RNN-based Online Learning: An Efficient First-Order Optimization Algorithm with a Convergence Guarantee

N. Mert Vural, Selim F. Yilmaz, Fatih Ilhan and Suleyman S. Kozat, *Senior Member, IEEE*

arXiv:2003.03601v1 [cs.LG] 7 Mar 2020

***Abstract*—We investigate online nonlinear regression with continually running recurrent neural network networks (RNNs), i.e., RNN-based online learning. For RNN-based online learning, we introduce an efficient first-order training algorithm that theoretically guarantees to converge to the optimum network parameters. Our algorithm is truly online such that it does not make any assumption on the learning environment to guarantee convergence. Through numerical simulations, we verify our theoretical results and illustrate significant performance improvements achieved by our algorithm with respect to the state-of-the-art RNN training methods.**

***Index Terms*—Online learning, neural network training, recurrent neural networks, sequential learning, regression, online gradient descent.**

I. Introduction

Prediction of individual sequences is one of the main subjects of interest in the contemporary online learning literature [1]. In this problem, we sequentially receive a data sequence related to a desired signal to predict the signal's next value [2]. This problem is also known as online regression and extensively studied in the neural network [3], [4], signal processing [5], [6], and machine learning literatures [1], [2]. In these studies, nonlinear approaches are generally employed since linear modeling is inadequate for a wide range of applications due to constraint of linearity [3], [6].

For online regression, there exists a wide range of nonlinear approaches in the fields of machine learning and signal processing [6]–[8]. However, these approaches usually suffer from prohibitive computational requirements and may provide poor performance due to overfitting and stability issues [5]. Adopting neural networks is another method for online nonlinear regression due to their success in approximating nonlinear functions [3], [9]. However, neural network-based regression algorithms are shown to be prone to such issues as overfitting or demonstrating inadequate performance in certain applications [10], [11]. To overcome the limitations of the shallow networks, neural networks composed of multiple layers, i.e., deep neural networks (DNNs), have recently been introduced. In DNNs, each layer performs a feature extraction based on the previous layers, which enables them to model highly nonlinear structures [12], [13]. However, this layered structure poorly performs in capturing time dependencies of the data due to its lack of temporal memory. Therefore, DNNs provide only limited performance in processing temporal data and modeling time series [14]. To remedy this issue, recurrent neural networks (RNNs) are used, as these networks have a feed-back connection that enables them to store past information [7], [15]. Through their many variants, recurrent neural networks have seen remarkable empirical success in a broad range of sequential learning domains [4], [16]–[18]. In this study, we are interested in online nonlinear regression with RNN-based networks due to their superior performance in capturing time dependencies.

For RNNs, there exists a wide range of online training methods to learn network parameters [4], [19]–[23]. Among them, the first-order methods are commonly preferred in practice due to their computational efficiency [19], [22], [23]. However, using gradient-based optimization methods for RNN-based online learning is challenging due to the divergence problems caused by the exploding gradient problem [24]. In addition to this, finding an effective learning rate for first-order methods requires time-consuming search algorithms [25], which costs a significant amount of time and effort in practical applications.

To resolve these problems, several heuristic methods have been proposed. For the divergence problems: Krueger et al. [26] penalized the squared distance between the norms of successive hidden states to keep the RNN model stable during training. Mikolov et al. [26] and Pascanu et al. [27] showed that gradient clipping helps to reduce exploding gradients in practice. To overcome the learning rate-related problems: Blier et al. [28] applied multiple learning rates with adaptive randomization, which is shown to perform close to the stochastic gradient descent (SGD) with the optimal learning rate. As an alternative approach, Orabona et al. [29] randomly performed gradient descent updates with a fixed learning rate, which is again empirically shown to be successful in neural network training.

While these heuristics are reasonably successful in practice, they are ad-hoc and based on empirical observations, which may not necessarily be applicable in online settings. There are also mathematical studies in the literature to provide theoretical performance guarantees for RNN-based learning: Hardt et al. [30] showed that the gradient descent algorithm learns the globally optimum weights when the learning model is a single input single output linear dynamical system. Oymak et al. [31] extended this result to the contractive nonlinear dynamical systems by assuming that the ground truth hidden

This work is supported in part by TUBITAK Contract No. 117E153.

N. M. Vural, S. F. Yilmaz, F. Ilhan and S. S. Kozat are with the Department of Electrical and Electronics Engineering, Bilkent University, Ankara 06800, Turkey, e-mail: vural@ee.bilkent.edu.tr, syilmaz@ee.bilkent.edu.tr, filhan@ee.bilkent.edu.tr, kozat@ee.bilkent.edu.tr.

state vectors are observed. Additionally, Allen-Zhu et al. [32] showed the Elman Network trained with SGD is capable of minimizing the regression loss with the assumption that the number of neurons is polynomial in the training data size.

Although these studies provide definite solutions for potential problems of the first-order methods, the conditions of their results are restrictive for online settings. Therefore, their results are usually inapplicable to practical online learning scenarios. In this study, differing from the previous works, we introduce a first-order optimization algorithm that theoretically guarantees to learn the locally optimum parameters without any assumption on the input/output sequences (except their boundedness) or the system dynamics of the model. We emphasize that our convergence guarantee is valid when our algorithm is used with commonly used RNN models, e.g., Elman networks [33] and LSTMs [34].

To obtain this result, we model RNN optimization as a sequential learning problem and assume each time step as a separate loss function, whose argument is the network weights (detailed in the following). By using the Lipschitz characteristics of these loss functions, we develop an online gradient descent algorithm that guarantees to converge to the weights with zero derivatives. In the simulations, we verify our theoretical results and show that our algorithm improves the error performance of the state-of-the-art methods [19], [22], [23] on several real-life datasets. Therefore, in this paper, we introduce a both practical and theoretically justified algorithm that can be used safely in RNN-based online learning settings.

Our contributions can be summarized as follows:

- To the best of our knowledge, we, as the first time in the literature, introduce an online first-order optimization algorithm that guarantees to learn the locally optimum parameters when used with practical RNN models. We note that previous studies in the literature either make ad-hoc assumptions in their results or give a theoretical justification for restricted settings that do not sufficiently describe the practical scenarios.
- Our algorithm is truly online such that it does not make any assumption on the desired data sequence to guarantee convergence. Therefore, it can be safely used in *any* practical application.
- Through simulations involving real-life datasets, we illustrate significant performance gains achieved by our algorithm with respect to the state-of-the-art methods [19], [22], [23].

This paper is organized as follows: In Section II, we formally introduce the online regression problem and describe our RNN model. In Section III, we develop a first-order optimization algorithm with a theoretical convergence guarantee. In Section IV, we verify our results and demonstrate the performance of our algorithm with numerical simulations. In Section V. we conclude the paper with final remarks.

II. Model and Problem Definition

All vectors are column vectors and denoted by boldface lower case letters. Matrices are represented by boldface capital letters. We use $\|\cdot\|$ (respectively $\|\cdot\|_\infty$) to denote the ℓ_2 (respectively ℓ_∞) vector or matrix norms depending on the argument. We use bracket notation $[n]$ to denote the set of the first n positive integers, i.e., $[n] = \{1, \cdots, n\}$.

We study online regression, where we observe a desired signal $\{d_t\}_{t\geq 1}$ and regression vectors $\{\mathbf{x}_t\}_{t\geq 1}$ to sequentially estimate d_t with $\hat{d}_t$. For mathematical convenience in the proofs, we assume $d_t \in [-\sqrt{n_h}, \sqrt{n_h}]$, with a user-dependent $n_h \in \mathbb{N}$, and $\mathbf{x}_t \in [-1, 1]^{n_x}$. However, our derivations hold for any bounded input/output sequence after proper shifting and scaling. Given our estimate $\hat{d}_t$, which is a function of $\{\cdots, \mathbf{x}_{t-1}, \mathbf{x}_t\}$ and $\{\cdots, d_{t-2}, d_{t-1}\}$, we suffer the loss $\ell(d_t, \hat{d}_t)$. Our aim is to optimize the network with respect to the loss function $\ell(\cdot, \cdot)$. In this study, we particularly work with the squared loss, i.e., $\ell(\hat{d}_t, d_t) = 0.5(d_t - \hat{d}_t)^2$.[1] However, since our work uses a generic approach for the gradient-based non-convex optimization, it can be extended for any continuous cost function. An extension for the cross-entropy loss is given in Appendix C.

In this paper, we study online regression with continually running RNNs [35]. For this task, we use the Elman network model, i.e.,

$$\mathbf{h}_{t+1} = \tanh(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t) \tag{1}$$

$$\hat{d}_t = \mathbf{c}^T \mathbf{h}_t. \tag{2}$$

Here, as the weight matrices, we have $\mathbf{W} \in \mathbb{R}^{n_h \times n_h}$, $\mathbf{U} \in \mathbb{R}^{n_h \times n_x}$ and $\mathbf{c} \in \mathbb{R}^{n_h}$, with $\|\mathbf{c}\| \leq 1$.[2] Moreover, $\mathbf{h}_t \in [-1, 1]^{n_h}$ is the hidden state vector, $\mathbf{x}_t \in [-1, 1]^{n_x}$ is the input vector, $\hat{d}_t \in [-\sqrt{n_h}, \sqrt{n_h}]$ is our estimation, and $\tanh$ applies to vectors point-wise. We note that we use the Elman network model with $\tanh$ and linear activations in (1)-(2) due to the wide use and simplicity of this model [36]. However, our derivations can be extended for any differentiable neural network architecture, given that the Lipschitz properties of the architecture can be explicitly derived. The sketch of such an extension for LSTMs is given in Appendix D. Additionally, although we do not explicitly write the bias terms, they can be included in (1)-(2) by augmenting the input vectors with a constant dimension.

III. Algorithm Development

In this section, we develop an online first-order algorithm that guarantees convergence to the locally optimum weights. We develop our algorithm in three subsections. In the first subsection, we describe our sequential learning approach and make definitions for the following analysis. In the second subsection, we present some auxiliary results that will be used in the development of the main algorithm. In the last subsection, we introduce our algorithm and mathematically prove its convergence guarantee.

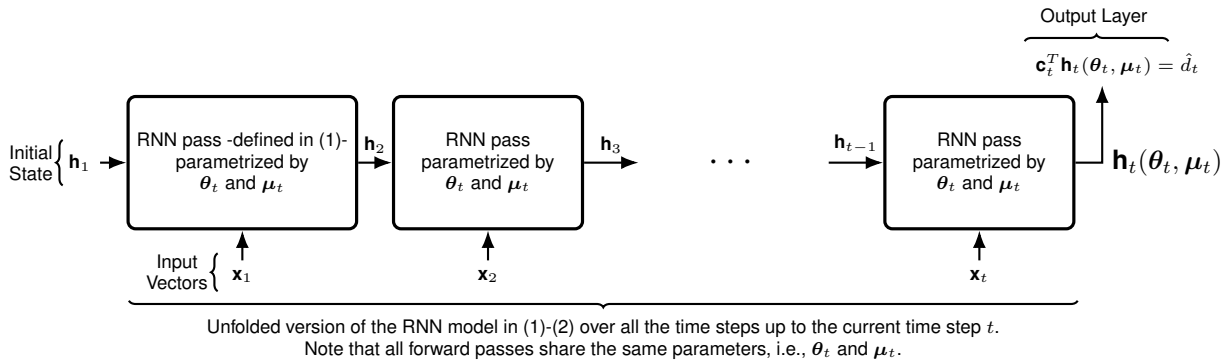

Fig. 1: In this figure, we visually describe $\mathbf{h}_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$, where $\mathbf{h}_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$ is defined as the hidden state vector obtained by running the model in (1)-(2) with $\boldsymbol{\theta}_t$ and $\boldsymbol{\mu}_t$ from the initial time step up to the current time step t. In the figure, the RNN sequence is initialized with a predetermined initial state $\mathbf{h}_1$, which is independent of the network weights. Then, the same RNN forward pass (given in (1)) is repeatedly applied to the input sequence $\{\mathbf{x}_t\}_{t\geq 1}$, where all the iterations are parametrized by $\boldsymbol{\theta}_t$ and $\boldsymbol{\mu}_t$. The resulting hidden vector after t iterations is defined to be $\mathbf{h}_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$. Here, we note that the dependence of $\mathbf{h}_t(\cdot,\cdot)$ on t is due to the increased length of the recursion at each time step.

A. Sequential Learning Approach

We investigate the RNN-based regression problem problem in the sequential learning framework. Here, we make no statistical assumptions on the data in order to model chaotic, non-stationary or even adversarial environments. Hence, we model our problem as a game between a learner and an adversary, where the learner is tasked with predicting the weight matrices from some convex sets in each round t. In the following, we use the vectorized forms of the weight matrices, i.e., $\boldsymbol{\theta}_t = \text{vec}(\mathbf{W}_t)$ and $\boldsymbol{\mu}_t = \text{vec}(\mathbf{U}_t)$, for mathematical convenience. Therefore, we construct the game as follows: At each round t, the learner declares his prediction $\boldsymbol{\theta}_t$ and $\boldsymbol{\mu}_t$; concurrently, the adversary chooses a target value $d_t \in [-\sqrt{n_h}, \sqrt{n_h}]$, an input $\mathbf{x}_t \in [-1,1]^{n_x}$, and a weight vector $\mathbf{c}_t$, $\|\mathbf{c}_t\| \leq 1$; then, the learner observes the loss function

$$\ell_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t) := 0.5\big(d_t - \underbrace{\mathbf{c}_t^T \mathbf{h}_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)}_{\hat{d}_t}\big)^2 \tag{3}$$

and suffers the loss $\ell_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$, where $\mathbf{h}_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$ denotes the hidden state vector obtained by running the model in (1)-(2) with $\boldsymbol{\theta}_t$ and $\boldsymbol{\mu}_t$ from the initial time step up to the current time step t (for the detailed description of $\mathbf{h}_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$ see Fig. 1). This procedure of play is repeated across T rounds, where T is the total number of input instances. Here, we note that we constructed our setting for adversarial $\mathbf{c}_t$ selections for mathematical convenience in our proofs. However, since the selected $\ell_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$ is convex with respect to $\mathbf{c}_t$, we will use the online gradient descent algorithm [37] to learn the optimal output layer weights (simultaneously with the hidden layer weights) during the training.

Since we assume no statistical assumptions on the input/output sequences, we analyze our performance with the notion of *regret*. However, we note that the standard regret definition for the convex problems is intractable in the non-convex settings due to the NP-hardness of the non-convex global optimization [38], [39]. Therefore, we use the notion of *local regret* recently introduced by Hazan et al. [39], which quantifies the objective of predicting points with a small gradient on average.

To formulate the local regret for our setting, we first define the projected partial derivatives of $\ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})$ with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\mu}$ as follows:

$$\frac{\partial_{\mathcal{K}_\theta,\eta_\theta} \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\theta}} := \frac{1}{\eta_\theta}\Big(\boldsymbol{\theta} - \Pi_{\mathcal{K}_\theta}\Big[\boldsymbol{\theta} - \eta_\theta \frac{\partial \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\theta}}\Big]\Big), \tag{4}$$

$$\frac{\partial_{\mathcal{K}_\mu,\eta_\mu} \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\mu}} := \frac{1}{\eta_\mu}\Big(\boldsymbol{\mu} - \Pi_{\mathcal{K}_\mu}\Big[\boldsymbol{\mu} - \eta_\mu \frac{\partial \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\mu}}\Big]\Big). \tag{5}$$

Here, $\partial_{\mathcal{K},\eta}$ denotes the projected partial derivative operator defined with some convex set $\mathcal{K}$ and some learning rate η, the learning rates η_θ and η_μ are used to update $\boldsymbol{\theta}_t$ and $\boldsymbol{\mu}_t$, the operators $\Pi_{\mathcal{K}_\theta}[\cdot]$ and $\Pi_{\mathcal{K}_\mu}[\cdot]$ denote the orthogonal projections onto $\mathcal{K}_\theta$ and $\mathcal{K}_\mu$.

We define the time-smoothed loss at time t, parametrized by some window size $w \in [T]$, as

$$L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu}) := \frac{1}{w}\sum_{i=0}^{w-1} \ell_{t-i}(\boldsymbol{\theta}, \boldsymbol{\mu}). \tag{6}$$

Then, we define the local regret as

$$R_w(T) := \sum_{t=1}^{T}\Big(\Big\|\frac{\partial_{\mathcal{K}_\theta,\eta_\theta} L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)}{\partial \boldsymbol{\theta}}\Big\|^2 + \Big\|\frac{\partial_{\mathcal{K}_\mu,\eta_\mu} L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)}{\partial \boldsymbol{\mu}}\Big\|^2\Big). \tag{7}$$

Our aim is to derive a sublinear upper bound for $R_w(T)$ in order to ensure the convergence of our algorithm to the locally optimum weights. However, before the derivations, we first present some auxiliary results, i.e., the Lipschitz and smoothness properties of $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu})$, which will be used in the convergence proof of our algorithm.

B. Lipschitz and Smoothness Properties

In this section, we derive the Lipschitz and smoothness properties of the time-smoothed loss function $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu})$. We

[1] We scale the squared loss with 0.5 for mathematical convenience when computing the derivatives.

[2] We note that the boundedness of $\mathbf{c}$ will be required in our proofs. However, our algorithm will guarantee to keep $\mathbf{c}$ bounded with a proper projection onto a convex set. Here, we assume in particular $\|\mathbf{c}\| \leq 1$ for mathematical convenience in the proofs.

note that $L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})$ is defined as the average of the most recent w instant loss functions (see (6)), where the loss function $\ell_t(\boldsymbol{\theta},\boldsymbol{\mu})$ recursively depends on $\boldsymbol{\theta}$ and $\boldsymbol{\mu}$ due to $\mathbf{h}_t(\boldsymbol{\theta},\boldsymbol{\mu})$ (see (3) and Fig. 1). We emphasize that since we are interested in online learning, this recursion can be infinitely long, which might cause $L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})$ to have *unboundedly large derivatives*. On the other hand, online algorithms naturally require *loss functions with bounded gradients* to guarantee convergence [40]. Therefore, in the following, we first analyze the (potentially infinite) recursive dependency of $L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})$ on $\boldsymbol{\mu}$ and $\boldsymbol{\theta}$, where we derive sufficient conditions for its derivatives to be bounded. Then, we use the results of this analysis to find the explicit formulations of the Lipschitz and smoothness constants of $L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})$ in terms of the model parameters defining our RNN model in (1)-(2).

To analyse the effect of recursion, we first write the hidden state update in (1) with the vectorized weight matrices as

$$\mathbf{h}_{t+1} = \tanh(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) \tag{8}$$

where $\mathbf{H}_t = \mathbf{I}\otimes\mathbf{h}_t^T$, $\mathbf{X}_t = \mathbf{I}\otimes\mathbf{x}_t^T$, and $\otimes$ is the Kronecker product.

We, then, provide the following lemma, where we derive the Lipschitz and smoothness properties of the single RNN iteration defined in (8).

Lemma 1. *Let $\boldsymbol{W}$ and $\boldsymbol{U}$ the hidden layer weight matrices in the model (1)-(2), which satisfy $\|\boldsymbol{W}\|\leq\lambda$, and $\|\boldsymbol{U}\|\leq\lambda$ for some $\lambda\in\mathbb{R}$. By using the equivalent hidden state update formula in (8), the Lipschitz and smoothness properties of the single RNN iteration can be written as:*

$$\left\|\frac{\partial\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{h}_t}\right\|\leq\lambda, \tag{9}$$

$$\left\|\frac{\partial^2\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{h}_t^2}\right\|\leq 2\lambda^2, \tag{10}$$

$$\left\|\frac{\partial^2\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{h}_t\partial\boldsymbol{\theta}}\right\|\leq 2\lambda\sqrt{n_h}, \tag{11}$$

$$\left\|\frac{\partial^2\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{h}_t\partial\boldsymbol{\mu}}\right\|\leq 2\lambda\sqrt{n_x}, \tag{12}$$

$$\left\|\frac{\partial^2\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{\theta}\partial\boldsymbol{\mu}}\right\|\leq 2\sqrt{n_x}\sqrt{n_h}. \tag{13}$$

$$\left\|\frac{\partial\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{\theta}}\right\|\leq\sqrt{n_h}, \tag{14}$$

$$\left\|\frac{\partial\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{\mu}}\right\|\leq\sqrt{n_x}, \tag{15}$$

$$\left\|\frac{\partial^2\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{\theta}^2}\right\|\leq 2n_h, \tag{16}$$

$$\left\|\frac{\partial^2\tanh(\boldsymbol{H}_t\boldsymbol{\theta}+\boldsymbol{X}_t\boldsymbol{\mu})}{\partial\boldsymbol{\mu}^2}\right\|\leq 2n_x, \tag{17}$$

Proof. See Appendix A. □

In the following lemma and remark, we derive the Lipschitz properties of $\mathbf{h}_t(\boldsymbol{\theta},\boldsymbol{\mu})$, and observe the effect of infinitely long recursion on the derivatives of $L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})$.

Lemma 2. *Let $\boldsymbol{W}$, $\boldsymbol{W}'$,$\boldsymbol{U}$, $\boldsymbol{U}'$ satisfy $\|\boldsymbol{W}\|,\|\boldsymbol{W}'\|\leq\lambda$, and $\|\boldsymbol{U}\|,\|\boldsymbol{U}'\|\leq\lambda$. By using the notation in (8), let $\boldsymbol{h}_t(\boldsymbol{\theta},\boldsymbol{\mu})$ and $\boldsymbol{h}_t(\boldsymbol{\theta}',\boldsymbol{\mu}')$ be the state vectors obtained at time t by running the model in (1)-(2) with the matrices $\boldsymbol{W}$, $\boldsymbol{U}$, and $\boldsymbol{W}'$, $\boldsymbol{U}'$ on common input sequence $\{\boldsymbol{x}_1,\boldsymbol{x}_2,\cdots,\boldsymbol{x}_{t-1}\}$, respectively. If $\boldsymbol{h}_1(\boldsymbol{\theta},\boldsymbol{\mu})=\boldsymbol{h}_1(\boldsymbol{\theta}',\boldsymbol{\mu}')$, then*

$$\|\boldsymbol{h}_t(\boldsymbol{\theta},\boldsymbol{\mu})-\boldsymbol{h}_t(\boldsymbol{\theta}',\boldsymbol{\mu}')\|\leq\sum_{i=0}^{t}\lambda^i\big(\sqrt{n_h}\|\boldsymbol{\theta}-\boldsymbol{\theta}'\|+\sqrt{n_x}\|\boldsymbol{\mu}-\boldsymbol{\mu}'\|\big). \tag{18}$$

Proof. See Appendix A. □

Remark 1. *We note that, by (18), to ensure $\boldsymbol{h}_t(\boldsymbol{\theta},\boldsymbol{\mu})$ has a bounded gradient with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\mu}$ in an infinite time horizon, λ should be in $[0,1)$, i.e., $\lambda\in[0,1)$. In this case, the right hand side of (18) becomes bounded, i.e.,*

$$\|\boldsymbol{h}_t(\boldsymbol{\theta},\boldsymbol{\mu})-\boldsymbol{h}_t(\boldsymbol{\theta}',\boldsymbol{\mu}')\|\leq\frac{\sqrt{n_h}}{1-\lambda}\|\boldsymbol{\theta}-\boldsymbol{\theta}'\|+\frac{\sqrt{n_x}}{1-\lambda}\|\boldsymbol{\mu}-\boldsymbol{\mu}'\| \tag{19}$$

for any $t\in[T]$.

Recall that $L_{t_w}(\boldsymbol{\theta},\boldsymbol{\mu})$ is dependent on $\boldsymbol{h}_t(\boldsymbol{\theta},\boldsymbol{\mu})$ due to (6) and (3). Hence, to ensure the derivatives of $L_{t_w}(\boldsymbol{\theta},\boldsymbol{\mu})$ stay bounded, we need to constrain our parameter space as $\mathcal{K}_\theta=\{vec(\boldsymbol{W}):\|\boldsymbol{W}\|\leq\lambda\}$ and $\mathcal{K}_\mu=\{vec(\boldsymbol{U}):\|\boldsymbol{U}\|\leq\lambda\}$ for some $\lambda\in[0,1)$. We note that since $\mathcal{K}_\theta$ and $\mathcal{K}_\mu$ are convex sets for any $\lambda\in[0,1)$, our constraint does not violate the setting described in the previous subsection.

Now that we have found $\lambda\in[0,1)$ is sufficient for $L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})$ to have bounded derivatives in any $t\in[T]$, in the following theorem, we provide its Lipschitz and smoothness constants.

Theorem 1. *Let $\boldsymbol{\theta}=vec(\boldsymbol{W})$ and $\boldsymbol{\mu}=vec(\boldsymbol{U})$, where $\boldsymbol{W}$ and $\boldsymbol{U}$ satisfy $\|\boldsymbol{W}\|\leq\lambda$, and $\|\boldsymbol{U}\|\leq\lambda$ for some $\lambda\in[0,1)$. Then, $L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})$ has the following Lipschitz and smoothness properties:*

$$1)\ \left\|\frac{\partial L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})}{\partial\boldsymbol{\theta}}\right\|\leq R_\theta,\ \textit{where}\ R_\theta=\frac{2n_h}{1-\lambda}. \tag{20}$$

$$2)\ \left\|\frac{\partial L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})}{\partial\boldsymbol{\mu}}\right\|\leq R_\mu,\ \textit{where}\ R_\mu=\frac{2\sqrt{n_hn_x}}{1-\lambda}. \tag{21}$$

$$3)\ \left\|\frac{\partial^2 L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})}{\partial\boldsymbol{\theta}^2}\right\|\leq\beta_\theta,\ \textit{where}\ \beta_\theta=\frac{4n_h\sqrt{n_h}}{(1-\lambda)^3}. \tag{22}$$

$$4)\ \left\|\frac{\partial^2 L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})}{\partial\boldsymbol{\mu}^2}\right\|\leq\beta_\mu,\ \textit{where}\ \beta_\mu=\frac{4n_x\sqrt{n_h}}{(1-\lambda)^3}. \tag{23}$$

$$5)\ \left\|\frac{\partial^2 L_{t,w}(\boldsymbol{\theta},\boldsymbol{\mu})}{\partial\boldsymbol{\theta}\partial\boldsymbol{\mu}}\right\|\leq\beta_{\theta\mu},\ \textit{where}\ \beta_{\theta\mu}=\frac{4n_h\sqrt{n_x}}{(1-\lambda)^3}. \tag{24}$$

Proof. See Appendix B. □

In the following section, we use these properties to derive an RNN learning algorithm with a convergence guarantee.

C. Main Algorithm

Here, we present our main algorithm, namely the Windowed Online Gradient Descent Algorithm (WOGD), shown in Algorithm 1.

In the algorithm, we take the window size $w\in[T]$ and $\lambda\in[0,1)$ as the inputs. We, then, define the parameter spaces $\mathcal{K}_\theta$, $\mathcal{K}_\mu$, and $\mathcal{K}_c$ in line 3. Here, we define $\mathcal{K}_\theta$ and $\mathcal{K}_\mu$ as given in Remark 1 to ensure that the derivatives of the loss functions are bounded. Furthermore, we define $\mathcal{K}_c$ as $\mathcal{K}_c=\{\mathbf{c}:\|\mathbf{c}\|\leq 1\}$ to satisfy our assumption of $\|\mathbf{c}\|\leq 1$.

Algorithm 1 Windowed Online Gradient Descent Algorithm (WOGD)

1: **Parameters:** $w \in [T]$, $\lambda \in [0, 1)$.
2: Initialize $\boldsymbol{\theta}_1$, $\boldsymbol{\mu}_1$, $\mathbf{c}_1$ and $\mathbf{h}_1$.
3: Let $\mathcal{K}_\theta = \{\text{vec}(\mathbf{W}) : \|\mathbf{W}\| < \lambda\}$
$\mathcal{K}_\mu = \{\text{vec}(\mathbf{U}) : \|\mathbf{U}\| \leq \lambda\}$
$\mathcal{K}_c = \{\mathbf{c} : \|\mathbf{c}\| \leq 1\}$
4: **for** $t = 1$ **to** T **do**
5: Predict $\boldsymbol{\theta}_t$, $\boldsymbol{\mu}_t$ and $\mathbf{c}_t$.
6: Receive $\mathbf{x}_t$ and generate $\hat{d}_t$.
7: Observe d_t and the cost function $\ell_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$.
8: Updates:

$$\mathbf{c}_{t+1} = \Pi_{\mathcal{K}_c}\Big[\mathbf{c}_t - \frac{1}{\sqrt{t}}\frac{\partial \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \mathbf{c}_t}\Big] \tag{25}$$

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta_\theta \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)}{\partial \boldsymbol{\theta}} \tag{26}$$

$$\boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t - \eta_\mu \frac{\partial_{\mathcal{K}_\mu, \eta_\mu} L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)}{\partial \boldsymbol{\mu}}. \tag{27}$$

9: **end for**

In the learning part, we first predict the hidden layer weight matrices in their vectorized forms, i.e., $\boldsymbol{\theta}_t$ and $\boldsymbol{\mu}_t$, and the output layer weights, i.e., $\mathbf{c}_t$ (see line 5). Then, we receive the input vector $\mathbf{x}_t$ and generate $\hat{d}_t$ by running the model in (1)-(2). We next observe ground truth value d_t and the loss function $\ell_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$ in line 7. Having observed the label, we update the weight matrices in line 8 (or in (25)-(27)). Here, we update the output layer weights $\mathbf{c}_t$ with the projected online gradient descent algorithm [37], since $\ell_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$ is convex with respect to $\mathbf{c}_t$. We update the hidden weights in (26)-(27) by using their projected partial derivatives defined with $(\mathcal{K}_\theta, \eta_\theta)$ and $(\mathcal{K}_\mu, \eta_\mu)$.

We note that since we constructed our setting for adversarial $\mathbf{c}_t$ selections, the update rule for the output layer in (25) does not contradict with our analysis. Moreover, by using [37, Theorem 1], we can prove that this update rule converges to the best possible output layer weights satisfying $\|\mathbf{c}\| \leq 1$. Therefore, in the following theorem, we provide the convergence guarantee of WOGD specifically for the hidden layer weights.

Theorem 2. *Let $\ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})$ and $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu})$ be the loss and time-smoothed loss functions defined in (3) and (6), respectively. Moreover, let β_θ and β_μ be the smoothness parameters defined in (22)-(23). Then, if WOGD is run with the parameters*

$$0 < \eta_\theta \leq \frac{1}{2\beta_\theta} \text{ and } 0 < \eta_\mu \leq \frac{1}{2\beta_\mu}, \tag{28}$$

it ensures that

$$R_w(T) \leq \frac{8\sqrt{n_h}}{\min\{\eta_\theta, \eta_\mu\}}\frac{T}{w}, \tag{29}$$

where $R_w(T)$ is the local regret defined in (7). By selecting a window size w such that $\frac{T}{w} = o(T)$, one can bound $R_w(T)$ with a sublinear bound, hence, guarantee convergence of the hidden layer weights to the locally optimum parameters.[3]

[3]We use little-o notation, i.e., $g(x) = o(f(x))$, to describe an upper-bound that cannot be tight, i.e., $\lim_{x\to\infty} g(x)/f(x) = 0$.

Proof. See Appendix B. □

Theorem 2 shows that with appropriate parameter selections, WOGD is guaranteed to converge to the locally optimum hidden layer weights without any assumption on the input/output sequences and output layer weights $\mathbf{c}_t$. Additionally, recall that by [37, Theorem 1], the output layer weights also converge to the best weights in hindsight. Therefore, we conclude that *WOGD with the learning rates in (28) guarantees to learn the locally optimum RNN parameters for any bounded input/output sequences*.

Now that we have proved the convergence guarantee of WOGD, in the following remark, we investigate the computational requirement of WOGD and comment on the window size selection for the algorithm.

Remark 2. *We note that the most expensive operation of WOGD is the computation of the projected partial derivatives in (26)-(27), which requires the computation of the partial derivatives of $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu})$ with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\mu}$, and their corresponding projections onto $\mathcal{K}_\theta$ and $\mathcal{K}_\mu$. To compute the partial derivatives, we use the Truncated Backpropagation Through Time algorithm [41], which has $O\big(hn_h(n_h + n_x)\big)$ computational complexity with a truncation length h. Since WOGD uses the partial derivatives of the last w losses, we can approximate these partial derivatives with a single back-propagation by using a truncation length $h = w + n$ for some $n \in \mathbb{N}$, which results in $O\big(wn_h(n_h + n_x)\big)$ computational requirement for computing the partial derivatives. In addition, the projection step can be performed by computing the singular value decomposition (SVD) of the weight matrices $\boldsymbol{W}$ and $\boldsymbol{U}$, and clipping their singular values with λ^2. Here, since performing SVD requires $O\big(\min\{n_h, n_x\}n_h(n_h+n_x)\big)$, the computational requirement of WOGD becomes $O\big((w + \min\{n_h, n_x\})n_h(n_h + n_x)\big)$ per time step.*

We highlight that WOGD introduces a tradeoff between the convergence speed and computational complexity. For example, one can choose a large window size w to ensure fast convergence (see (29)) by increasing the computational requirement of WOGD, or vice versa. However, as we will show in the following section, selecting $w = \lceil\sqrt{T}\rceil$ usually provides the best trade-off between performance and efficiency.

In the next remark, we discuss the effect of choosing higher learning rates than the theoretically guaranteed ones given in Theorem 2.

Remark 3. *We note that WOGD is constructed by assuming the worst-case Lipschitz constants derived in Theorem 1. On the other hand, our experiments suggest that in practice, the landscape of the objective function is generally nicer what is predicted by our theoretical development. For example, in the simulations, it is observed that the smoothness parameters of the error surface, i.e., β_θ and β_μ, is usually 500 to 1000 times smaller than the values given in (22)-(23). Therefore, it is practically possible to obtain vanishing gradient with WOGD by using much higher learning rates than theoretically guaranteed learning rates in (28). As the regret bound of WOGD is inversely proportional with the learning rates (see (29)), in*

the following, we use WOGD with the higher learning rates than suggested in Theorem 2 to obtain faster convergence.

IV. Simulations

In this section, we verify our theoretical analysis and illustrate the performance of our algorithm on different real-life datasets. In particular, we consider the regression performance for elevators [42], and pumadyn [43] datasets. To demonstrate the performance improvements of our algorithm, we compare it with the three widely used first-order optimization algorithms: Adam [19], RmsProp [22] and SGD [23].

For all the simulations, we randomly draw the initial RNN weights from a Gaussian distribution with zero mean and standard deviation of 0.1, and set the initial values of all internal state variables to 0. For a fair comparison, in each experiment, we choose the hyper-parameters such that all the compared algorithms reach their maximum performance in that setup. We run each experiment 10 times and provide the mean performances.

A. Elevators Dataset

In the first simulation, we consider the regression performance on the elevators dataset [42], which includes 10000 input/output pairs obtained from a procedure that is related to controlling an F16 aircraft, i.e., $T = 10000$. Here, our aim is to predict the scalar output that expresses the actions of the F16-aircraft. For this, we use 18-dimensional input vectors of the dataset with an additional bias dimension, i.e., $n_x = 19$. To get small loss values with relatively lower run-time, we use 16-dimensional state vectors in RNNs, i.e., $n_h = 16$. In WOGD, we use $\eta_\theta = 0.006$, $\eta_\mu = 0.0075$, $\lambda = 0.9$, $w = 100$. In Adam, RmsProp and SGD, we respectively choose the learning rates as 0.002, 0.004 and 0.004. We note that we do not use momentum in SGD.

In Fig. 2a, we demonstrate the temporal loss of the compared algorithms for the elevators dataset. Here, we observe that despite the differences in the performances at the beginning, RmsProp, Adam and SGD converge to the similar loss values. We also observe that since WOGD obtains small loss values much faster than the other algorithms, it can track the desired data sequence with lower error values throughout the simulations.

To observe how window size affects the performance, we run WOGD with six different window sizes, namely $w \in \{10, 30, 50, 100, 300, 500\}$, and plot their resulting temporal performance in Fig. 2b. Here, we see that as consistent with our regret bound in Theorem 2, the final loss values obtained by the algorithms are inversely proportional to their window sizes. Moreover, we observe that in the initial part of the experiment, the algorithms with larger window sizes tend to suffer larger losses, as averaging gradients in a large window prevents them to overfit to the small amount of data observed in the earlier stage. From the figure, we can observe that the WOGD with $w = 100$ provides comparable performance with the others at both initial and final parts of the experiments. As the computational requirement of our algorithm is linear in w (see Remark bla), selecting $w = 100$ provides a highly preferable trade-off for practical purposes.

In Figs. 2c and 2d, we compare the empirical smoothness parameters of the smoothed loss functions with the theoretical upper bounds derived in Theorem 1. To observe the behaviour of the empirical smoothness parameters without calculating the exact Hessian matrix, we use the empirical Lipschitz constants, i.e.,

$$\beta_{\theta,t}^{emp} = \frac{\left\| \frac{\partial L_{t,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1})}{\partial \boldsymbol{\theta}} - \frac{\partial L_{t,w}(\boldsymbol{\theta}_{t}, \boldsymbol{\mu}_{t})}{\partial \boldsymbol{\theta}} \right\|}{\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t\|} \tag{30}$$

$$\beta_{\mu,t}^{emp} = \frac{\left\| \frac{\partial L_{t,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1})}{\partial \boldsymbol{\mu}} - \frac{\partial L_{t,w}(\boldsymbol{\theta}_{t}, \boldsymbol{\mu}_{t})}{\partial \boldsymbol{\mu}} \right\|}{\|\boldsymbol{\mu}_{t+1} - \boldsymbol{\mu}_t\|}. \tag{31}$$

In the figures, we observe that $\beta_{\theta,t}^{emp}$ and $\beta_{\mu,t}^{emp}$ are practically much lower than the theoretical upper bounds, where the theoretical values are given in the titles of the plots. We note that the difference between the theoretical upper-bounds and empirical smoothness parameters are expected, since we derived the worst-case upper bounds by considering the saturation region of RNNs, which is rarely encountered in practice due to variations in real-world datasets. Furthermore, in Figs. 2b and 2c, we see that the learning rates we used satisfy the requirement of Theorem 2 – stated in (28).

To verify our theoretical analysis in Theorem 2, we plot the normalized regret of WOGD, i.e., $R_w(t)/t$ for $t \in [T]$, and the regret bound in (29) (scaled with 0.15) in Fig. 2d. Here, we see that as consistent with our theoretical derivation, the normalized regret vanishes. Moreover, it is bounded by the regret bound scaled with 0.15. We believe that the gap between the normalized regret and actual regret bound is due to our adversarial model for the output layer weights, which is mainly used for mathematical convenience. Since we update the output layer weights simultaneously with the hidden layer weights, deriving a tighter regret bound requires a cooperative learning model for neural network optimization, which, to the best of our knowledge, has not been studied in the literature. The construction of such model and its analysis is left as an open problem for the future studies.

B. Pumaydn Dataset

In our second simulation, we consider the pumaydn dataset [43], which includes 8000 input/output pairs obtained from the simulation of Unimation Puma 560 robotic arm, i.e., $T = 8000$. Here, our aim is to estimate the angular acceleration of the arm by using the angular position and angular velocity of the links. For this simulation, we use 8-dimensional input vectors of the dataset with an additional bias dimension, i.e., $n_x = 9$, and 16-dimensional state vectors in RNN, i.e., $n_h = 16$. In WOGD, we use $\eta_\theta = 0.04$, $\eta_\mu = 0.08$, $\lambda = 0.9$, $w = 90$. In Adam, RmsProp, and SGD, we choose the learning rates as 0.004, 0.008, and 0.01, respectively. As in the previous experiment, we do not use momentum in SGD.

In Fig. 3a, we demonstrate the temporal loss of the compared algorithms for the pumaydn dataset. Here, we observe that while all algorithms provide comparable performances, WOGD enjoys a smaller error value at the end of the

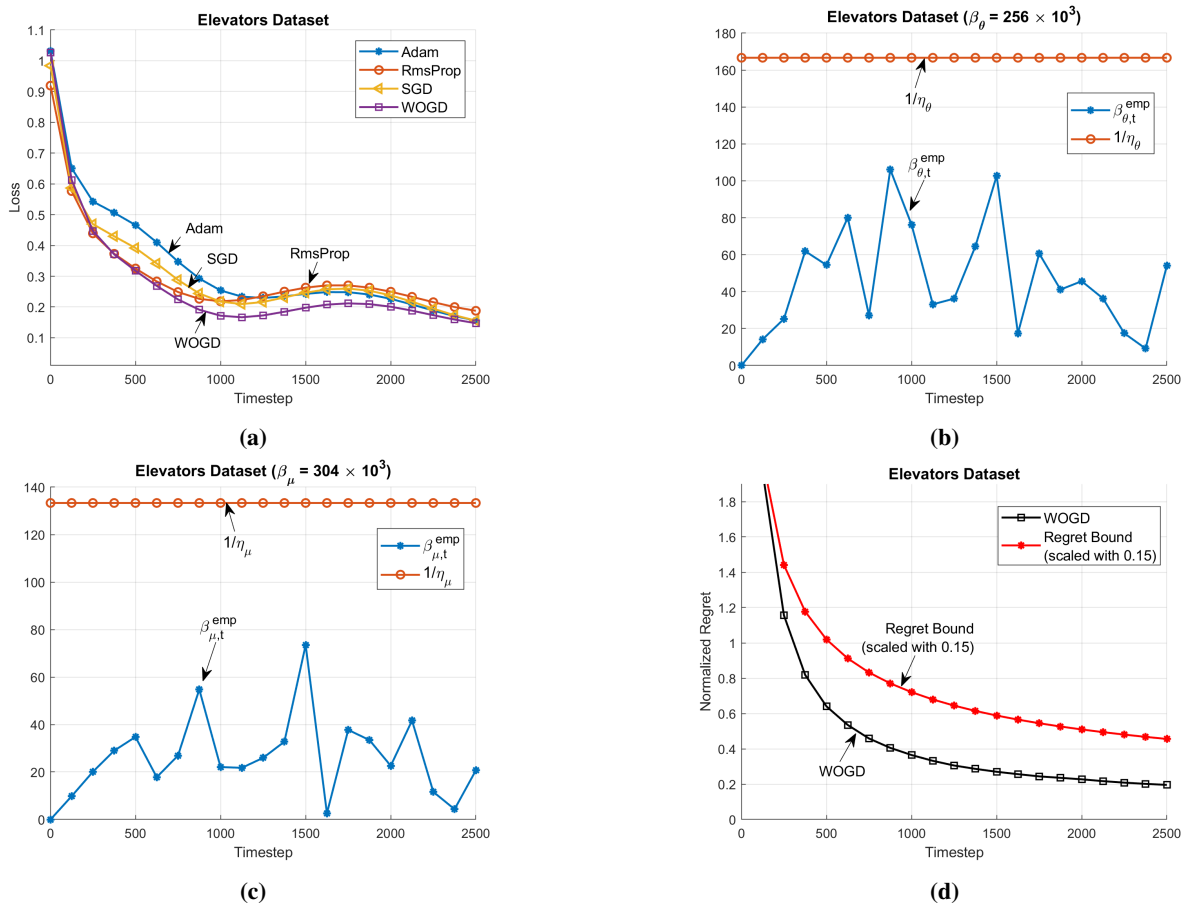

Fig. 2: **(a)** Sequential prediction performances of the algorithms for the elevators dataset **(b)-(c)** Comparison between the empirical smothness parameteres of $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu})$ and their theoretical upper bounds given in (22) and (23) **(d)** Comparison between the normalized regret bound of WOGD, i.e., $R_w(t)/t$ for $t \in [T]$, and its theoretical upper bound given in (29).

simulation. In Figs. 3b and 3c, we compare the empirical smoothness parameters (as defined in (30)-(31)) and their theoretical upper-bound derived in Theorem 1. As in the previous experiment, we observe that the empirical smoothness parameters, namely $\beta^{emp}_{\theta,t}$ and $\beta^{emp}_{\mu,t}$, are considerably smaller than the theoretical upper bounds and our learning rate selection satisfies the requirement in (28). In Fig. 3d, we plot the normalized regret of WOGD, i.e., $R_w(t)/t$ for $t \in [T]$, and the regret bound in (29) scaled with 0.01. As in the previous experiment, here also, we see that the normalized regret vanishes and the regret bound scaled with 0.01 bounds the normalized regret by above, which is parallel with our derivations in Theorem 2. Since our algorithm enjoys this guarantee by providing comparable performance with the state-of-the-art methods, it can be a theoretically grounded alternative of the widely used heuristics [19], [22], [23] in RNN-based online learning settings.

V. Conclusion

We studied online nonlinear regression with continually running RNNs, i.e., RNN-based online learning. For this problem, we introduced a first-order gradient-based optimization algorithm that provides convergence guarantee to the locally optimum parameters. We emphasize that unlike previous theoretical results, which holds in restricted settings [30], [31], our algorithm is generic such that it guarantees convergence with any smooth RNN architecture, e.g., the Elman networks [33] or LSTMs [34], without any assumption on the input/output sequences.

To achieve this result, we model the RNN-based online regression problem as a sequential learning problem, where we assumed each time step as a separate loss function assigned by an adversary. We characterize the Lipschitz properties of these loss functions with respect to the network weights and derived sufficient conditions for our model to have bounded derivatives. Then, by using these results, we introduce an online gradient descent algorithm that is guaranteed to converge to the locally optimum parameters. In the simulations, we verify our theoretical analysis. Here, we also demonstrate that our algorithm achieves considerable performance improvements with respect to the state-of-the-art gradient descent methods [19], [22], [23].

Appendix A

For the following, we denote the derivative of $\tanh$ as $\tanh'$, where $\tanh'(x) = 1 - \tanh(x)^2$. Due to the space restrictions, we denote the elementary row scaling operation with $\odot$, i.e., $\mathbf{x} \odot \mathbf{W} = \text{diag}(\mathbf{x})\mathbf{W}$, where $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{W} \in \mathbb{R}^{n \times m}$ and $\text{diag}(\mathbf{x}) \in \mathbb{R}^{n \times n}$ is the elementary scaling matrix whose

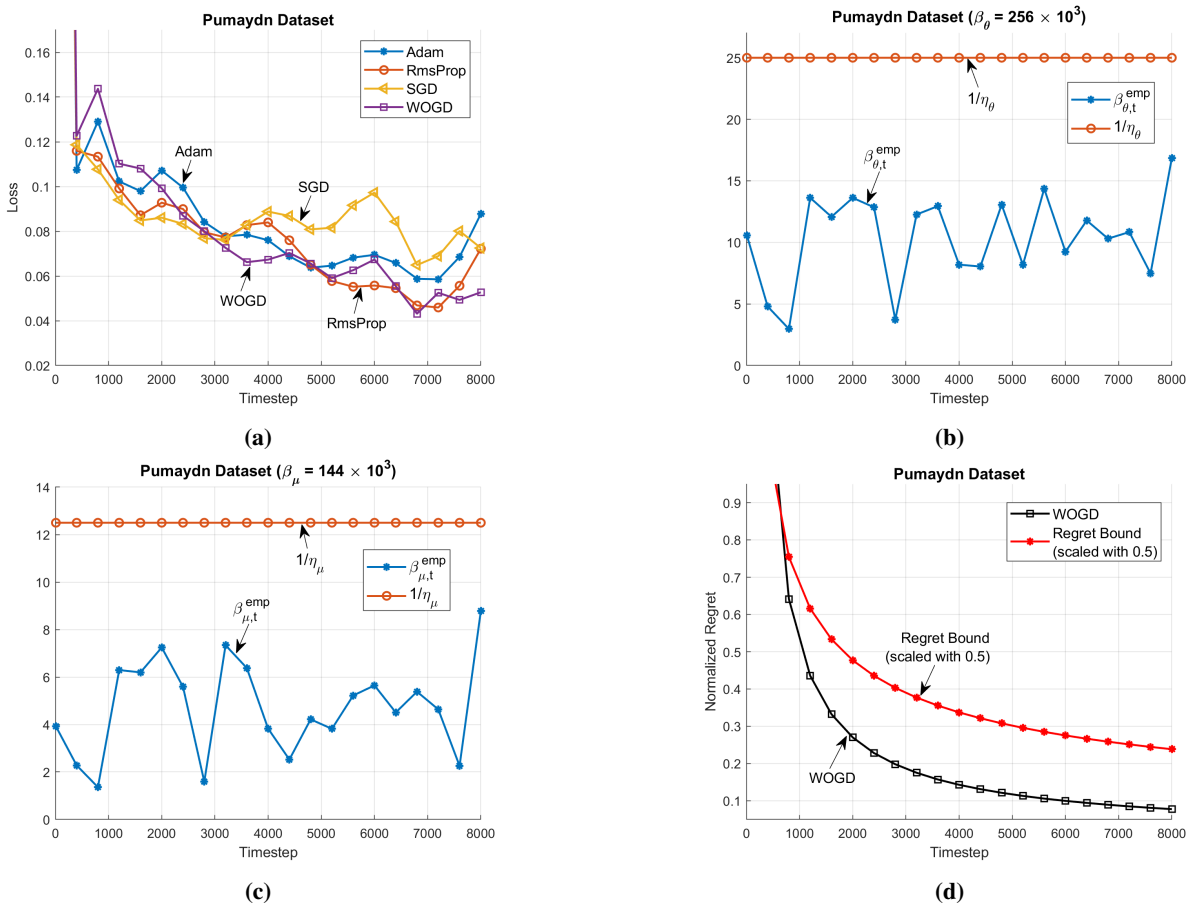

Fig. 3: **(a)** Sequential prediction performances of the algorithms for the pumaydn dataset **(b)-(c)** Comparison between the empirical smothness parameteres of $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu})$ and their theoretical upper bounds given in (22) and (23) **(d)** Comparison between the normalized regret bound of WOGD, i.e., $R_w(t)/t$ for $t \in [T]$, and its theoretical upper bound given in (29).

diagonal elements are the components of $\mathbf{x}$. Before the proofs, we give three inequalities that will be used frequently in the following:

Lemma 3. *For any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, $\boldsymbol{W} \in \mathbb{R}^{n \times m}$, where $n, m \in \mathbb{N}$, the following statements hold:*

1) $\|\boldsymbol{x} \odot \boldsymbol{W}\| \leq \|\boldsymbol{x}\|_\infty \|\boldsymbol{W}\|$ (32)
2) $\|\tanh'(\boldsymbol{x}) - \tanh'(\boldsymbol{y})\|_\infty \leq 2\|\boldsymbol{x} - \boldsymbol{y}\|$ (33)
3) $\|\boldsymbol{I} \otimes \boldsymbol{x}^T\| = \|\boldsymbol{x}\|$. (34)

Proof. 1) Since $\mathbf{x} \odot \mathbf{W} = \text{diag}(\mathbf{x})\mathbf{W}$, we have $\|\mathbf{x} \odot \mathbf{W}\| \leq \|\text{diag}(\mathbf{x})\|\|\mathbf{W}\|$, where we use the Cauchy-Schwarz inequality for bounding. Since by definition $\|\text{diag}(\mathbf{x})\| = \|\mathbf{x}\|_\infty$, $\|\mathbf{x} \odot \mathbf{W}\| \leq \|\text{diag}(\mathbf{x})\|\|\mathbf{W}\| = \|\mathbf{x}\|_\infty \|\mathbf{W}\|$.

2) Recall that $\tanh'(x) = 1 - \tanh(x)^2$. Since $\tanh(x) \in [-1, 1]$, tanh is 1-Lipschitz 2-smooth. Then, by using $\|\mathbf{x}\|_\infty \leq \|\mathbf{x}\|$ for any $\mathbf{x} \in \mathbb{R}^n$, we have $\|\tanh'(\mathbf{x}) - \tanh'(\mathbf{y})\|_\infty \leq \|\tanh'(\mathbf{x}) - \tanh'(\mathbf{y})\|$. Since tanh is 2-smooth, we have $\|\tanh'(\mathbf{x}) - \tanh'(\mathbf{y})\| \leq 2\|\mathbf{x} - \mathbf{y}\|$.

3) See [44, Theorem 8].

□

Proof of Lemma 1. In the following, we prove each statement separately:

1) *Proof of (9)*: We note that (1) and (8) are equivalent. By using (32) and $\tanh'(x) \leq 1$ on (1), we write

$$\left\| \frac{\partial \tanh(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t)}{\partial \mathbf{h}_t} \right\| = \|\tanh'(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t) \odot \mathbf{W}\| \leq \|\tanh'(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t)\|_\infty \|\mathbf{W}\| \leq \lambda.$$

2) *Proof of (10)*: By using (32) and (33), we write

$$\begin{aligned} &\left\| \frac{\partial \tanh(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t)}{\partial \mathbf{h}_t} - \frac{\partial \tanh(\mathbf{W}'\mathbf{h}_t + \mathbf{U}\mathbf{x}_t)}{\partial \mathbf{h}_t} \right\| \\ &= \|\tanh'(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t) \odot \mathbf{W} - \tanh'(\mathbf{W}\mathbf{h}'_t + \mathbf{U}\mathbf{x}_t) \odot \mathbf{W}\| \\ &\leq \|\tanh'(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t) - \tanh'(\mathbf{W}\mathbf{h}'_t + \mathbf{U}\mathbf{x}_t)\|_\infty \|\mathbf{W}\| \\ &\leq 2\|\mathbf{W}\|\|\mathbf{h}_t - \mathbf{h}'_t\|\|\mathbf{W}\| \leq 2\lambda^2 \|\mathbf{h}_t - \mathbf{h}'_t\|. \end{aligned}$$

3) *Proof of (11)*: We note that since tanh is twice differentiable, the order of partial derivatives is not important. Then,

by using (32), (33), (34), and $\|\mathbf{h}_t\| \leq \sqrt{n_h}$, we write

$$\begin{aligned}
&\Big\|\frac{\partial \tanh(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu})}{\partial \boldsymbol{\theta}} - \frac{\partial \tanh(\mathbf{H}_t'\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu})}{\partial \boldsymbol{\theta}}\Big\| \\
&\quad = \|\tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) \odot \mathbf{H}_t - \tanh'(\mathbf{H}_t'\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) \odot \mathbf{H}_t'\| && (35)\\
&\quad \leq \|\tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) - \tanh'(\mathbf{H}_t'\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu})\|_\infty \|\mathbf{H}_t\| \\
&\qquad + \|\tanh'(\mathbf{H}_t'\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu})\|_\infty \|\mathbf{H}_t - \mathbf{H}_t'\| && (36)\\
&\quad \leq \|\tanh'(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t) - \tanh'(\mathbf{W}\mathbf{h}_t' + \mathbf{U}\mathbf{x}_t)\|_\infty \|\mathbf{H}_t\| \\
&\qquad + \|\mathbf{h}_t - \mathbf{h}_t'\| && (37)\\
&\quad \leq \big(2\|\mathbf{W}\|\sqrt{n_h} + 1\big)\|\mathbf{h}_t - \mathbf{h}_t'\| \leq (2\lambda\sqrt{n_h} + 1)\|\mathbf{h}_t - \mathbf{h}_t'\|,
\end{aligned}$$

where we add $\pm \tanh'(\mathbf{H}_t'\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) \odot \mathbf{H}_t$ inside of the norm in (35), and use the triangle inequality for (36). Here, we omit $+1$ term for mathematical convenience in the following derivations.

4) *Proof of (12)*: This can be obtained by repeating the steps in the proof of (11) for $\boldsymbol{\mu}$ and $\mathbf{h}_t$.

5) *Proof of (13)*: By using (32), (33), (34), $\|\mathbf{h}_t\| \leq \sqrt{n_h}$ and $\|\mathbf{x}_t\| \leq \sqrt{n_x}$, we write

$$\begin{aligned}
&\Big\|\frac{\partial \tanh(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu})}{\partial \boldsymbol{\theta}} - \frac{\partial \tanh(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}')}{\partial \boldsymbol{\theta}}\Big\| \\
&\quad = \|\tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) \odot \mathbf{H}_t - \tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}') \odot \mathbf{H}_t\| \\
&\quad \leq \|\tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) - \tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}')\|_\infty \|\mathbf{H}_t\| \\
&\quad \leq 2\|\mathbf{X}_t\|\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|\|\mathbf{H}_t\| \leq 2\sqrt{n_h}\sqrt{n_x}\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|.
\end{aligned}$$

6) *Proof of (14)*: By using (32), $\tanh'(x) \leq 1$, and $\|\mathbf{h}_t\| \leq \sqrt{n_h}$,

$$\begin{aligned}
\Big\|\frac{\partial \tanh(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu})}{\partial \boldsymbol{\theta}}\Big\| &= \|\tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) \odot \mathbf{H}_t\| \\
&\leq \|\tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu})\|_\infty \|\mathbf{H}_t\| \leq \sqrt{n_h}.
\end{aligned}$$

7) *Proof of (15)*: This can be obtained by repeating the steps in the proof of (14) for $\boldsymbol{\mu}$.

8) *Proof of (16)*: By using (32), (33), (34), and $\|\mathbf{h}_t\| \leq \sqrt{n_h}$, we write

$$\begin{aligned}
&\Big\|\frac{\partial \tanh(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu})}{\partial \boldsymbol{\theta}} - \frac{\partial \tanh(\mathbf{H}_t\boldsymbol{\theta}' + \mathbf{X}_t\boldsymbol{\mu})}{\partial \boldsymbol{\theta}}\Big\| \\
&\quad = \|\tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) \odot \mathbf{H}_t - \tanh'(\mathbf{H}_t\boldsymbol{\theta}' + \mathbf{X}_t\boldsymbol{\mu}) \odot \mathbf{H}_t\| \\
&\quad \leq \|\tanh'(\mathbf{H}_t\boldsymbol{\theta} + \mathbf{X}_t\boldsymbol{\mu}) - \tanh'(\mathbf{H}_t\boldsymbol{\theta}' + \mathbf{X}_t\boldsymbol{\mu})\|_\infty \|\mathbf{H}_t\| \\
&\quad \leq 2\|\mathbf{H}_t\|\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|\|\mathbf{H}_t\| \leq 2n_h\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|.
\end{aligned}$$

9) *Proof of (17)*: This can be obtained by repeating the steps in the proof of (16) for $\boldsymbol{\mu}$.

□

Proof of Lemma 2. Before the proof, let $\mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu})$ be the state vector obtained at time t by running the model in (1) with the matrices $\mathbf{W}'$, $\mathbf{U}$, input sequence $\{\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_{t-1}\}$, and the initial condition $\mathbf{h}_1(\boldsymbol{\theta}', \boldsymbol{\mu}) = \mathbf{h}_1(\boldsymbol{\theta}, \boldsymbol{\mu}) = \mathbf{h}_1(\boldsymbol{\theta}', \boldsymbol{\mu}')$. Then,

$$\begin{aligned}
&\|\mathbf{h}_t(\boldsymbol{\theta}, \boldsymbol{\mu}) - \mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu}')\| && (38)\\
&\quad \leq \|\mathbf{h}_t(\boldsymbol{\theta}, \boldsymbol{\mu}) - \mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu})\| + \|\mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu}) - \mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu}')\|, && (39)
\end{aligned}$$

where we add $\pm\mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu})$ inside of the norm in (38), and use the triangle inequality. We will bound the terms in (39) separately. We begin with the first term. Since $\mathbf{h}_t(\boldsymbol{\theta}, \boldsymbol{\mu})$ and $\mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu})$ share the same $\boldsymbol{\mu}$, in the following (between (40)-(44)), we abbreviate them as $\mathbf{h}_t(\boldsymbol{\theta})$ and $\mathbf{h}_t(\boldsymbol{\theta}')$:

$$\begin{aligned}
&\|\mathbf{h}_t(\boldsymbol{\theta}) - \mathbf{h}_t(\boldsymbol{\theta}')\| = && (40)\\
&= \|\tanh(\mathbf{W}\mathbf{h}_{t-1}(\boldsymbol{\theta}) + \mathbf{U}\mathbf{x}_t) - \tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\theta}') + \mathbf{U}\mathbf{x}_t)\| && (41)\\
&\leq \|\tanh(\mathbf{W}\mathbf{h}_{t-1}(\boldsymbol{\theta}) + \mathbf{U}\mathbf{x}_t) - \tanh(\mathbf{W}\mathbf{h}_{t-1}(\boldsymbol{\theta}') + \mathbf{U}\mathbf{x}_t)\| \\
&\quad + \|\tanh(\mathbf{W}\mathbf{h}_{t-1}(\boldsymbol{\theta}') + \mathbf{U}\mathbf{x}_t) - \tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\theta}') + \mathbf{U}\mathbf{x}_t)\| && (42)\\
&\leq \lambda\|\mathbf{h}_{t-1}(\boldsymbol{\theta}) - \mathbf{h}_{t-1}(\boldsymbol{\theta}')\| + \sqrt{n_h}\|\boldsymbol{\theta} - \boldsymbol{\theta}'\| && (43)\\
&\leq \sum_{i=0}^{t-2}\Big(\lambda^i\sqrt{n_h}\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|\Big) && (44)\\
& && (45)
\end{aligned}$$

Here, to obtain (42), we add $\pm\tanh(\mathbf{W}\mathbf{h}_{t-1}(\boldsymbol{\theta}') + \mathbf{U}\mathbf{x}_t)$ inside of the norm in (41), and use the triangle inequality. Then, we use (9) and (14) to get (43). Until we reach (44), we repeatedly apply the same bounding technique to bound the norm of the differences between the state vectors.

Now, we bound the second term in (39). Since $\mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu})$ and $\mathbf{h}_t(\boldsymbol{\theta}', \boldsymbol{\mu}')$ share the same $\boldsymbol{\mu}'$, in the following (between (46)-(50)), we abbreviate them as $\mathbf{h}_t(\boldsymbol{\mu})$ and $\mathbf{h}_t(\boldsymbol{\mu}')$:

$$\begin{aligned}
&\|\mathbf{h}_t(\boldsymbol{\mu}) - \mathbf{h}_t(\boldsymbol{\mu}')\| = && (46)\\
&= \|\tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\mu}) + \mathbf{U}\mathbf{x}_t) - \tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\mu}') + \mathbf{U}'\mathbf{x}_t)\| && (47)\\
&\leq \|\tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\mu}) + \mathbf{U}\mathbf{x}_t) - \tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\mu}') + \mathbf{U}\mathbf{x}_t)\| \\
&\quad + \|\tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\mu}') + \mathbf{U}\mathbf{x}_t) - \tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\mu}') + \mathbf{U}'\mathbf{x}_t)\| && (48)\\
&\leq \lambda\|\mathbf{h}_{t-1}(\boldsymbol{\mu}) - \mathbf{h}_{t-1}(\boldsymbol{\mu}')\| + \sqrt{n_x}\|\boldsymbol{\mu} - \boldsymbol{\mu}'\| && (49)\\
&\leq \sum_{i=0}^{t-2}\Big(\lambda^i\sqrt{n_x}\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|\Big) && (50)\\
& && (51)
\end{aligned}$$

Here, for (48), we add $\pm\tanh(\mathbf{W}'\mathbf{h}_{t-1}(\boldsymbol{\mu}') + \mathbf{U}\mathbf{x}_t)$ and use the triangle inequality. We, then, use (9) and (15) to get (49). Until we reach (50), we repeatedly apply the same technique to bound the norm of the differences between state vectors. In the end, we use (44) and (50) to bound (39), which yields the statement in the lemma. □

APPENDIX B

Proof of Theorem 1. Recall that $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu}) = \frac{1}{w}\sum_{i=0}^{w-1}\ell_{t-i}(\boldsymbol{\theta}, \boldsymbol{\mu})$. Hence, if we can bound the derivative of $\ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})$ for an arbitrary $t \in [T]$, the resulting bound will be valid for $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu})$. Therefore, in the following, we analyse the Lipschitz properties of $\ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})$ for an arbitrary $t \in [T]$ and extend the result for $L_{t,w}(\boldsymbol{\theta}, \boldsymbol{\mu})$. In addition, in the following, we note that since $d_t, \hat{d}_t \in [-\sqrt{n_h}, \sqrt{n_h}]$ and $\|\mathbf{c}_t\| \leq 1$, the ℓ_2 norm of $-(d_t - \hat{d}_t)\mathbf{c}$ is bounded by $2\sqrt{n_h}$, i.e., $\|-(d_t - \hat{d}_t)\mathbf{c}\| \leq 2\sqrt{n_h}$.

1) We write

$$\Big\|\frac{\partial \ell_t(\boldsymbol{\theta},\boldsymbol{\mu})}{\partial \boldsymbol{\theta}}\Big\| = \Big\|-(d_t-\hat{d}_t)\frac{\partial \hat{d}_t}{\partial \boldsymbol{\theta}}\Big\| \tag{52}$$

$$= \Big\|-(d_t-\hat{d}_t)\frac{\partial \hat{d}_t}{\partial \mathbf{h}_t}\Big(\sum_{\tau=1}^{t}\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}}\Big)\Big\|$$

$$\leq \Big\|-(d_t-\hat{d}_t)\mathbf{c}\Big\|\Big\|\sum_{\tau=1}^{t}\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}}\Big\|$$

$$\leq 2\sqrt{n_h}\Big\|\sum_{\tau=1}^{t}\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}}\Big\|$$

$$= 2\sqrt{n_h}\Big\|\sum_{\tau=1}^{t}\Big(\prod_{i=\tau+1}^{t}\frac{\partial \mathbf{h}_i}{\partial \mathbf{h}_{i-1}}\Big)\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}}\Big\|$$

$$\leq 2\sqrt{n_h}\sum_{\tau=1}^{t}\Big\|\prod_{i=\tau+1}^{t}\frac{\partial \mathbf{h}_i}{\partial \mathbf{h}_{i-1}}\Big\|\Big\|\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}}\Big\|$$

$$\leq 2\sqrt{n_h}\sqrt{n_h}\sum_{\tau=1}^{t}\lambda^{t-\tau} \tag{53}$$

$$\leq \frac{2n_h}{1-\lambda}, \tag{54}$$

where we use (9) and (14) to get (53). By realizing that (54) holds for an arbitrary t, the statement in the theorem can be obtained.

2) By using similar steps in (52)-(54), we write

$$\Big\|\frac{\partial \ell_t(\boldsymbol{\theta},\boldsymbol{\mu})}{\partial \boldsymbol{\mu}}\Big\| = \Big\|-(d_t-\hat{d}_t)\frac{\partial \hat{d}_t}{\partial \boldsymbol{\mu}}\Big\|$$

$$= \Big\|-(d_t-\hat{d}_t)\frac{\partial \hat{d}_t}{\partial \mathbf{h}_t}\Big(\sum_{\tau=1}^{t}\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\mu}}\Big)\Big\|$$

$$\leq \Big\|-(d_t-\hat{d}_t)\mathbf{c}\Big\|\Big\|\sum_{\tau=1}^{t}\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\mu}}\Big\|$$

$$\leq 2\sqrt{n_h}\sqrt{n_x}\sum_{\tau=1}^{t}\lambda^{t-\tau} \tag{55}$$

$$\leq \frac{2\sqrt{n_h n_x}}{1-\lambda}, \tag{56}$$

where we use (9) and (15) to get (55). By realizing that (56) holds for an arbitrary t, the statement in the theorem can be obtained.

3) Let us use $\mathbf{h}_t$ and $\mathbf{h}'_t$ for the state vectors obtained by running the model in (1) from the initial step up to current time step t with the same initial condition, same input layer matrix $\boldsymbol{\mu}$, common input sequence $\{\mathbf{x}_1,\cdots,\mathbf{x}_{t-1}\}$ but different $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$, respectively. Let us also say $\ell_t(\boldsymbol{\theta}',\boldsymbol{\mu}) = 0.5(d_t-\hat{d}'_t)^2$, where $\hat{d}'_t$ is the prediction of the second model producing $\mathbf{h}'_t$. Then,

$$\Big\|\frac{\partial \ell_t(\boldsymbol{\theta},\boldsymbol{\mu})}{\partial \boldsymbol{\theta}} - \frac{\partial \ell_t(\boldsymbol{\theta}',\boldsymbol{\mu})}{\partial \boldsymbol{\theta}}\Big\|$$

$$= \Big\|(d_t-\hat{d}'_t)\mathbf{c}^T\Big(\sum_{\tau=1}^{t}\frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau}\frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\theta}}\Big) - (d_t-\hat{d}_t)\mathbf{c}^T\Big(\sum_{\tau=1}^{t}\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}}\Big)\Big\| \tag{57}$$

$$\leq 2\sqrt{n_h}\sum_{\tau=1}^{t}\Big\|\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau}\frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\theta}}\Big\| \tag{58}$$

$$\leq 2\sqrt{n_h}\sum_{\tau=1}^{t}\Big(\Big\|\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\Big\|\Big\|\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}} - \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\theta}}\Big\| + \Big\|\frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\theta}}\Big\|\Big\|\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau}\Big\|\Big) \tag{59}$$

$$\leq 2\sqrt{n_h}\sum_{\tau=1}^{t}\lambda^{t-\tau}\Big\|\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}} - \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\theta}}\Big\| + 2n_h\sum_{\tau=1}^{t}\Big\|\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau}\Big\|. \tag{60}$$

Here, to get (59) from (58), we add $\pm\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau}\frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\theta}}$ inside the norm, and use the triangle inequality. To get (60), we use (9) and (14).

In the following, we will bound the terms in (60) separately. We begin with the first term. Note that $\mathbf{h}_\tau = \tanh(\mathbf{W}\mathbf{h}_{\tau-1} + \mathbf{U}\mathbf{x}_{\tau-1})$, and $\mathbf{h}'_\tau = \tanh(\mathbf{W}'\mathbf{h}'_{\tau-1} + \mathbf{U}\mathbf{x}_{\tau-1})$. Then,

$$2\sqrt{n_h}\sum_{\tau=1}^{t}\lambda^{t-\tau}\Big\|\frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\theta}} - \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\theta}}\Big\| \tag{61}$$

$$\leq 2\sqrt{n_h}\sum_{\tau=1}^{t}\lambda^{t-\tau}\big(2\lambda\sqrt{n_h}\|\mathbf{h}_{\tau-1}-\mathbf{h}'_{\tau-1}\| + 2n_h\|\boldsymbol{\theta}-\boldsymbol{\theta}'\|\big) \tag{62}$$

$$\leq 2\sqrt{n_h}\Big(\sum_{\tau=1}^{t}\lambda^{t-\tau}\Big)\Big(\frac{2\lambda n_h}{1-\lambda} + 2n_h\Big)\|\boldsymbol{\theta}-\boldsymbol{\theta}'\| \tag{63}$$

$$\leq \frac{4n_h\sqrt{n_h}}{1-\lambda}\Big(\frac{\lambda}{1-\lambda}+1\Big)\|\boldsymbol{\theta}-\boldsymbol{\theta}'\|, \tag{64}$$

where we add $\pm\frac{\partial \tanh(\mathbf{W}\mathbf{h}'_{\tau-1}+\mathbf{U}\mathbf{x}_{\tau-1})}{\partial \boldsymbol{\theta}}$, and use the triangle inequality, (11) and (16) for (62). We use Lemma 2 for (63). Now, we bound the second term in (60). To bound the term, we first focus on the term inside of the sum, i.e., $\Big\|\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau}\Big\|$:

$$\Big\|\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau}\Big\| \leq \Big\|\frac{\partial \mathbf{h}_{t-1}}{\partial \mathbf{h}_\tau}\Big\|\Big\|\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_{t-1}} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_{t-1}}\Big\| + \Big\|\frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_{t-1}}\Big\|\Big\|\frac{\partial \mathbf{h}_{t-1}}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_{t-1}}{\partial \mathbf{h}_\tau}\Big\| \tag{65}$$

$$\leq \lambda^{t-\tau-1}\big(2\lambda^2\|\mathbf{h}_{t-1}-\mathbf{h}'_{t-1}\| + 2\lambda\sqrt{n_h}\|\boldsymbol{\theta}-\boldsymbol{\theta}'\|\big) + \lambda\Big\|\frac{\partial \mathbf{h}_{t-1}}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_{t-1}}{\partial \mathbf{h}_\tau}\Big\| \tag{66}$$

$$\leq \lambda^{t-\tau-1}\Big(\frac{2\lambda^2\sqrt{n_h}}{1-\lambda} + 2\lambda\sqrt{n_h}\Big)\|\boldsymbol{\theta}-\boldsymbol{\theta}'\| + \lambda\Big\|\frac{\partial \mathbf{h}_{t-1}}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_{t-1}}{\partial \mathbf{h}_\tau}\Big\| \tag{67}$$

$$\leq (t-\tau)\lambda^{t-\tau-1}\Big(\frac{2\lambda^2\sqrt{n_h}}{1-\lambda} + 2\lambda\sqrt{n_h}\Big)\|\boldsymbol{\theta}-\boldsymbol{\theta}'\|, \tag{68}$$

where we add $\pm\frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_{t-1}}\frac{\partial \mathbf{h}_{t-1}}{\partial \mathbf{h}_\tau}$, and use the triangle inequality for (65), utilize (10) and (11) for (66), and Lemma 2 for (67). We, then, repeat the same manipulations in (65)-(67) to bound the

terms with partial derivatives. Then, the second term in (60) can be bound as:

$$2n_h \sum_{\tau=1}^{t} \Big\| \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau} \Big\| \tag{69}$$

$$\leq 2n_h \sum_{\tau=1}^{t} (t-\tau)\lambda^{t-\tau-1} \Big(\frac{2\lambda^2\sqrt{n_h}}{1-\lambda} + 2\lambda\sqrt{n_h} \Big) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| \tag{70}$$

$$= \frac{4n_h\sqrt{n_h}}{1-\lambda} \Big(\frac{\lambda^2}{(1-\lambda)^2} + \frac{\lambda}{1-\lambda} \Big) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|, \tag{71}$$

where we use the upper bound of the series $\sum_{\tau=1}^{t}(t-\tau)\lambda^{t-\tau-1}$, i.e., $1/(1-\lambda)^2$, to get (71). Then, by using (64) and (71), we can bound (60) as

$$\Big\| \frac{\partial \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\theta}} - \frac{\partial \ell_t(\boldsymbol{\theta}', \boldsymbol{\mu})}{\partial \boldsymbol{\theta}} \Big\| \tag{72}$$

$$\leq \frac{4n_h\sqrt{n_h}}{1-\lambda} \Big(\frac{\lambda^2}{(1-\lambda)^2} + \frac{2\lambda}{1-\lambda} + 1 \Big) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| \tag{73}$$

$$= \frac{4n_h\sqrt{n_h}}{(1-\lambda)^3} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|. \tag{74}$$

By realizing that (74) holds for an arbitrary t, the statement in the theorem can be obtained.

4) This part can be obtained by adapting the steps in the previous proof for $\boldsymbol{\mu}$, and use the Lipschitz conditions in (12), (15), (17) and Lemma 2 accordingly.

5) We use the same notation in the proof of the 3rd statement. Then,

$$\Big\| \frac{\partial \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\mu}} - \frac{\partial \ell_t(\boldsymbol{\theta}', \boldsymbol{\mu})}{\partial \boldsymbol{\mu}} \Big\|$$

$$= \Big\| (d_t - \hat{d}'_t)\mathbf{c}^T \Big(\sum_{\tau=1}^{t} \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau} \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\mu}} \Big) - (d_t - \hat{d}_t)\mathbf{c}^T \Big(\sum_{\tau=1}^{t} \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} \frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\mu}} \Big) \Big\| \tag{75}$$

$$\leq 2\sqrt{n_h} \sum_{\tau=1}^{t} \Big\| \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} \frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\mu}} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau} \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\mu}} \Big\| \tag{76}$$

$$\leq 2\sqrt{n_h} \sum_{\tau=1}^{t} \Big(\Big\| \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} \Big\| \Big\| \frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\mu}} - \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\mu}} \Big\| + \Big\| \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\mu}} \Big\| \Big\| \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau} \Big\| \Big) \tag{77}$$

$$\leq 2\sqrt{n_h} \sum_{\tau=1}^{t} \lambda^{t-\tau} \Big\| \frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\mu}} - \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\mu}} \Big\| + 2\sqrt{n_h n_x} \sum_{\tau=1}^{t} \Big\| \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau} \Big\|. \tag{78}$$

Here, to get (77) from (76), we add $\pm \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\mu}}$ inside the norm, and use the triangle inequality. To get (78), we use (9) and (15).

We bound the terms in (80) separately. We begin with the first term. Note that $\mathbf{h}_\tau = \tanh(\mathbf{W}\mathbf{h}_{\tau-1} + \mathbf{U}\mathbf{x}_{\tau-1})$, and $\mathbf{h}'_\tau = \tanh(\mathbf{W}'\mathbf{h}'_{\tau-1} + \mathbf{U}\mathbf{x}_{\tau-1})$. Then,

$$2\sqrt{n_h} \sum_{\tau=1}^{t} \lambda^{t-\tau} \Big\| \frac{\partial \mathbf{h}_\tau}{\partial \boldsymbol{\mu}} - \frac{\partial \mathbf{h}'_\tau}{\partial \boldsymbol{\mu}} \Big\| \tag{79}$$

$$\leq 2\sqrt{n_h} \sum_{\tau=1}^{t} \lambda^{t-\tau} \big(2\lambda\sqrt{n_x} \|\mathbf{h}_{\tau-1} - \mathbf{h}'_{\tau-1}\| + 2\sqrt{n_x n_h} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| \big) \tag{80}$$

$$\leq 2\sqrt{n_h} \Big(\sum_{\tau=1}^{t} \lambda^{t-\tau} \Big) \Big(\frac{2\lambda\sqrt{n_x n_h}}{1-\lambda} + 2\sqrt{n_x n_h} \Big) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| \tag{81}$$

$$\leq \frac{4n_h\sqrt{n_x}}{1-\lambda} \Big(\frac{\lambda}{1-\lambda} + 1 \Big) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|, \tag{82}$$

where we add $\pm \frac{\partial \tanh(\mathbf{W}\mathbf{h}'_{\tau-1} + \mathbf{U}\mathbf{x}_{\tau-1})}{\partial \boldsymbol{\theta}}$, and use the triangle inequality, (12) and (13) for (80). We use Lemma 2 for (81).

Now, we bound the second term in (78):

$$2\sqrt{n_h n_x} \sum_{\tau=1}^{t} \Big\| \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau} \Big\| \tag{83}$$

$$\leq 2\sqrt{n_h n_x} \sum_{\tau=1}^{t} (t-\tau)\lambda^{t-\tau-1} \Big(\frac{2\lambda^2\sqrt{n_h}}{1-\lambda} + 2\lambda\sqrt{n_h} \Big) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| \tag{84}$$

$$= \frac{4n_h\sqrt{n_x}}{1-\lambda} \Big(\frac{\lambda^2}{(1-\lambda)^2} + \frac{\lambda}{1-\lambda} \Big) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|, \tag{85}$$

where we use (68) to bound the terms $\Big\| \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_\tau} - \frac{\partial \mathbf{h}'_t}{\partial \mathbf{h}'_\tau} \Big\|$. Then, by using (82) and (85), we bound (78) as follows:

$$\Big\| \frac{\partial \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\mu}} - \frac{\partial \ell_t(\boldsymbol{\theta}', \boldsymbol{\mu})}{\partial \boldsymbol{\mu}} \Big\| \tag{86}$$

$$\leq \frac{4n_h\sqrt{n_x}}{1-\lambda} \Big(\frac{\lambda^2}{(1-\lambda)^2} + \frac{2\lambda}{1-\lambda} + 1 \Big) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| \tag{87}$$

$$= \frac{4n_h\sqrt{n_x}}{(1-\lambda)^3} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|. \tag{88}$$

By realizing that (88) holds for an arbitrary t, the statement in the theorem can be obtained.

□

Proof of Theorem 2. In the following, we use $\langle \cdot, \cdot \rangle$ to denote the inner product. Due to the space constraints, we omit the arguments in the partial derivative terms, i.e.,

$$\frac{\partial L_{t,w}}{\partial \boldsymbol{\theta}} := \frac{\partial \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\theta}}, \qquad \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} := \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)}{\partial \boldsymbol{\theta}},$$
$$\frac{\partial L_{t,w}}{\partial \boldsymbol{\mu}} := \frac{\partial \ell_t(\boldsymbol{\theta}, \boldsymbol{\mu})}{\partial \boldsymbol{\mu}}, \qquad \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} := \frac{\partial_{\mathcal{K}_\mu, \eta_\mu} L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)}{\partial \boldsymbol{\mu}}.$$

We start our proof by bounding the term $L_{t,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1}) -$

$L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$:

$$
\begin{aligned}
&L_{t,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1}) - L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t) \\
&\leq \Big\langle \frac{\partial L_{t,w}}{\partial \boldsymbol{\theta}}, \boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t \Big\rangle + \frac{\beta_\theta}{2} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t\|^2 \\
&\quad \Big\langle \frac{\partial L_{t,w}}{\partial \boldsymbol{\mu}}, \boldsymbol{\mu}_{t+1} - \boldsymbol{\mu}_t \Big\rangle + \frac{\beta_\mu}{2} \|\boldsymbol{\mu}_{t+1} - \boldsymbol{\mu}_t\|^2 \\
&\quad + \beta_{\theta\mu} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t\| \|\boldsymbol{\mu}_{t+1} - \boldsymbol{\mu}_t\| && (89) \\
&= -\eta_\theta \Big\langle \frac{\partial L_{t,w}}{\partial \boldsymbol{\theta}}, \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} \Big\rangle - \eta_\mu \Big\langle \frac{\partial L_{t,w}}{\partial \boldsymbol{\mu}}, \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} \Big\rangle \\
&\quad + \frac{\beta_\theta \eta_\theta^2}{2} \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} \Big\|^2 + \frac{\beta_\mu \eta_\mu^2}{2} \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} \Big\|^2 \\
&\quad + \beta_{\theta\mu} \eta_\theta \eta_\mu \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} \Big\| \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} \Big\| && (90) \\
&= -\eta_\theta \Big\langle \frac{\partial L_{t,w}}{\partial \boldsymbol{\theta}}, \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} \Big\rangle - \eta_\mu \Big\langle \frac{\partial L_{t,w}}{\partial \boldsymbol{\mu}}, \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} \Big\rangle \\
&\quad + 0.5 \Big(\eta_\theta \sqrt{\beta_\theta} \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} \Big\| + \eta_\mu \sqrt{\beta_\mu} \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} \Big\| \Big)^2 && (91) \\
&\leq -\big(\eta_\theta - \eta_\theta^2 \beta_\theta\big) \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} \Big\|^2 \\
&\quad - \big(\eta_\mu - \eta_\mu^2 \beta_\mu\big) \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} \Big\|^2, && (92)
\end{aligned}
$$

where we use [45, Lemma 3.4] for (89), the update rules in (26)-(27) for (90), and the fact that $\beta_{\theta\mu} = \sqrt{\beta_\theta \beta_\mu}$ for (91) (see (22)-(24)). Moreover, for (92), we use [39, Lemma 3.2] and the fact $2(a^2 + b^2) \leq (a + b)^2$ for any $a, b \in \mathbb{R}$.

For the following, we define $l_t(\boldsymbol{\theta}, \boldsymbol{\mu}) = 0$ for $t \leq 0$. We continue our proof by bounding $L_{T+1,w}(\boldsymbol{\theta}_{T+1}, \boldsymbol{\mu}_{T+1})$ as follows:

$$
\begin{aligned}
&L_{T+1,w}(\boldsymbol{\theta}_{T+1}, \boldsymbol{\mu}_{T+1}) \\
&\quad = \sum_{t=0}^{T} L_{t+1,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1}) - L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t) && (93) \\
&\quad = \sum_{t=1}^{T} \Big(L_{t,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1}) - L_{t,w}(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t) \Big) \\
&\qquad + \sum_{t=0}^{T} \Big(L_{t+1,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1}) - L_{t,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1}) \Big) && (94) \\
&\quad \leq -\big(\eta_\theta - \eta_\theta^2 \beta_\theta\big) \sum_{t=1}^{T} \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} \Big\|^2 \\
&\qquad - \big(\eta_\mu - \eta_\mu^2 \beta_\mu\big) \sum_{t=1}^{T} \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} \Big\|^2 + \frac{4\sqrt{n_h} T}{w}, && (95)
\end{aligned}
$$

where we add $\pm L_{t,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1})$ to (93) to obtain (94), and use (93) and $d_t, \hat{d}_t \in [-\sqrt{n_h}, \sqrt{n_h}]$ for any $t \in [T]$, which implies $L_{t+1,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1}) - L_{t,w}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\mu}_{t+1}) \leq 4\sqrt{n_h}/w$, to obtain (95).

We note that since the $\ell_t(\boldsymbol{\theta}_t, \boldsymbol{\mu}_t)$ is defined as the square loss between the ground truth value and our prediction, it is non-negative for all $t \in [T]$, which implies $L_{T+1,w}(\boldsymbol{\theta}_{T+1}, \boldsymbol{\mu}_{T+1}) \geq 0$. Then, by using (95), we write

$$
(\eta_\theta - \eta_\theta^2 \beta_\theta) \sum_{t=1}^{T} \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\theta}} \Big\|^2 + (\eta_\mu - \eta_\mu^2 \beta_\mu) \sum_{t=1}^{T} \Big\| \frac{\partial_{\mathcal{K}_\theta, \eta_\theta} L_{t,w}}{\partial \boldsymbol{\mu}} \Big\|^2 \leq 4\sqrt{n_h} \frac{T}{w}. \tag{96}
$$

By choosing $0 < \eta_\theta \leq 1/(2\beta_\theta)$, $0 < \eta_\mu \leq 1/(2\beta_\mu)$ and dividing both sides of (96) with $\min\{\eta_\theta, \eta_\mu\}/2$, we can obtain the statement in the theorem. □

Appendix C

In this part, we describe how to extend our work for the cross-entropy loss, which is denoted as $RE(\cdot||\cdot)$. Since the cross-entropy loss is mainly used for classification, we describe our extension by using the following RNN architecture:

$$
\begin{aligned}
&\mathbf{h}_{t+1} = \tanh(\mathbf{W}\mathbf{h}_t + \mathbf{U}\mathbf{x}_t) \\
&\hat{d}_t = f(\mathbf{c}^T \mathbf{h}_t) \\
&E_t = RE(d_t || \hat{d}_t).
\end{aligned}
$$

Here, f is assumed to be the sigmoid or softmax function depending on the dimension of the desired sequence. As in (1)-(2), $\mathbf{h}_t \in [-1, 1]^{n_h}$ is the hidden state vector, $\mathbf{x}_t \in [-1, 1]^{n_x}$ is the input vector, and $d_t, \hat{d}_t \in [0, 1]$ is our estimation. Moreover, E_t denotes the cross-entropy loss at time instance t.

As in the squared loss, the cross-entropy is convex with respect to output layer weights, i.e., $\mathbf{c}$. Therefore, we can use the projected online gradient descent – as in (25)– to ensure the convergence of the output layer update. Moreover, since the formula for the derivative of the cross-entropy function with respect to $\mathbf{c}$ is the same with that of the squared loss, i.e., $\frac{0.5(d_t - \hat{d}_t)^2}{\partial \mathbf{c}} = \frac{\partial RE(d_t || \hat{d}_t)}{\partial \mathbf{c}} = (d_t - \hat{d}_t)\mathbf{c}$, the Lipschitz properties derived in Theorem 1 applies to the the cross-entropy loss as well. Since Theorem 2 uses only the Lipschitz properties of the RNN architecture to prove convergence, it can be extended for the cross-entropy with the same update-projection steps in (26)-(27). Therefore, Algorithm 1 can be used for the cross-entropy case without any change as well.

Appendix D

In this part, we describe how to extend our work for LSTMs. The equations of LSTM and the loss function is given as:

$$
\begin{aligned}
\mathbf{z}_t &= \tanh(\mathbf{W}_1 \mathbf{y}_{t-1} + \mathbf{U}_1 \mathbf{x}_t) && (97) \\
\mathbf{i}_t &= \sigma(\mathbf{W}_2 \mathbf{y}_{t-1} + \mathbf{U}_2 \mathbf{x}_t) && (98) \\
\mathbf{f}_t &= \sigma(\mathbf{W}_3 \mathbf{y}_{t-1} + \mathbf{U}_3 \mathbf{x}_t) && (99) \\
\mathbf{c}_t &= \mathbf{i}_t \odot \mathbf{z}_t + \mathbf{f}_t \odot \mathbf{c}_{t-1} && (100) \\
\mathbf{o}_t &= \sigma(\mathbf{W}_4 \mathbf{y}_{t-1} + \mathbf{U}_4 \mathbf{x}_t) && (101) \\
\mathbf{y}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t) && (102) \\
\hat{d}_t &= \mathbf{c}^T \mathbf{y}_t && (103) \\
E_t &= 0.5(d_t - \hat{d}_t)^2. && (104)
\end{aligned}
$$

where $\odot$ denotes the element-wise multiplication, $\mathbf{c}_t \in \mathbb{R}^{n_h}$ is the state vector, $\mathbf{x}_t \in [-1, 1]^{n_x}$ is the input vector, and $\mathbf{y}_t \in \mathbb{R}^{n_h}$ is the output vector, and $\hat{d}_t \in [-\sqrt{n_h}, \sqrt{n_h}]$ is our final estimation. Furthermore, the sigmoid function $\sigma(.)$ and

the hyperbolic tangent function $\tanh(.)$ applies point wise to the vector elements. The weight matrices are $\mathbf{W}_i \in \mathbb{R}^{n_h \times n_h}$, $\mathbf{U}_i \in \mathbb{R}^{n_h \times n_x}$ for $i = 1, \cdots, 4$, and $\mathbf{c} \in \mathbb{R}^{n_h}$, with $\|\mathbf{c}\| \leq 1$. As in the vanilla RNN case, the boundedness of $\mathbf{c}$ can be guaranteed with a proper projection onto a convex set. Note that although we do not explicitly write the bias terms, they can be included in (97)-(103) by augmenting the input vector with a constant dimension.

Similar to the vanilla RNN case, the loss function E_t is convex with respect to the output weights $\mathbf{c}$. Therefore, we can use the projected gradient descent to ensure the convergence of the output layer update. For the hidden layer weights, note that we define the projected gradients as in (4)-(5) and the regret as as in (7) to ensure to find a stationary point for the gradient-descent updates. Then, by using the same intuition, we can extend the projected gradient definition for LSTM as

$$\frac{\partial_{\mathcal{K}_\theta,\eta_\theta}\ell_t(\boldsymbol{\theta}_i,\boldsymbol{\mu}_i)}{\partial\boldsymbol{\theta}_i} := \frac{1}{\eta_\theta}\Big(\boldsymbol{\theta}_i - \Pi_{\mathcal{K}_\theta}\Big[\boldsymbol{\theta}_i - \eta_\theta\frac{\partial\ell_t(\boldsymbol{\theta}_i,\boldsymbol{\mu}_i)}{\partial\boldsymbol{\theta}_i}\Big]\Big),$$
$$\frac{\partial_{\mathcal{K}_\mu,\eta_\mu}\ell_t(\boldsymbol{\theta}_i,\boldsymbol{\mu}_i)}{\partial\boldsymbol{\mu}_i} := \frac{1}{\eta_\mu}\Big(\boldsymbol{\mu}_i - \Pi_{\mathcal{K}_\mu}\Big[\boldsymbol{\mu}_i - \eta_\mu\frac{\partial\ell_t(\boldsymbol{\theta}_i,\boldsymbol{\mu}_i)}{\partial\boldsymbol{\mu}_i}\Big]\Big)$$

and the regret definition as

$$R_w(T) := \sum_{t=1}^{T}\sum_{i=1}^{4}\Big(\Big\|\frac{\partial_{\mathcal{K}_\theta,\eta_\theta}L_{t,w}(\boldsymbol{\theta}_{i,t},\boldsymbol{\mu}_{i,t})}{\partial\boldsymbol{\theta}_i}\Big\|^2 + \Big\|\frac{\partial_{\mathcal{K}_\mu,\eta_\mu}L_{t,w}(\boldsymbol{\theta}_{i,t},\boldsymbol{\mu}_{i,t})}{\partial\boldsymbol{\mu}_i}\Big\|^2\Big), \tag{105}$$

where $\boldsymbol{\theta}_{i,t}$and $\boldsymbol{\mu}_{i,t}$ are the vectorized forms of the weight matrices $\mathbf{W}_{i,t}$ and $\mathbf{U}_{i,t}$ at time t, i.e., $\boldsymbol{\theta}_{i,t} = \text{vec}(\mathbf{W}_{i,t})$ and $\boldsymbol{\mu}_{i,t} = \text{vec}(\mathbf{U}_{i,t})$.

We note that the convergence analysis for our algorithm requires the Lipschitz properties of the architecture of interest. To this end, we can use the Lipschitz properties of LSTMs derived in [46, Proposition 2]. Then, we can use these results in Theorem 2 to upper-bound the regret defined in (105). Accordingly, Algorithm 1 can be extended by for the LSTM optimization.

References

[1] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. New York, NY, USA: Cambridge University Press, 2006.

[2] V. G. Vovk, "Aggregating strategies," in *Proceedings of the Third Annual Workshop on Computational Learning Theory*, ser. COLT '90. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1990, pp. 371–386.

[3] D. F. Specht, "A general regression neural network," *IEEE Transactions on Neural Networks*, vol. 2, no. 6, pp. 568–576, Nov 1991.

[4] T. Ergen and S. S. Kozat, "Efficient online learning algorithms based on lstm neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 8, pp. 3772–3783, Aug 2018.

[5] A. C. Singer, G. W. Wornell, and A. V. Oppenheim, "Nonlinear autoregressive modeling and estimation in the presence of noise," *Digital Signal Processing*, vol. 4, no. 4, pp. 207 – 221, 1994. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1051200484710219

[6] Y. Engel, S. Mannor, and R. Meir, "The kernel recursive least-squares algorithm," *IEEE Transactions on Signal Processing*, vol. 52, no. 8, pp. 2275–2285, Aug 2004.

[7] S. Haykin, *Neural Networks: A Comprehensive Foundation*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1998.

[8] Tsungnan Lin, B. G. Horne, P. Tino, and C. L. Giles, "Learning long-term dependencies in narx recurrent neural networks," *IEEE Transactions on Neural Networks*, vol. 7, no. 6, pp. 1329–1338, Nov 1996.

[9] J. Y. Goulermas, P. Liatsis, X. Zeng, and P. Cook, "Density-driven generalized regression neural networks (dd-grnn) for function approximation," *IEEE Transactions on Neural Networks*, vol. 18, no. 6, pp. 1683–1696, Nov 2007.

[10] N. D. Vanli, M. O. Sayin, I. Delibalta, and S. S. Kozat, "Sequential nonlinear learning for distributed multiagent systems via extreme learning machines," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 3, pp. 546–558, March 2017.

[11] S. Lawrence and C. L. Giles, "Overfitting and neural networks: conjugate gradient and backpropagation," in *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, vol. 1, July 2000, pp. 114–119 vol.1.

[12] I. Arel, D. C. Rose, and T. P. Karnowski, "Deep machine learning - a new frontier in artificial intelligence research [research frontier]," *IEEE Computational Intelligence Magazine*, vol. 5, no. 4, pp. 13–18, Nov 2010.

[13] L. Shao, D. Wu, and X. Li, "Learning deep and wide: A spectral method for learning deep networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 25, no. 12, pp. 2303–2308, Dec 2014.

[14] M. Hermans and B. Schrauwen, "Training and analysing deep recurrent neural networks," in *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2013, pp. 190–198.

[15] K. S. Narendra and K. Parthasarathy, "Identification and control of dynamical systems using neural networks," *IEEE Transactions on Neural Networks*, vol. 1, no. 1, pp. 4–27, March 1990.

[16] P. Liu, X. Qiu, and X. Huang, "Recurrent neural network for text classification with multi-task learning," 2016.

[17] N. Laptev, J. Yosinski, L. E. Li, and S. Smyl, "Time-series extreme event forecasting with neural networks at uber," in *International Conference on Machine Learning*, no. 34, 2017, pp. 1–5.

[18] W. D. Mulder, S. Bethard, and M.-F. Moens, "A survey on the application of recurrent neural networks to statistical language modeling," *Computer Speech & Language*, vol. 30, no. 1, pp. 61 – 98, 2015.

[19] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014.

[20] J. Martens and I. Sutskever, "Learning recurrent neural networks with hessian-free optimization," in *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ser. ICML'11. USA: Omnipress, 2011, pp. 1033–1040. [Online]. Available: http://dl.acm.org/citation.cfm?id=3104482.3104612

[21] G. V. Puskorius and L. A. Feldkamp, "Neurocontrol of nonlinear dynamical systems with kalman filter trained recurrent networks," *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 279–297, March 1994.

[22] G. Hinton, "Neural networks for machine learning," 2012.

[23] S. Ruder, "An overview of gradient descent optimization algorithms," 2016.

[24] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 157–166, March 1994.

[25] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "LSTM: A search space odyssey," *CoRR*, vol. abs/1503.04069, 2015. [Online]. Available: http://arxiv.org/abs/1503.04069

[26] D. Krueger and R. Memisevic, "Regularizing rnns by stabilizing activations," 2015.

[27] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," 2012.

[28] L. Blier, P. Wolinski, and Y. Ollivier, "Learning with random learning rates," 2018.

[29] F. Orabona and T. Tommasi, "Training deep networks without learning rates through coin betting," 2017.

[30] M. Hardt, T. Ma, and B. Recht, "Gradient descent learns linear dynamical systems," 2016.

[31] S. Oymak, "Stochastic gradient descent learns state equations with nonlinear activations," 2018.

[32] Z. Allen-Zhu, Y. Li, and Z. Song, "On the convergence rate of training recurrent neural networks," *CoRR*, vol. abs/1810.12065, 2018. [Online]. Available: http://arxiv.org/abs/1810.12065

[33] J. L. Elman, "Finding structure in time," *COGNITIVE SCIENCE*, vol. 14, no. 2, pp. 179–211, 1990.

[34] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

[35] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural Comput.*, vol. 1, no. 2, pp. 270–280, Jun. 1989. [Online]. Available: http://dx.doi.org/10.1162/neco.1989.1.2.270

[36] Z. C. Lipton, J. Berkowitz, and C. Elkan, "A critical review of recurrent neural networks for sequence learning," 2015.

[37] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," vol. 2, 04 2003.

[38] S. Aydore, L. Dicker, and D. Foster, "A local regret in nonconvex online learning," 2018.

[39] E. Hazan, K. Singh, and C. Zhang, "Efficient regret minimization in non-convex games," *CoRR*, vol. abs/1708.00075, 2017. [Online]. Available: http://arxiv.org/abs/1708.00075

[40] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.

[41] R. J. Williams and D. Zipser, "Backpropagation," Y. Chauvin and D. E. Rumelhart, Eds. Hillsdale, NJ, USA: L. Erlbaum Associates Inc., 1995, ch. Gradient-based Learning Algorithms for Recurrent Networks and Their Computational Complexity, pp. 433–486. [Online]. Available: http://dl.acm.org/citation.cfm?id=201784.201801

[42] J. Alcal-Fdez, A. Fernndez, J. Luengo, J. Derrac, and S. Garca, "Keel data-mining software tool: Data set repository, integration of algorithms and experimental analysis framework." *Multiple-Valued Logic and Soft Computing*, vol. 17, no. 2-3, pp. 255–287, 2011.

[43] C. E. Rasmussen, "Delve data sets," http://www.cs.toronto.edu/~delve/data/datasets.html, accessed: 2019-07-21.

[44] P. Lancaster and H. Farahat, "Norms on direct sums and tensor products," *Mathematics of Computation - Math. Comput.*, vol. 26, 05 1972.

[45] S. Bubeck, "Convex optimization: Algorithms and complexity," *Found. Trends Mach. Learn.*, vol. 8, no. 3-4, pp. 231–357, Nov. 2015. [Online]. Available: http://dx.doi.org/10.1561/2200000050

[46] J. Miller and M. Hardt, "Stable recurrent models." in *ICLR (Poster)*. OpenReview.net, 2019.