})}\|J[\boldsymbol{x}]\|_{2}} \tag{117}
\end{equation*}
$$

where $\operatorname{conv}(\mathcal{X})$ denotes the convex hull of $\mathcal{X}$ and $J[\boldsymbol{x}]=d \mathcal{F}[\boldsymbol{x}] / d \boldsymbol{x}$, with $\mathcal{F}[\boldsymbol{x}]=p\left(V \bar{h}_{M}^{\delta}[\boldsymbol{x}]\right)$, is the input-output Jacobian associated to the RNN.

Proof. The proof is essentially identical to that of Theorem 4 in [27]. We provide the full detail here for completeness.
Denote $o\left(s_{i}\right)=o\left(\boldsymbol{x}^{(i)}, y^{(i)}\right)$, where $\boldsymbol{x}^{(i)}:=\left(x_{0}^{(i)}, \cdots, x_{M-1}^{(i)}\right) \in \mathcal{X} \subset \mathbb{R}^{d_{x} M}$, and $v_{i j}=\sqrt{2}\left(e_{i}-\right.$ $e_{j}$ ), where $e_{i} \in \mathbb{R}^{d_{y}}$ denotes the Kronecker delta vector.
The classification margin of the training sample $s_{i}$ is:

$$
\begin{align*}
\gamma^{d}\left(s_{i}\right) & =\sup \left\{a:\left\|\boldsymbol{x}^{(i)}-\boldsymbol{x}\right\|_{2} \leq a \Longrightarrow g^{\delta}(\boldsymbol{x})=y^{(i)} \forall \boldsymbol{x}\right\}  \tag{118}\\
& =\sup \left\{a:\left\|\boldsymbol{x}^{(i)}-\boldsymbol{x}\right\|_{2} \leq a \Longrightarrow o\left(\boldsymbol{x}, y^{(i)}\right)>0 \forall \boldsymbol{x}\right\} . \tag{119}
\end{align*}
$$

By Definition 7, $o\left(\boldsymbol{x}, y^{(i)}\right)>0$ if and only if $\min _{j \neq y^{(i)}} v_{y^{(i)} j}^{T} \mathcal{F}[\boldsymbol{x}]>0$.
On the other hand,

$$
\begin{align*}
\min _{j \neq y^{(i)}} v_{y^{(i)} j}^{T} \mathcal{F}[\boldsymbol{x}] & =\min _{j \neq y^{(i)}}\left(v_{y^{(i)} j}^{T} \mathcal{F}\left[\boldsymbol{x}^{(i)}\right]+v_{y^{(i)} j}^{T}\left(\mathcal{F}[\boldsymbol{x}]-\mathcal{F}\left[\boldsymbol{x}^{(i)}\right]\right)\right)  \tag{120}\\
& \geq \min _{j \neq y^{(i)}} v_{y^{(i)} j}^{T} \mathcal{F}\left[\boldsymbol{x}^{(i)}\right]+\min _{j \neq y^{(i)}} v_{y^{(i) j}}^{T}\left(\mathcal{F}[\boldsymbol{x}]-\mathcal{F}\left[\boldsymbol{x}^{(i)}\right]\right)  \tag{121}\\
& =o\left(\boldsymbol{x}^{(i)}, y^{(i)}\right)+\min _{j \neq y^{(i)}} v_{y^{(i)} j}^{T}\left(\mathcal{F}[\boldsymbol{x}]-\mathcal{F}\left[\boldsymbol{x}^{(i)}\right]\right) \tag{122}
\end{align*}
$$

Therefore, $o\left(\boldsymbol{x}^{(i)}, y^{(i)}\right)+\min _{j \neq y^{(i)}} v_{y^{(i)} j}^{T}\left(\mathcal{F}[\boldsymbol{x}]-\mathcal{F}\left[\boldsymbol{x}^{(i)}\right]\right)>0$ implies that $o\left(\boldsymbol{x}, y^{(i)}\right)>0$ and so

$$
\begin{align*}
\gamma^{d}\left(s_{i}\right) & \geq \sup \left\{a:\left\|\boldsymbol{x}^{(i)}-\boldsymbol{x}\right\|_{2} \leq a \Longrightarrow o\left(\boldsymbol{x}^{(i)}, y^{(i)}\right)+\min _{j \neq y^{(i)}} v_{y^{(i)} j}^{T}\left(\mathcal{F}[\boldsymbol{x}]-\mathcal{F}\left[\boldsymbol{x}^{(i)}\right]\right)>0 \quad \forall \boldsymbol{x}\right\}  \tag{123}\\
& =\sup \left\{a:\left\|\boldsymbol{x}^{(i)}-\boldsymbol{x}\right\|_{2} \leq a \Longrightarrow o\left(\boldsymbol{x}^{(i)}, y^{(i)}\right)-\max _{j \neq y^{(i)}} v_{y^{(i)} j}^{T}\left(\mathcal{F}\left[\boldsymbol{x}^{(i)}\right]-\mathcal{F}[\boldsymbol{x}]\right)>0 \quad \forall \boldsymbol{x}\right\}  \tag{124}\\
& =\sup \left\{a:\left\|\boldsymbol{x}^{(i)}-\boldsymbol{x}\right\|_{2} \leq a \Longrightarrow o\left(\boldsymbol{x}^{(i)}, y^{(i)}\right)>\max _{j \neq \boldsymbol{y}^{(i)}} v_{y^{(i)} j}^{T}\left(\mathcal{F}\left[\boldsymbol{x}^{(i)}\right]-\mathcal{F}[\boldsymbol{x}]\right) \forall \boldsymbol{x}\right\} \tag{125}
\end{align*}
$$

Now, using the fact that $\left\|v_{y^{(i)} j}\right\|_{2}=1$ and Lemma 6, we have:

$$
\begin{equation*}
\max _{j \neq y^{(i)}} v_{y^{(i)} j}^{T}\left(\mathcal{F}\left[\boldsymbol{x}^{(i)}\right]-\mathcal{F}[\boldsymbol{x}]\right) \leq \sup _{\overline{\boldsymbol{x}} \in \operatorname{conv}(\mathcal{X})}\|J[\overline{\boldsymbol{x}}]\|_{2} \cdot\left\|\boldsymbol{x}^{(i)}-\boldsymbol{x}\right\|_{2} \tag{126}
\end{equation*}
$$

Using this inequality gives:

$$
\begin{align*}
\gamma^{d}\left(s_{i}\right) & \geq \sup \left\{a:\left\|\boldsymbol{x}^{(i)}-\boldsymbol{x}\right\|_{2} \leq a \Longrightarrow o\left(\boldsymbol{x}^{(i)}, y^{(i)}\right)>\sup _{\overline{\boldsymbol{x}} \in \operatorname{conv}(\mathcal{X})}\|J[\overline{\boldsymbol{x}}]\|_{2} \cdot\left\|\boldsymbol{x}^{(i)}-\boldsymbol{x}\right\|_{2} \forall \boldsymbol{x}\right\}  \tag{127}\\
& \geq \frac{o\left(\boldsymbol{x}^{(i)}, y^{(i)}\right)}{\sup _{\overline{\boldsymbol{x}} \in \operatorname{conv}(\mathcal{X})}\|J[\overline{\boldsymbol{x}}]\|_{2}} \tag{128}
\end{align*}
$$

The proof is done.

We now have all the needed ingredients to prove Theorem 10 and Theorem 11 .
Proof of Theorem 10 and Theorem 11 By Proposition 2, Lemma 5, and our assumption on complexity of the sample space, the RNN classifier is $\left(d_{y} \cdot\left(2 C_{M} / \gamma\right)^{k}, 0\right)$-robust, for some constant $C_{M}>0$. Due to Theorem 12 (with $M:=L_{g}$ there), it remains to prove the upper bound (102) for the classification margin of a training sample to complete the proof. Theorem 11 then follows from Theorem (with $M:=L_{g}$ there) and the inequality (103) follows immediately from Theorem 12

By Proposition 3, we have

$$
\begin{equation*}
\gamma^{d}\left(s_{i}\right) \geq \frac{o\left(s_{i}\right)}{\sup _{\hat{\boldsymbol{x}} \in \operatorname{conv}(\mathcal{X})}\|J[\hat{\boldsymbol{x}}]\|_{2}} \tag{129}
\end{equation*}
$$

where $J[\hat{\boldsymbol{x}}]:=d \mathcal{F}[\hat{\boldsymbol{x}}] / d \hat{\boldsymbol{x}}$ is the input-output Jacobian associated to the RNN. Therefore, to complete the proof it suffices to show that

$$
\begin{equation*}
\|J[\hat{\boldsymbol{x}}]\|_{2} \leq C \sum_{m=0}^{M-1} \delta_{m}\left\|\hat{\Phi}_{M, m+1}[\hat{\boldsymbol{x}}]\right\|_{2} \tag{130}
\end{equation*}
$$

where $C$ is the constant from the theorem and $\hat{\Phi}_{m+1, k}, 0 \leq k \leq m \leq M-1$ satisfies:

$$
\begin{align*}
\hat{\Phi}_{k, k} & =I  \tag{131}\\
\hat{\Phi}_{m+1, k} & =\hat{J}_{m} \hat{\Phi}_{m, k} \tag{132}
\end{align*}
$$

where $\hat{J}_{m}=I+\delta_{m} f^{\prime}\left(\hat{h}_{m}^{0}, \hat{x}_{m}\right)$ (with the $\hat{h}_{m}^{(0)}$ satisfying Eq. 78), recalling that we are replacing the superscript $\delta$ by hat when denoting the $\delta$-dependent approximating solutions for the sake of notation cleanliness) and the $\delta_{m}>0$ are the step sizes.
Iterating (132) up to the $(m+1)$ th step, for $m \geq k$, gives:

$$
\begin{equation*}
\hat{\Phi}_{m+1, k}=\hat{J}_{m} \hat{J}_{m-1} \cdots \hat{J}_{k}=: \prod_{l=k}^{m} \hat{J}_{l} \tag{133}
\end{equation*}
$$

Note that

$$
\begin{equation*}
\hat{\Phi}_{m+1, k}=\frac{\partial \hat{h}_{m+1}^{(0)}}{\partial \hat{h}_{m}^{(0)}} \frac{\partial \hat{h}_{m}^{(0)}}{\partial \hat{h}_{m-1}^{(0)}} \cdots \frac{\partial \hat{h}_{k+1}^{(0)}}{\partial \hat{h}_{k}^{(0)}}=\frac{d \hat{h}_{m+1}^{(0)}}{d \hat{h}_{k}^{(0)}} \tag{134}
\end{equation*}
$$

Now, applying chain rule:

$$
\begin{equation*}
J[\hat{\boldsymbol{x}}]=\frac{\partial p\left(V \hat{h}_{M}^{(0)}\right)}{\partial \hat{h}_{M}^{(0)}} \sum_{j=0}^{M-1} \frac{\partial \hat{h}_{M}^{(0)}}{\partial \hat{h}_{M-1}^{(0)}} \cdots \frac{\partial \hat{h}_{j+2}^{(0)}}{\partial \hat{h}_{j+1}^{(0)}} \frac{\partial \hat{h}_{j+1}^{(0)}}{\partial \hat{x}_{j}} \tag{135}
\end{equation*}
$$

where $p$ is the softmax function.
We compute:

$$
\begin{equation*}
\frac{\partial p\left(V \hat{h}_{M}^{(0)}\right)}{\partial \hat{h}_{M}^{(0)}}=V E \tag{136}
\end{equation*}
$$

where $E^{i j}=p^{i}\left(e^{i j}-p^{j}\right)$. From 134), we have

$$
\begin{equation*}
\frac{\partial \hat{h}_{M}^{(0)}}{\partial \hat{h}_{M-1}^{(0)}} \cdots \frac{\partial \hat{h}_{j+2}^{(0)}}{\partial \hat{h}_{j+1}^{(0)}}=\hat{\Phi}_{M, j+1} \tag{137}
\end{equation*}
$$

On the other hand,

$$
\begin{equation*}
\frac{\partial \hat{h}_{j+1}^{(0)}}{\partial \hat{x}_{j}}=\delta_{j} \frac{\partial f\left(\hat{h}_{j}^{(0)}, \hat{x}_{j}\right)}{\partial \hat{x}_{j}} \tag{138}
\end{equation*}
$$

for $j=0,1, \ldots, M-1$. Note that for Lipschitz RNNs, we have $\frac{\partial \hat{h}_{j+1}^{(0)}}{\partial \hat{x}_{j}}=\delta_{j} D_{j} U$, where $D_{l}^{i j}=$ $a^{\prime}\left(\left[W \hat{h}_{l}^{(0)}+U \hat{x}_{l}+b\right]^{i}\right) e_{i j}$.

Using the results of the above computations gives:

$$
\begin{equation*}
J[\hat{\boldsymbol{x}}]=V E \sum_{m=0}^{M-1} \delta_{m} \hat{\Phi}_{M, m+1} \frac{\partial f\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right)}{\partial \hat{x}_{m}} \tag{139}
\end{equation*}
$$

Therefore,

$$
\begin{align*}
\|J[\hat{\boldsymbol{x}}]\|_{2} & \leq\|V E\|_{2} \sum_{m=0}^{M-1}\left\|\delta_{m} \hat{\Phi}_{M, m+1}\right\|_{2}\left\|\frac{\partial f\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right)}{\partial \hat{x}_{m}}\right\|_{2}  \tag{140}\\
& \leq\|V\|_{2}\left(\max _{m=0,1, \ldots, M-1}\left\|\frac{\partial f\left(\hat{h}_{m}^{(0)}, \hat{x}_{m}\right)}{\partial \hat{x}_{m}}\right\|_{2}\right) \sum_{m=0}^{M-1} \delta_{m}\left\|\hat{\Phi}_{M, m+1}\right\|_{2}  \tag{141}\\
& =C \sum_{m=0}^{M-1} \delta_{m}\left\|\hat{\Phi}_{M, m+1}\right\|_{2} . \tag{142}
\end{align*}
$$

For the Lipschitz RNN, we have $C:=\|V\|\left(\max _{m=0,1, \ldots, M-1}\left\|D_{m} U\right\|_{2}\right)$. The proof is done.
It follows immediately from Eq. (142) that we have the following sufficient condition for stability with respect to hidden states of deterministic RNN to guarantee stability with respect to input sequence.
Corollary 1. Fix a $M$ and assume that $C \sum_{m=0}^{M-1} \delta_{m}<1$. Then, $\left\|\hat{\Phi}_{M, m+1}\right\|_{2} \leq 1$ for $m=$ $0, \ldots, M-1$ implies that $\|J[\hat{\boldsymbol{x}}]\|_{2}<1$.

## F Stability and Noise-Induced Stabilization for NRNNs: Proof of Theorem 3 in the Main Paper

We begin by discussing stochastic stability for SDEs, which are the underlying continuous-time models for our NRNNs.

Although the additional complexities of SDEs over ODEs often necessitate more involved analyses, many of the same ideas typically carry across. This is also true for stability. A typical approach for proving stability of ODEs involves Lyapunov functions - in Chapter 4 of [23], such approaches are extended for SDEs. This gives way to three notions of stability: (1) stability in probability; (2) moment stability; and (3) almost sure stability. Their definitions are provided in Definitions 4.2.1, 4.3.1, 4.4.1 in [23], and are repeated below for convenience.

To preface the definition, consider initializing (23) at two different random variables $h_{0}$ and $h_{0}^{\prime}:=$ $h_{0}+\epsilon_{0}$, where $\epsilon_{0} \in \mathbb{R}^{d_{h}}$ is a constant non-random perturbation with $\left\|\epsilon_{0}\right\| \leq \delta$. The resulting hidden states, $h_{t}$ and $h_{t}^{\prime}$, are set to satisfy (23) with the same Brownian motion $B_{t}$, starting from their initial values $h_{0}$ and $h_{0}^{\prime}$, respectively. The evolution of $\epsilon_{t}=h_{t}^{\prime}-h_{t}$ satisfies

$$
\begin{equation*}
\mathrm{d} \epsilon_{t}=A \epsilon_{t} \mathrm{~d} t+\Delta a_{t}\left(\epsilon_{t}\right) \mathrm{d} t+\Delta \sigma_{t}\left(\epsilon_{t}\right) \mathrm{d} B_{t} \tag{143}
\end{equation*}
$$

where $\Delta a_{t}\left(\epsilon_{t}\right)=a\left(W h_{t}^{\prime}+U x_{t}+b\right)-a\left(W h_{t}+U x_{t}+b\right)$ and $\Delta \sigma_{t}\left(\epsilon_{t}\right)=\sigma_{t}\left(h_{t}+\epsilon_{t}\right)-\sigma_{t}\left(h_{t}\right)$. Since $\Delta a_{t}(0)=0, \Delta \sigma_{t}(0)=0$ for all $t \in[0, T], \epsilon_{t}=0$ admits a trivial equilibrium for (143).
Definition 8 (Stability for SDEs). The trivial solution of the SDE (143) is
(i) stochastically stable (or, stable in probability) if for every $\epsilon \in(0,1), r>0$, there exists a $\delta=\delta(\epsilon, r)>0$ such that $\mathbb{P}\left(\left\|\epsilon_{t}\right\|<r\right.$ for all $\left.t \geq 0\right) \geq 1-\epsilon$ whenever $\left\|\epsilon_{0}\right\|<\delta$.
(ii) stochastically asymptotically stable if it is stochastically stable and, moreover, for every $\epsilon \in(0,1)$, there exists a $\delta_{0}=\delta_{0}(\epsilon)>0$ such that $\mathbb{P}\left(\lim _{t \rightarrow \infty} \epsilon_{t}=0\right) \geq 1-\epsilon$ whenever $\left\|\epsilon_{0}\right\|<\delta_{0}$.
(iii) almost surely exponentially stable if $\limsup _{t \rightarrow \infty} t^{-1} \log \left\|\epsilon_{t}\right\|<0$ with probability one whenever $\left\|\epsilon_{0}\right\|<\delta_{1}$.
(iv) p-th moment exponentially stable if there exists $\lambda, C>0$ such that $\mathbb{E}\left\|\epsilon_{t}\right\|^{p} \leq$ $C\left\|\epsilon_{0}\right\|^{p} e^{-\lambda\left(t-t_{0}\right)}$ for all $t \geq t_{0}$.

The properties in Definition 8 are said to hold globally if they also hold under no restrictions on $\epsilon_{0}$. Stability in probability neglects to quantify rates of convergence, and is implied by almost sure exponential stability. On the other hand, for our class of SDEs, $p$-th moment exponential stability would imply almost sure exponential stability (see Theorem 4.2 in [23]).

One critical difference between Lyapunov stability theory for ODEs and SDEs lies in the stochastic stabilization phenomenon. Let $L$ be the infinitesimal generator (for a given input signal $x_{t}$ ) of the diffusion process described by the SDE (143):

$$
\begin{equation*}
L=\frac{\partial}{\partial t}+\sum_{i}\left((A \epsilon)^{i}+\Delta a_{t}^{i}(\epsilon)\right) \frac{\partial}{\partial \epsilon^{i}}+\frac{1}{2} \sum_{i, j}\left[\Delta \sigma_{t}(\epsilon) \Delta \sigma_{t}(\epsilon)^{\top}\right]^{i j} \frac{\partial^{2}}{\partial \epsilon^{i} \partial \epsilon^{j}} \tag{144}
\end{equation*}
$$

The generator for the corresponding ODE arises by taking $\Delta \sigma_{t} \equiv 0$. In classical Lyapunov theory for ODEs, the existence of a non-negative Lyapunov function $V$ satisfying $L V \leq 0$ in some neighbourhood of the equilibrium is both necessary and sufficient for stability (see Chapter 4 of [15]) For SDEs, it has been shown that this condition is sufficient, but no longer necessary [23, 22]. This is by the nature of stochastic stabilization - the addition of noise can can have the surprising effect of increased stability over its deterministic counterpart. Of course, this is not universally the case as some forms of noise can be sufficiently extreme to induce instability; see Section 4.5 in [23].
Identifying sufficient conditions which quantify the stochastic stabilization phenomenon are especially useful in our setting, and as it turns out (see also [19]), these are most easily obtained for almost sure exponential stability. Therefore, our stability analysis will focus on establishing almost sure exponential stability. The objective is to analyze such stability of the solution $\epsilon_{t}=0$, that is, to see how the final state $\epsilon_{T}$ (and hence the output $y_{T}^{\prime}-y_{T}=V \epsilon_{T}$ of the RNN) changes for an arbitrarily small initial perturbation $\epsilon_{0} \neq 0$.

To this end, we consider an extension of the Lyapunov exponent to SDEs at the level of sample path [23].

Definition 9 (Almost sure global exponential stability). The sample (or pathwise) Lyapunov exponent of the trivial solution of (143) is $\Lambda=\lim \sup _{t \rightarrow \infty} t^{-1} \log \left\|\epsilon_{t}\right\|$. The trivial solution $\epsilon_{t}=0$ is almost surely globally exponentially stable if $\Lambda$ is almost surely negative for all $\epsilon_{0} \in \mathbb{R}^{d_{h}}$.

For the sample Lyapunov exponent $\Lambda(\omega)$, there is a constant $C>0$ and a random variable $0 \leq$ $\tau(\omega)<\infty$ such that for all $t>\tau(\omega),\left\|\epsilon_{t}\right\|=\left\|h_{t}^{\prime}-h_{t}\right\| \leq C e^{\Lambda t}$ almost surely. Therefore, almost sure exponential stability implies that almost all sample paths of (143) will tend to the equilibrium solution $\epsilon=0$ exponentially fast. With this definition in tow, we state and prove our primary stability result, which is equivalent to Theorem 3 in the main paper.
Theorem 13 (Bounds for sample Lyapunov exponent of the trivial solution). Assume that Assumption $C$ holds. Suppose that a is $L_{a}$-Lipschitz, $0 \leq a_{\Delta}^{T}(\epsilon, t) \epsilon \leq L_{a}\|\epsilon\|_{2}^{2}$ and $0 \leq \sigma_{1}\|\epsilon\| \leq\left\|\Delta \sigma_{t}(\epsilon)\right\|_{F} \leq$ $\sigma_{2}\|\epsilon\|$ for all nonzero $\epsilon \in \mathbb{R}^{d_{h}}, t \in[0, T]$. Then, with probability one,

$$
\begin{equation*}
-\sigma_{2}^{2}+\frac{\sigma_{1}^{2}}{2}+\lambda_{\min }\left(A^{\mathrm{sym}}\right) \leq \Lambda \leq-\sigma_{1}^{2}+\frac{\sigma_{2}^{2}}{2}+L_{a} \sigma_{\max }(W)+\lambda_{\max }\left(A^{\mathrm{sym}}\right) \tag{145}
\end{equation*}
$$

for any $\epsilon_{0} \in \mathbb{R}^{d_{h}}$.
To establish the bounds in Theorem 13, we appeal to the following theorem, which arises from combining Theorems 4.3.3 and 4.3.5 in [23] in the case $p=2$. Here, for a function $V$, we let $V_{\epsilon}=\partial V / \partial \epsilon$.

Theorem 14 (Stochastic Lyapunov theorem). If there exists a function $V \in \mathcal{C}^{2,1}\left(\mathbb{R}^{d_{h}} \times \mathbb{R}^{+} ; \mathbb{R}^{+}\right)$ and $c_{1}, C_{1}>0, c_{2}, C_{2} \in \mathbb{R}, c_{3}, C_{3} \geq 0$ such that for all $\epsilon \neq 0$ and $t \geq t_{0}$,
(i) $c_{1}\|\epsilon\|^{2} \leq V(\epsilon, t) \leq C_{1}\|\epsilon\|^{2}$,
(ii) $c_{2} V(\epsilon, t) \leq L V(\epsilon, t) \leq C_{2} V(\epsilon, t)$, and
(iii) $c_{3} V(\epsilon, t)^{2} \leq\left\|V_{\epsilon}(\epsilon, t) \Delta \sigma_{t}(\epsilon)\right\|_{F}^{2} \leq C_{3} V(\epsilon, t)^{2}$,
then, with probability one, the Lyapunov exponent $\Lambda$ lies in the interval

$$
\begin{equation*}
\frac{2 c_{2}-C_{3}}{4} \leq \Lambda \leq-\frac{c_{3}-2 C_{2}}{4} . \tag{146}
\end{equation*}
$$

The proof of Theorem 14 involves the Itô formula, an exponential martingale inequality and a BorelCantelli type argument. The functions $V$ above are called stochastic Lyapunov functions and the use of the theorem involves construction of these functions. We are now in a position to prove Theorem 13. and will find that the choice $V(\epsilon, t)=\|\epsilon\|^{2}$ will suffice.

Proof of Theorem 13. It suffices to verify the conditions of Theorem 2 with $V(\epsilon, t)=V(\epsilon)=\|\epsilon\|^{2}$. Clearly (i) is satisfied. To show (iii), by the conditions on $\Delta \sigma_{t}$, we have that $4 \sigma_{1}^{2}\|\epsilon\|^{4} \leq$ $\left\|V_{\epsilon}(\epsilon) \Delta \sigma_{t}(\epsilon)\right\|_{F}^{2} \leq 4 \sigma_{2}^{2}\|\epsilon\|^{4}$. It remains only to show (ii). Observe that

$$
L V(\epsilon)=\epsilon^{\top}\left(A+A^{\top}\right) \epsilon+2 \Delta a_{t}(\epsilon) \epsilon+\operatorname{tr}\left(\Delta \sigma_{t}(\epsilon) \Delta \sigma_{t}(\epsilon)^{\top}\right)
$$

Since $0 \leq \Delta a_{t}(\epsilon) \epsilon$ and

$$
\begin{aligned}
\left|\Delta a_{t}(\epsilon) \epsilon\right| & \leq\left\|a\left(W h_{t}^{\prime}+U x_{t}+b\right)-a\left(W h_{t}+U x_{t}+b\right)\right\|\|\epsilon\| \\
& \leq L_{a}\|W \epsilon\|\|\epsilon\| \leq L_{a} \sigma_{\max }(W)\|\epsilon\|^{2},
\end{aligned}
$$

it follows that

$$
L V(\epsilon) \leq\left(2 \lambda_{\max }\left(A^{s y m}\right)+2 L_{a} \sigma_{\max }(W)+\sigma_{2}^{2}\right)\|\epsilon\|^{2}
$$

and

$$
L V(\epsilon) \geq\left(2 \lambda_{\min }\left(A^{\text {sym }}\right)+\sigma_{1}^{2}\right)\|\epsilon\|^{2} .
$$

The bound (146) now follows from Theorem 14 with $c_{1}=C_{1}=1, c_{2}=2 \lambda_{\min }\left(A^{\text {sym }}\right)+\sigma_{1}^{2}$, $C_{2}=2 \lambda_{\max }\left(A^{\text {sym }}\right)+2 L_{a} \sigma_{\max }(W)+\sigma_{2}^{2}, c_{3}=4 \sigma_{1}^{2}$, and $C_{3}=4 \sigma_{2}^{2}$.
Remark 5. To see if the bounds in Theorem 13 are indeed sharp (at least for certain cases), consider the linear $\operatorname{SDE} d H_{t}=A H_{t} d t+B H_{t} d W_{t}$, where $A \in \mathbb{R}^{d_{h} \times d_{h}}, B=\sigma I, \sigma \in \mathbb{R}$ and $W_{t}$ is a scalar Wiener process. Then, since $A$ and $B$ commute, they can be simultaneously diagonalized, and so the linear SDE can be reduced via transformation to a set of independent one-dimensional linear SDEs. In particular, one can show that $\left.H_{t}=\exp \left(\left(A-B^{2} / 2\right) t+B W_{t}\right)\right) H_{0}$ and the Lyapunov exponents $\Lambda$ of this system are the real part of the eigenvalues of $A-B^{2} / 2$. Note that $\lambda_{\min }\left(A_{\text {sym }}-B^{2} / 2\right) \leq$ $\Lambda \leq \lambda_{\max }\left(A_{\text {sym }}-B^{2} / 2\right)$ a.s.. Since $B=\sigma I$, this inequality implies Eqn. (145) with $L_{a}:=0$. The bounds are tight in the scalar case ( $d_{h}=1$ and $A$ is a scalar), with the inequality becoming an equality.
Remark 6. Even in the additive noise setting, however, the Lyapunov exponents of the CT-NRNN driven by additive noise are not generally the same as those of the corresponding deterministic CTRNN. Oseledets multiplicative ergodic theorem implies they will be the same if the data generating process $x_{t}$ is ergodic [1]. Characterizing Lyapunov exponents for SDEs is a non-trivial affair in general - we refer to, for instance, [2] for details on this.

## G Experimental Details

## G. 1 Experimental Results Presented in the Main Paper

Following [6], we construct the hidden-to-hidden weight matrices $A$ and $W$ as

$$
\begin{align*}
A & =T\left(B, \beta_{a}, \gamma_{a}\right):=\left(1-\beta_{a}\right) \cdot\left(B+B^{T}\right)+\beta_{a} \cdot\left(B-B^{T}\right)-\gamma_{a} I  \tag{147}\\
W & =T\left(C, \beta_{w}, \gamma_{w}\right):=\left(1-\beta_{w}\right) \cdot\left(C+C^{T}\right)+\beta_{w} \cdot\left(C-C^{T}\right)-\gamma_{w} I \tag{148}
\end{align*}
$$

Here, $B$ and $C$ denote weight matrices that have the same dimensions as $A$ and $W$. The tuning parameters $\gamma_{a}$ and $\gamma_{w}$ can be used to increase dampening. We initialize the weight matrices by sampling weights from the normal distribution $\mathcal{N}\left(0, \sigma_{\text {init }}^{2}\right)$, where $\sigma_{\text {init }}^{2}$ is the variance. Table 4 summarizes the tuning parameters that we have used in our experiments. We train our models for 100 epochs, with scheduled learning rate decays at epochs $\{90\}$. We use Adam with default parameters for minimizing the objective.
We performed a random search to obtain the tuning parameters. Since our model is closely related to the Lipschitz RNN, we started with the tuning parameters proposed in [6]. We evaluated different noise levels, both for multiplicative and additive noise, in the range $[0.01,0.1]$. We tuned the levels of noise-injection so that the models achieve state-of-the-art performance on clean input data. Further, we observed that the robustness of the model is not significantly improving when trained with

Table 4: Tuning parameters used to train the NRNN.

| Name | $\mathrm{d} \_\mathrm{h}$ | lr | decay | $\beta$ | $\gamma_{a}$ | $\gamma_{w}$ | $\epsilon$ | $\sigma_{\text {init }}^{2}$ | add. noise | mult. noise |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Ordered MNIST | 128 | 0.001 | 0.1 | 0.75 | 0.001 | 0.001 | 0.01 | $0.1 / 128$ | 0.02 | 0.02 |
| Ordered MNIST | 128 | 0.001 | 0.1 | 0.75 | 0.001 | 0.001 | 0.01 | $0.1 / 128$ | 0.05 | 0.02 |
| Permuted MNIST | 128 | 0.001 | 0.1 | 0.75 | 0.001 | 0.001 | 0.01 | $0.1 / 128$ | 0.02 | 0.02 |
| Permuted MNIST | 128 | 0.001 | 0.1 | 0.75 | 0.001 | 0.001 | 0.01 | $0.1 / 128$ | 0.05 | 0.02 |
| ECG | 128 | 0.001 | 0.1 | 0.9 | 0.001 | 0.001 | 0.1 | $0.1 / 128$ | 0.06 | 0.03 |



Figure 3: Hessian loss landscapes for deterministic (left) and noisy (right) model, computed using PyHessian.
increased levels of noise-injections. Overall, our experiments indicated that the model is relatively insensitive to the particular amount of additive and multiplicative noise level in the small noise regime.

Further, we need to note that we only considered models that used a combination of additive and multiplicative noise-injections. One could also train models using either only additive or multiplicative noise-injections. We did not investigate in detail the trade-offs between the different strategies. The motivation for our experiments was to demonstrate that (i) models trained with noise-injections can achieve state-of-the-art performance on clean input data, and (ii) such models are also more resilient to input perturbations.

Figure 3 shows that NRNN exhibits a smoother Hessian landscape than that of the deterministic counterpart.

For establishing a fair set of baselines, we used the following implementations and prescribed tuning parameters for the other models that we considered.

- Exponential RNN. We used the following implementation: https://github.com/ Lezcano/expRNN. We used the default parameters. We trained the model, with hidden dimension $d_{h}=128$, for 100 epochs.
- CoRNN. We used the following implementation, provided as part of the Supplementary Material: https://openreview.net/forum?id=F3s69XzWOia We used the default parameters proposed by the authors for training the model with hidden dimension $d_{h}=128$. We trained the model for 100 epochs with learning rate decay at epoch 90 .
- Lipschitz RNN. We used the following implementation, provided as part of the Supplementary Material: https://openreview.net/forum?id=-N7PBXqOUJZ. We used the default parameters proposed by the authors for training the model with hidden dimension $d_{h}=128$. We trained the model for 100 epochs with learning rate decay at epoch 90 .
- Antisymmetric RNN. To our best knowledge, there is no public implementation by the authors for this model. However, the Antisymmetric RNN can be seen as a special case of the Lipschitz RNN or the NRNN, without the stabilizing term $A$ and without noise-injection. We trained this model by using our implementation and the following tuning parameters: $\beta=1.0, \gamma=0.001, \operatorname{lr}=0.002, \epsilon=0.01$. We trained the model for 100 epochs with learning rate decay at epoch 90 .

Table 5: Robustness w.r.t. white noise $(\sigma)$ and $\mathrm{S} \& \mathrm{P}(\alpha)$ perturbations on the permuted MNIST task.

| Name | clean | $\sigma=0.1$ | $\sigma=0.2$ | $\sigma=0.3$ | $\alpha=0.03$ | $\alpha=0.05$ | $\alpha=0.1$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Antisymmetric RNN [5] | $92.8 \%$ | $92.4 \%$ | $89.5 \%$ | $81.9 \%$ | $90.5 \%$ | $87.9 \%$ | $72.6 \%$ |
| CoRNN [24] | $96.05 \%$ | $65.1 \%$ | $38.25 \%$ | $29.1 \%$ | $84.8 \%$ | $73.8 \%$ | $52.6 \%$ |
| Exponential RNN [18] | $93.3 \%$ | $90.6 \%$ | $78.4 \%$ | $61.6 \%$ | $80.4 \%$ | $70.6 \%$ | $51.6 \%$ |
| Lipschitz RNN [6] | $\mathbf{9 5 . 9 \%}$ | $95.4 \%$ | $93.5 \%$ | $83.7 \%$ | $93.7 \%$ | $90.2 \%$ | $70.8 \%$ |
| NRNN (mult./add. noise: $0.02 / 0.02)$ | $94.9 \%$ | $\mathbf{9 4 . 8} \%$ | $\mathbf{9 4 . 6 \%}$ | $94.3 \%$ | $\mathbf{9 4 . 0 \%}$ | $93.1 \%$ | $88.6 \%$ |
| NRNN (mult./add. noise: $0.02 / 0.05$ ) | $94.7 \%$ | $94.6 \%$ | $\mathbf{9 4 . 6 \%}$ | $\mathbf{9 4 . 4 \%}$ | $\mathbf{9 4 . 0 \%}$ | $\mathbf{9 3 . 2} \%$ | $\mathbf{9 0 . 5 \%}$ |


(a) White noise perturbations.

(b) Salt and pepper perturbations.

Figure 4: Test accuracy for the permuted MNIST task as function of the strength of input perturbations.

## G. 2 Additional Results for Permuted Pixel-by-Pixel MNIST Classification

Here we consider the permuted pixel-by-pixel MNIST classification task. This task sequentially presents a scrambled sequence of the 784 pixels to the model and uses the final hidden state to predict the class membership probability of the input image.
Table 5 shows the average test accuracy (evaluated for models that are trained with 10 different seed values). Here we present results for white noise and salt and pepper (S\&P) perturbations. Again, the NRNNs show an improved resilience to input perturbations. Figure 4 summarizes the performance of different models with respect to white noise and salt and pepper perturbations.

## References

[1] Ludwig Arnold, W Kliemann, and E Oeljeklaus. Lyapunov exponents of linear stochastic systems. In Lyapunov exponents, pages 85-125. Springer, 1986.
[2] Ludwig Arnold and Wolfgang Kliemann. Large deviations of linear stochastic differential equations. In Stochastic differential systems, pages 115-151. Springer, 1987.
[3] Yurii Nikolaevich Blagoveshchenskii and Mark Iosifovich Freidlin. Some properties of diffusion processes depending on a parameter. In Doklady Akademii Nauk, volume 138, pages 508-511. Russian Academy of Sciences, 1961.
[4] Alexander Camuto, Matthew Willetts, Umut Simsekli, Stephen J Roberts, and Chris C Holmes. Explicit regularisation in Gaussian noise injections. In Advances in Neural Information Processing Systems, volume 33, pages 16603-16614, 2020.
[5] Bo Chang, Minmin Chen, Eldad Haber, and Ed H. Chi. AntisymmetricRNN: A dynamical system view on recurrent neural networks. In International Conference on Learning Representations, 2019.
[6] N. Benjamin Erichson, Omri Azencot, Alejandro Queiruga, Liam Hodgkinson, and Michael W. Mahoney. Lipschitz recurrent neural networks. In International Conference on Learning Representations, 2021.
[7] Wei Fang. Adaptive timestepping for SDEs with non-globally Lipschitz drift. PhD thesis, University of Oxford, 2019.
[8] Wei Fang, Michael B Giles, et al. Adaptive Euler-Maruyama method for SDEs with nonglobally Lipschitz drift. Annals of Applied Probability, 30(2):526-560, 2020.
[9] Mark Iosifovich Freidlin and Alexander D Wentzell. Random Perturbations of Dynamical Systems. Springer, 1998.
[10] J.F.L. Gall. Brownian Motion, Martingales, and Stochastic Calculus. Graduate Texts in Mathematics. Springer International Publishing, 2016.
[11] Desmond J Higham, Xuerong Mao, and Andrew M Stuart. Strong convergence of Euler-type methods for nonlinear stochastic differential equations. SIAM Journal on Numerical Analysis, 40(3):1041-1063, 2002.
[12] Martin Hutzenthaler and Arnulf Jentzen. Numerical approximations of stochastic differential equations with non-globally Lipschitz continuous coefficients, volume 236. American Mathematical Society, 2015.
[13] Martin Hutzenthaler, Arnulf Jentzen, Peter E Kloeden, et al. Strong convergence of an explicit numerical method for SDEs with nonglobally Lipschitz continuous coefficients. The Annals of Applied Probability, 22(4):1611-1641, 2012.
[14] Ioannis Karatzas and Steven E Shreve. Brownian Motion. Springer, 1998.
[15] Hassan K Khalil and Jessy W Grizzle. Nonlinear Systems, volume 3. Prentice hall Upper Saddle River, NJ, 2002.
[16] Peter E Kloeden and Eckhard Platen. Numerical Solution of Stochastic Differential Equations, volume 23. Springer Science \& Business Media, 2013.
[17] Hiroshi Kunita. Stochastic differential equations and stochastic flows of diffeomorphisms. In Ecole d'été de Probabilités de Saint-Flour XII-1982, pages 143-303. Springer, 1984.
[18] Mario Lezcano-Casado and David Martınez-Rubio. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. In International Conference on Machine Learning, pages 3794-3803, 2019.
[19] Xuanqing Liu, Tesi Xiao, Si Si, Qin Cao, Sanjiv Kumar, and Cho-Jui Hsieh. How does noise help robustness? Explanation and exploration under the neural SDE framework. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 282-290, 2020.
[20] Yiping Lu, Aoxiao Zhong, Quanzheng Li, and Bin Dong. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In International Conference on Machine Learning, pages 3276-3285. PMLR, 2018.
[21] Simon JA Malham and Anke Wiese. An introduction to SDE simulation. arXiv preprint arXiv:1004.0646, 2010.
[22] Xuerong Mao. Exponential Stability of Stochastic Differential Equations. Marcel Dekker, 1994.
[23] Xuerong Mao. Stochastic Differential Equations and Applications. Elsevier, 2007.
[24] T. Konstantin Rusch and Siddhartha 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.
[25] Simo Särkkä and Arno Solin. Applied Stochastic Differential Equations, volume 10. Cambridge University Press, 2019.
[26] Shai Shalev-Shwartz and Shai Ben-David. Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press, 2014.
[27] Jure Sokolić, Raja Giryes, Guillermo Sapiro, and Miguel RD Rodrigues. Generalization error of deep neural networks: Role of classification margin and data structure. In 2017 International Conference on Sampling Theory and Applications (SampTA), pages 147-151. IEEE, 2017.
[28] Colin Wei and Tengyu Ma. Data-dependent sample complexity of deep neural networks via Lipschitz augmentation. In Advances in Neural Information Processing Systems, volume 32, 2019.
[29] E Weinan. A proposal on machine learning via dynamical systems. Communications in Mathematics and Statistics, 5(1):1-11, 2017.
[30] Huan Xu and Shie Mannor. Robustness and generalization. Machine Learning, 86(3):391-423, 2012.

