



OPEN ACCESS

EDITED BY

Amirreza Yousefzadeh,
University of Twente, Netherlands

REVIEWED BY

Guangzhi Tang,
Maastricht University, Netherlands
Neelesh Kumar,
Procter & Gamble, United States

*CORRESPONDENCE

Jesús García Fernández
✉ jesus.garciafernandez@donders.ru.nl

RECEIVED 27 May 2024

ACCEPTED 25 June 2024

PUBLISHED 10 July 2024

CITATION

Fernández JG, Keemink S and van Gerven M (2024) Gradient-free training of recurrent neural networks using random perturbations. *Front. Neurosci.* 18:1439155. doi: 10.3389/fnins.2024.1439155

COPYRIGHT

© 2024 Fernández, Keemink and van Gerven. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Gradient-free training of recurrent neural networks using random perturbations

Jesús García Fernández*, Sander Keemink and Marcel van Gerven

Department of Machine Learning and Neural Computing, Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, Netherlands

Recurrent neural networks (RNNs) hold immense potential for computations due to their Turing completeness and sequential processing capabilities, yet existing methods for their training encounter efficiency challenges. Backpropagation through time (BPTT), the prevailing method, extends the backpropagation (BP) algorithm by unrolling the RNN over time. However, this approach suffers from significant drawbacks, including the need to interleave forward and backward phases and store exact gradient information. Furthermore, BPTT has been shown to struggle to propagate gradient information for long sequences, leading to vanishing gradients. An alternative strategy to using gradient-based methods like BPTT involves stochastically approximating gradients through perturbation-based methods. This learning approach is exceptionally simple, necessitating only forward passes in the network and a global reinforcement signal as feedback. Despite its simplicity, the random nature of its updates typically leads to inefficient optimization, limiting its effectiveness in training neural networks. In this study, we present a new approach to perturbation-based learning in RNNs whose performance is competitive with BPTT, while maintaining the inherent advantages over gradient-based learning. To this end, we extend the recently introduced activity-based node perturbation (ANP) method to operate in the time domain, leading to more efficient learning and generalization. We subsequently conduct a range of experiments to validate our approach. Our results show similar performance, convergence time and scalability when compared to BPTT, strongly outperforming standard node perturbation and weight perturbation methods. These findings suggest that perturbation-based learning methods offer a versatile alternative to gradient-based methods for training RNNs which can be ideally suited for neuromorphic computing applications.

KEYWORDS

recurrent neural network, artificial neural network, gradient approximation, BPTT, node perturbation learning

1 Introduction

Recurrent neural networks (RNNs), with their ability to process sequential data and capture temporal dependencies, have found applications in tasks such as natural language processing (Yao et al., 2013; Cho et al., 2014; Sutskever et al., 2014) and time series prediction (Hewamalage et al., 2021). They hold immense potential for computation thanks to their Turing completeness (Chung and Siegelmann, 2021). Furthermore, due to their sequential processing capabilities, they offer high versatility to process variable-sequence length inputs and fast inference on long sequences (Orvieto et al., 2023). Nevertheless, traditional training methods like backpropagation through time

(BPTT) (Werbos, 1990) are challenging to apply (Bengio et al., 1994; Lillicrap and Santoro, 2019), particularly with long sequences. Unrolling the RNN over time for gradient propagation and weight updating proves to be computationally demanding and difficult to parallelize with variable-length sequences. Additionally, employing BPTT can result in issues like vanishing or exploding gradients (Pascanu et al., 2013). Moreover, the non-locality of their updates can pose significant challenges when implemented on unconventional computing platforms (Kaspar et al., 2021).

An alternative approach to training neural networks is stochastically approximating gradients through perturbation-based methods (Widrow and Lehr, 1990; Spall, 1992; Werfel et al., 2003). In this type of learning, synaptic weights are adjusted based on the impact of introducing perturbations to the network. When the perturbation enhances performance, the weights are strengthened, driven by a global reinforcement signal, and vice versa. This method is computationally simple, relying solely on forward passes and a reinforcement signal distribution across the network. It differs from gradient-based methods, such as BP, which needs a specific feedback circuit to propagate specific signals and a dedicated backward pass to compute explicit errors. This simplicity is especially beneficial for RNNs, as it eliminates the need to unroll the network over time during training (Werbos, 1990). Examples of gradient-based and perturbation-based approaches to update the neural weights are visually depicted in Figure 1. Standard perturbation methods include node perturbation (NP), where the perturbations are added into the neurons, and weight perturbation (WP), where the perturbations are added into the synaptic weights (Werfel et al., 2003; Züge et al., 2023).

Nevertheless, the stochastic nature of perturbation-based updates can lead to inefficiencies in optimizing the loss landscape (Lillicrap et al., 2020), resulting in prolonged convergence times (Werfel et al., 2003). Additionally, perturbation-based often exhibit poor scalability, leading to inferior performance when compared to gradient-based methods (Hiratani et al., 2022). This performance gap increases as the network size increases, and in large networks, instability often arises manifesting as extreme weight growth (Hiratani et al., 2022). Due to these challenges, a number of recent implementations have aimed to establish perturbation-based learning as an effective gradient-free method for neural network training. In Lansdell et al. (2019), NP is utilized to approximate the feedback system in feedforward and convolutional networks, mitigating the error computation aspect of the BP algorithm. However, a dedicated feedback system is still necessary to propagate synapse-specific errors. In Züge et al. (2023), standard NP and WP are employed on temporally extended tasks in RNNs, though direct comparisons with gradient-based algorithms like BP are lacking. Their results suggest that WP may outperform NP in specific cases.

In this study, we present an implementation of perturbation-based learning in RNNs whose performance is competitive with BP, while maintaining the inherent advantages over gradient-based learning. To this end, we extend the activity-based node perturbation (ANP) approach (Dalm et al., 2023) to operate in the time domain using RNNs. This approach relies solely on neural activities, eliminating the need for direct access to the noise process. The resulting updates align more closely with directional derivatives, compared to standard NP, approximating SGD more

accurately. As a result, this approach significantly outperforms the standard NP in practical tasks. Furthermore, we extend standard implementations of NP and WP to operate in the time domain, using them as baselines along with BPTT, referred to as BP in this study for simplicity. In addition, we augment our time-extended implementations of ANP, NP, and WP (as well as BP) with a decorrelation mechanism, as done in the original ANP work (Dalm et al., 2023). This decorrelation mechanism has been shown to significantly accelerate the training of deep neural networks (Dalm et al., 2024) by aligning SGD updates with the natural gradient through the use of uncorrelated input variables (Desjardins et al., 2015). We name the resulting decorrelated methods DANP, DNP, DWP and DBP.

We assess the efficiency of our approach across various tasks. Firstly, we evaluate learning performance using three common machine learning benchmarks. Secondly, we examine the scalability of our approach to larger networks in terms of stability and task performance. Results indicate similar learning performance, convergence time, generalization and scalability compared to BP, with significant superiority over standard NP and WP. In contrast to gradient-based methods, the proposed method also offers increased versatility, with its local computations potentially rendering it compatible with neuromorphic hardware (Schuman et al., 2022).

In the following sections, we detail our adaptation of the different perturbation-based methods utilized here, namely NP, WP and ANP. For standard approaches like NP and WP, we describe how our time-based extensions diverge from the more commonly employed temporal extensions. Additionally, we analyse the operation of the decorrelation mechanism within our RNNs. Subsequently, we describe the series of experiments conducted in this study and present the resulting outcomes.

2 Methods

In this section, we describe our original time-extended implementations of existing perturbation-based algorithms for feedforward networks, adapted for compatibility with RNNs and time-extended tasks. We highlight how our implementations differ from other time-extended methods in the literature. For clarity and ease of understanding, we present them in the following order: Node perturbation (NP), activity-based node perturbation (ANP), and weight perturbation (WP). These three methods will throughout be compared to backpropagation (BP). BP, WP and NP will throughout be used as baseline methods, with ANP (extended to the time-domain) being our novel contribution. Additionally, we describe the decorrelation mechanism proposed by Dalm et al. (2023) for feedforward networks, which we incorporate into our RNN model to enhance the learning efficiency of the algorithms under consideration.

2.1 Recurrent neural network model

In all our experiments, we employ RNNs with one hidden layer containing a large number N of units with learnable weights and non-linearities in the neural outputs. The structure of the neural

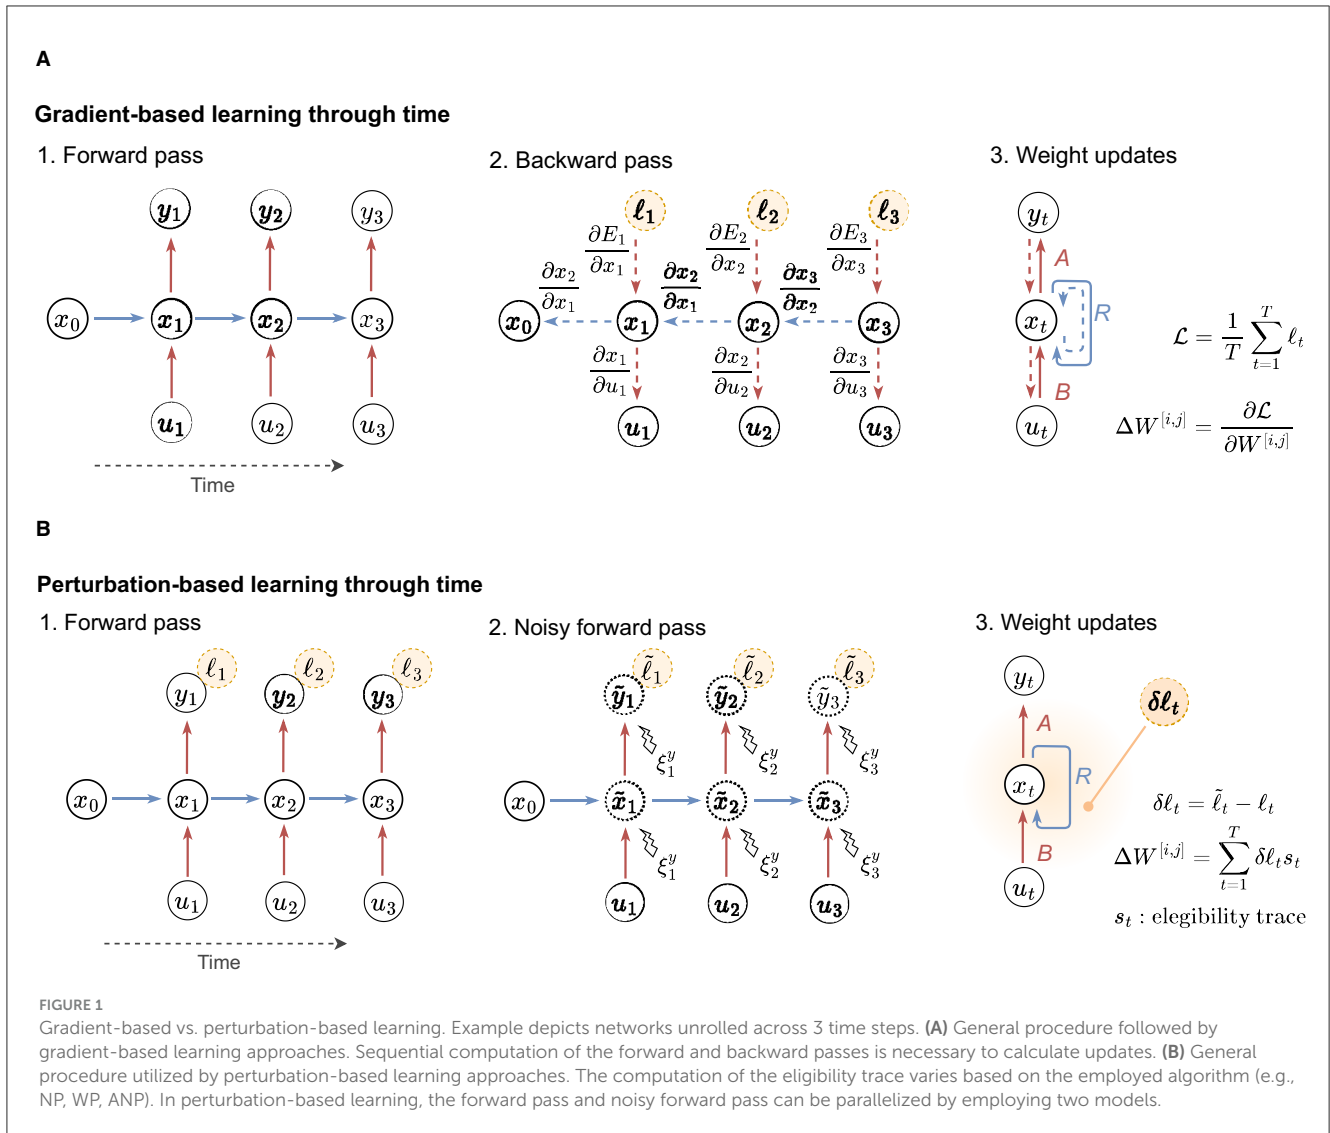


FIGURE 1 Gradient-based vs. perturbation-based learning. Example depicts networks unrolled across 3 time steps. **(A)** General procedure followed by gradient-based learning approaches. Sequential computation of the forward and backward passes is necessary to calculate updates. **(B)** General procedure utilized by perturbation-based learning approaches. The computation of the eligibility trace varies based on the employed algorithm (e.g., NP, WP, ANP). In perturbation-based learning, the forward pass and noisy forward pass can be parallelized by employing two models.

networks with recurrently connected hidden units is depicted in Figure 2, with the forward pass defined as

$$x_t = f(Au_t + Rx_{t-1})$$

$$y_t = Bx_t$$

where u_t , x_t and y_t denote input, hidden and output activations at time t , respectively, and $f(\cdot)$ represent a non-linear activation function. In our networks, we use the hyperbolic tangent activation function $\tanh(x) = (e^x - e^{-x}) / (e^x + e^{-x})$. The parameters of the network are given by the input weights A , the output weights B and the recurrent weights R .

The output error at each time step is defined as $\ell_t = \|y_t - y_t^*\|^2$ where y_t^* denotes the target output at time t . The loss associated with output $y = (y_1, \dots, y_T)$, with T the time horizon, is defined as the mean of the output errors over the entire sequence:

$$\mathcal{L} = \frac{1}{T} \sum_{t=1}^T \ell_t.$$

During learning, the network's weights undergo incremental and iterative updates to minimize the loss according to $W \leftarrow W - \eta \Delta W$ with $W \in \{A, R, B\}$ the input, output and recurrent weight

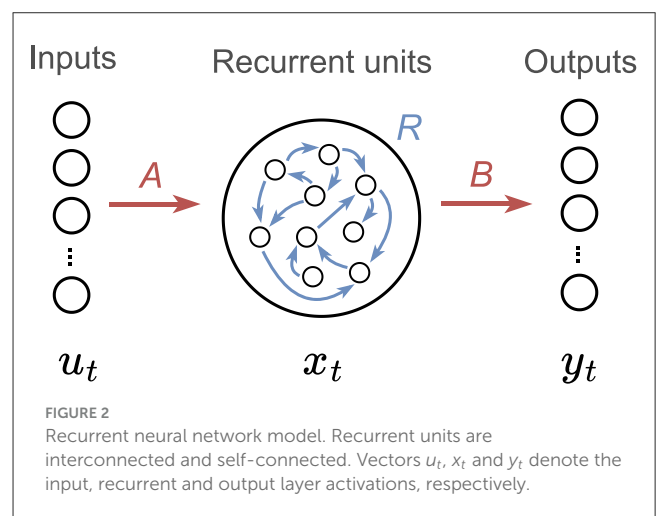


FIGURE 2 Recurrent neural network model. Recurrent units are interconnected and self-connected. Vectors u_t , x_t and y_t denote the input, recurrent and output layer activations, respectively.

matrices, respectively, and ΔW their corresponding updates. Here, η denotes the learning rate for both forward and recurrent weights.

The updates can be computed using different learning algorithms. In the case of backpropagation through time, the updates are given by the gradient $\nabla \mathcal{L}$ of the loss with respect to the

parameters, averaged over multiple trajectories. Below, we describe the different perturbation-based learning methods considered in this study. Additionally, we describe the employed decorrelation mechanism and how it is incorporated into the networks.

2.2 Node perturbation through time

The node perturbation approach involves two forward passes: A standard forward pass and a noisy pass. These passes can take place either concurrently or sequentially, and the loss is computed afterwards. The weights are updated in the direction of the noise if the loss decreases, and in the opposite direction if the loss increases. During the noisy pass, noise is added to the pre-activation of each neuron as follows:

$$\begin{aligned} \tilde{x}_t &= f(Au_t + R\tilde{x}_{t-1} + \xi_t) \\ \tilde{y}_t &= B\tilde{x}_t + v_t \end{aligned}$$

where \tilde{x}_t and \tilde{y}_t denote the noisy neural outputs at time t . The noise added in the hidden and output layers is generated from zero-mean, uncorrelated Gaussian random variables $\xi_t \sim \mathcal{N}(0, \sigma^2 I_x)$ and $v_t \sim \mathcal{N}(0, \sigma^2 I_y)$, respectively, where I_x and I_y are identity matrices with the dimensions of the hidden and output layers. The noise injected is different in every timestep t .

In typical implementations of node perturbation in the time domain, as seen in works like (Fiete and Seung, 2006; Züge et al., 2023), the reinforcement signal is derived from the difference in loss, $\mathcal{L} = \tilde{\mathcal{L}} - \mathcal{L}$, where $\tilde{\mathcal{L}}$ and \mathcal{L} represent the loss of the noisy pass and the clean pass, respectively. This signal captures the network's overall performance across the entire sequence. Additionally, an eligibility trace, computed as the sum over time of the pre-synaptic neuron's output multiplied by the injected perturbation ($\xi_t x_t^\top$, $v_t u_t^\top$ or $\xi_t x_{t-1}^\top$, depending on the synaptic weights being updated), is utilized. According to this method, the learning signals for updating the weights are defined in Equation 1.

$$\begin{aligned} \Delta A &= \sigma^{-2} \delta \mathcal{L} \sum_{t=1}^T \xi_t x_t^\top, & \Delta B &= \sigma^{-2} \delta \mathcal{L} \sum_{t=1}^T v_t u_t^\top, \\ \Delta R &= \sigma^{-2} \delta \mathcal{L} \sum_{t=1}^T \xi_t x_{t-1}^\top. \end{aligned} \tag{1}$$

While this approach offers benefits such as enhanced compatibility with delayed rewards, we here consider a local approach, where we only consider the local loss difference $\delta \ell_t = \tilde{\ell}_t - \ell_t$ in individual timesteps when computing updates in contrast to using the total loss difference \mathcal{L} . This reward-per-time step is employed alongside eligibility traces local on time defined as the product of the pre-synaptic neuron's output and the injected perturbation at each time step ($\xi_t x_t^\top$, $v_t u_t^\top$ or $\xi_t x_{t-1}^\top$, depending on the synaptic weights being updated). According to our approach, the learning signals to update the weights over time are defined in Equation 2.

$$\begin{aligned} \Delta A &= \sigma^{-2} \sum_{t=1}^T \delta \ell_t \xi_t x_t^\top, & \Delta B &= \sigma^{-2} \sum_{t=1}^T \delta \ell_t v_t u_t^\top, \\ \Delta R &= \sigma^{-2} \sum_{t=1}^T \delta \ell_t \xi_t x_{t-1}^\top. \end{aligned} \tag{2}$$

Here, both the clean standard pass and the noisy pass could run concurrently using two identical copies of the model, which enables compatibility with online learning setups. This technique allows our method to compute and implement updates online at every time step.

Zenke and Nefci (2020) investigate setups similar to the one proposed here, seeking to bridge the real-time recurrent learning (RTRL) algorithm (Williams and Zipser, 1989), which is more effective for online setups than BPTT but computationally demanding, with biologically plausible learning rules. They demonstrate that by combining learning algorithms that approximate RTRL with temporally local losses, effective approximations can be achieved. These approximations notably decrease RTRL's computational cost while preserving strong learning performance. A similar rationale is used by Bellec et al. (2020) in the context of recurrent spiking neural networks. Supplementary materials 1 provides a comparison between the conventional implementation and our implementation of NP through time.

2.3 Activity-based node perturbation through time

Activity-based node perturbation (ANP) is a variant of the node perturbation approach, proposed by Dalm et al. (2023), which has been exclusively applied in feedforward networks. This approach approximates the directional derivatives across the network, resulting in a closer alignment between the updates generated by this method and those provided by BP (also see Baydin et al., 2022). Additionally, it does not require direct access to the noise process itself as it operates solely by measuring changes in neural activity. Given that NP can be interpreted as a noisy variant of SGD (Hiratani et al., 2022), ANP can be seen as a more precise approximation of SGD. For a detailed derivation of the link between ANP and SGD, we refer to (Dalm et al., 2023).

Similar to the node perturbation approach, the noisy pass is the same as the one performed in the standard node perturbation approach. Consistent with our node perturbation implementation extended over time, we calculate reinforcement signals at each time step to drive synaptic changes. Let $\alpha_t = Au_t + Rx_{t-1}$ and $\beta_t = Bx_t$ are the pre-activations in the clean pass. Similarly, let $\tilde{\alpha}_t = Au_t + R\tilde{x}_{t-1} + \xi_t$ and $\tilde{\beta}_t = B\tilde{x}_t + v_t$ denote the pre-activations in the noisy pass. Define N as the total number of neurons in the network. We compute the learning signals responsible for weight updating as defined in Equation 3.

$$\begin{aligned} \Delta A &= N \sum_{t=1}^T \ell_t \frac{\delta \alpha_t}{\|\delta \alpha_t\|^2} x_t^\top, & \Delta B &= N \sum_{t=1}^T \ell_t \frac{\delta \beta_t}{\|\delta \beta_t\|^2} u_t^\top, \\ \Delta R &= N \sum_{t=1}^T \ell_t \frac{\delta \alpha_t}{\|\delta \alpha_t\|^2} x_{t-1}^\top, \end{aligned} \tag{3}$$

where $\delta \alpha_t = \tilde{\alpha}_t - \alpha_t$ and $\delta \beta_t = \tilde{\beta}_t - \beta_t$ are the pre-activation differences between the forward passes.

2.4 Weight perturbation through time

Weight perturbation is an approach akin to node perturbation, where noise is injected in a second forward pass, and adjustments to the weights are made based on the resulting increase or decrease in loss. The key distinction lies in the injection of noise into the weights rather than the neural pre-activation. The noisy pass is defined as

$$\begin{aligned} \tilde{x}_t &= f((A + \xi_t) u_t + (R + \nu_t) \tilde{x}_{t-1}) \\ \tilde{y}_t &= (B + \zeta_t) \tilde{x}_t \end{aligned}$$

where \tilde{x}_t and \tilde{y}_t denote the noisy neural outputs at time t . As in node perturbation, the noise is denoted by the zero-mean, uncorrelated Gaussian random variables $\xi_t \sim \mathcal{N}(0, \sigma^2 I_A)$, $\nu_t \sim \mathcal{N}(0, \sigma^2 I_R)$ and $\zeta_t \sim \mathcal{N}(0, \sigma^2 I_B)$, where I_A , I_R and I_B are identity matrices with the dimensions of A , R and B , with distinct values for each timestep t .

Similar to node perturbation in the time domain, typical implementations of weight perturbation in the time domain, such as those described by (Cauwenberghs, 1992; Züge et al., 2023), derive the reinforcement signal from the difference in loss $\delta\mathcal{L} = \tilde{\mathcal{L}} - \mathcal{L}$, computed over the entire sequence. According to this method, the learning signals for updating the weights are defined as in Equation 4.

$$\Delta A = \sigma^{-2} \delta\mathcal{L} \xi_t, \quad \Delta B = \sigma^{-2} \delta\mathcal{L} \zeta_t, \quad \Delta R = \sigma^{-2} \delta\mathcal{L} \nu_t. \quad (4)$$

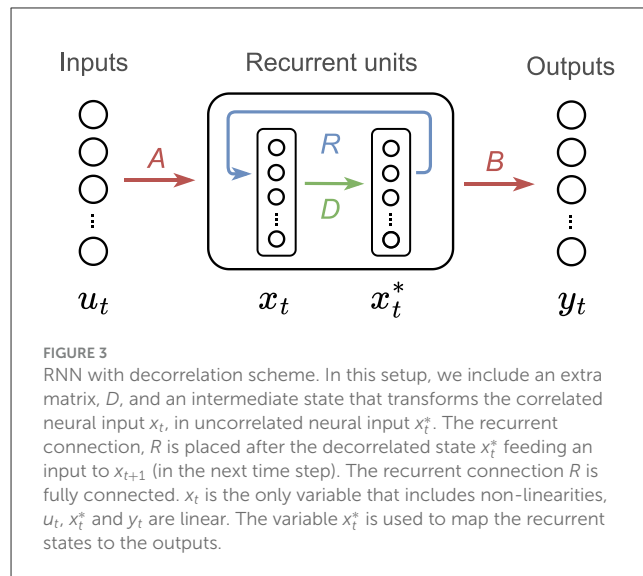
This approach inherits the same set of drawbacks and benefits as seen in node perturbation. Hence, we again employ reinforcement signals, $\delta\ell_t$, computed at each time step, to drive synaptic changes. We define the learning signals to update the weights in the time domain as in Equation 5.

$$\begin{aligned} \Delta A &= \sigma^{-2} \sum_{t=1}^T \delta\ell_t \xi_t, & \Delta B &= \sigma^{-2} \sum_{t=1}^T \delta\ell_t \zeta_t, \\ \Delta R &= \sigma^{-2} \sum_{t=1}^T \delta\ell_t \nu_t. \end{aligned} \quad (5)$$

2.5 Decorrelation of neural inputs

Decorrelating neural input allows for more efficient neural representation by reducing the redundancy in neural activity, leading to improved efficiency in learning and faster learning rates. This phenomenon has found support in both biological studies (Wiechert et al., 2010; Cayco-Gajic et al., 2017) and artificial neural network research (Desjardins et al., 2015; Luo, 2017; Huang et al., 2018; Ahmad et al., 2022). It has also been shown that it significantly accelerates the training of deep neural networks (Dalm et al., 2024) by aligning SGD updates with the natural gradient through the use of uncorrelated input variables (Desjardins et al., 2015). Additionally, in the context of global reinforcement methods such as WP or NP, where weight updates introduce substantial noise, the addition of decorrelation proves beneficial as it makes neural networks less sensitive to noise (Tetzlaff et al., 2012).

Given these benefits, we investigate the integration of a decorrelation mechanism into our networks, presented in Dalm et al. (2023). We apply this mechanism solely to the hidden



units, as the inputs fed into the network are low-dimensional, thus decorrelating them would not yield any noticeable changes and would increase computational costs. This involves the transformation of correlated hidden layer input x_{t-1} into decorrelated input x_{t-1}^* via a linear transform $x_{t-1}^* = Dx_{t-1}$ with D the decorrelation matrix. The update of recurrent units in the network is in this case defined as $x_t = f(Au_t + Rx_{t-1}^*)$ and the resulting neural connectivity is depicted in Figure 3.

The updates of the decorrelation weights are performed using a particularly efficient learning rule proposed by Ahmad et al. (2022), which aims at reducing the cross-correlation between the neural outputs. The update is given by $D \leftarrow D - \epsilon \Delta D$ with learning rate ϵ and update defined in Equation 6.

$$\Delta D = \left(x_t^* (x_t^*)^\top - \text{diag} \left((x_t^*)^2 \right) \right) D \quad (6)$$

for $1 \leq t \leq T$. The updates of the decorrelation weights are performed in an unsupervised manner, in parallel with learning of the forward weights. We use DANP to refer to the combination of decorrelation with activity-based node perturbation.

2.6 Experimental validation

We evaluate the effectiveness of the described methods in training RNNs using a series of experiments encompassing several objectives. Firstly, we evaluate the performance of the networks using three standard machine learning benchmarks: Mackey-Glass time series prediction, the copying memory task and a weather prediction task. These tasks are commonly employed in the literature to evaluate the performance of RNNs and other time-series prediction models. Secondly, we assess the scalability of the considered networks when incorporating an increased number of units. Lastly, we investigate in more detail the functioning of the decorrelation mechanism. Five different runs with random seeds are carried out for each experiment. The averages of these runs are then depicted with error bars indicating maximum and minimum values.

Throughout these evaluations, we assess the perturbation-based learning methods NP, WP, and ANP, alongside the gradient-based BP learning method in conjunction with the Adam optimizer (Kingma and Ba, 2014) for comparative analysis. Please note that for simplicity, we use the term BP in this manuscript to refer to both backpropagation and its time-domain application, backpropagation through time (BPTT) (Werbos, 1990), as they are essentially the same algorithm, with the distinction being the unrolling of the network over time. Similarly, the terms NP, WP, and ANP are used in this context to refer to their application in the time domain. Subsequently, we enhance all the methods by incorporating the decorrelation mechanism previously described into the hidden units of our networks. The resulting extended methods are named decorrelated node perturbation (DNP), decorrelated weight perturbation (DWP), decorrelated activity-based node perturbation (DANP) and decorrelated backpropagation (DBP). Detailed hyperparameters for each experiment can be found in [Supplementary materials 2](#).

Our networks and experimental setups are developed using the programming language Python (Van Rossum et al., 1995) and the deep learning framework PyTorch (Paszke et al., 2019). The BPTT algorithm is implemented using the automatic differentiation tool from PyTorch. For reproducibility and to access the implementation code, please visit our GitHub repository.¹

3 Results

3.1 Mackey-Glass time series task

The Mackey-Glass time series task is a classic benchmark for assessing the ability of neural networks to capture and predict chaotic dynamical systems. The data is a sequence of one-dimensional observations generated using the Mackey-Glass delay differential equations (Mackey and Glass, 1977), resulting in a nonlinear, delayed, and chaotic time series. We reproduce the setup of Voelker et al. (2019), where the model is tasked with predicting 15 time steps into the future with a time constant of 17 steps and a sequence length of 5000 time steps.

Figure 4 depicts the results for the Mackey-Glass experiment. Figure 4A visualizes 500 time steps of a synthetically generated Mackey-Glass time series. Additionally, we depict the predictions made by a BP-trained model both before and after training. This example provides a visual understanding of the dataset used in our study. Figure 4B shows the performance during training over the train and test set for the different methods. Figure 4C shows the final performance, computed as the mean performance over the last 50 epochs, facilitating a quantitative comparison between methods.

The outcomes of this experiment reveal that standard perturbation-based methods NP and WP, adapted to operate in the time domain, exhibit significantly inferior performance compared to the BP baseline. The convergence time and final performance of ANP, in contrast, closely approach those of the BP baseline, especially when augmented with the decorrelation mechanism. Introducing the decorrelation mechanism to BP does not appear

to result in a pronounced difference in convergence and final performance.

3.2 Copying memory task

The copying memory task is another well-established task for evaluating the memorization capabilities of recurrent models/units. In this task, the model must memorize a sequence of bits and return the exact same sequence after a specified delay period. We build on the setup of Arjovsky et al. (2016), where the sequence comprises 8 distinct bit values, with 1 extra bit serving to mark the delay period and another extra bit indicating to the network when to reproduce the sequence. In our experiments, we use a length sequence of 100 and a delay period equal to 10.

Figure 5 depicts the results for the copying memory task. In Figure 5A, we present a visualization of a sample of synthetically generated data for this task along with annotations. Additionally, we depict the predictions made by a BP-trained model both before and after training. This example serves to provide a visual understanding of the dataset used in our study. In Figure 5B, we present the performance during training over the train and test sets. Figure 5C shows the final performance, computed as the mean performance over the last 50 epochs, facilitating a quantitative comparison between methods.

The outcomes of this experiment reveal that standard perturbation-based methods NP and WP, adapted to operate in the time domain, exhibit significantly inferior performance compared to the BP baseline. The convergence time and final performance of ANP closely approach those of the BP baseline. In this experiment, the addition of the decorrelation mechanism does not lead to a pronounced difference in convergence and final performance, in both BP and ANP.

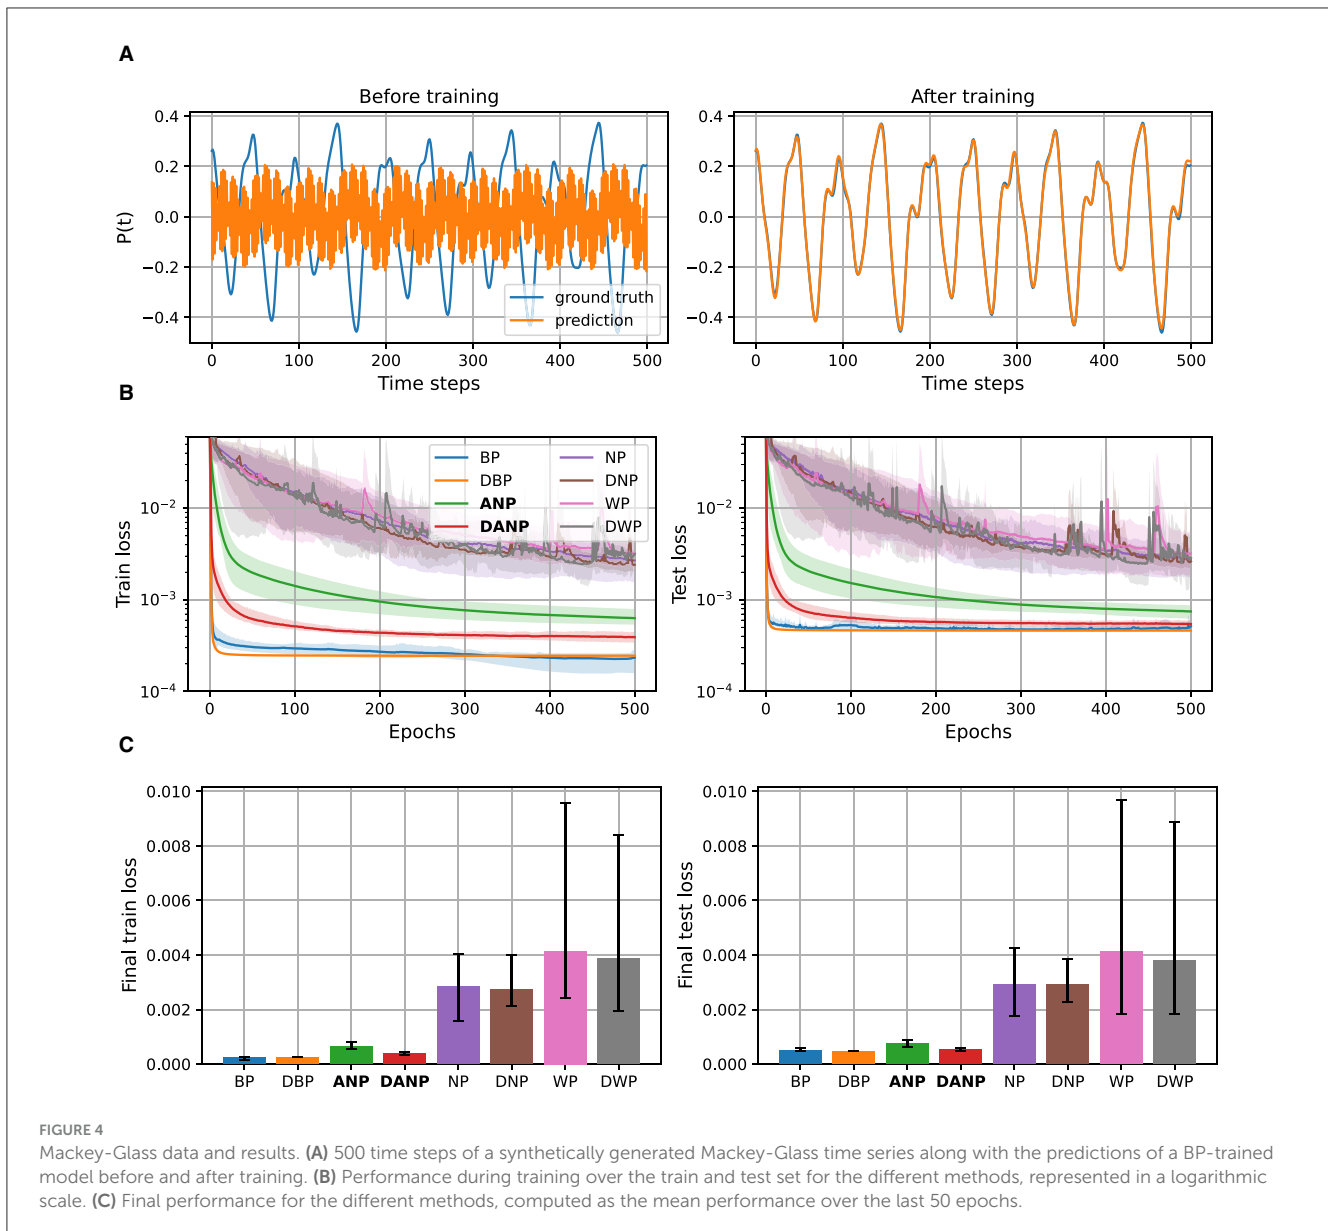
3.3 Weather prediction task

In contrast to the other benchmarks used in this paper, this task relies on real-world rather than synthetically generated data. The dataset used in this task contains climatological data spanning 1,600 U.S. locations from 2010 to 2013.² We build on the setup of Zhou et al. (2021), where each data point consists of one single target value to be predicted 1, 24 and 48 hours in advance, and various input climate features. In our specific configuration, we exclude duplicated features with different units of measurement (retaining Fahrenheit-measured features) and select the 'Dry bulb' feature, which is a synonym for air temperature, as the target variable. The training data encompasses the initial 28 months, while the last 2 months are used for testing the model's performance.

Figure 6 depicts the results for this task. In Figure 6A, we present a visualization of 2000 time steps of the target feature, "Dry bulb", from the weather dataset. Additionally, we depict the predictions 48 hours ahead made by a BP-trained model both before and after training. This example serves to provide a visual

¹ https://github.com/jesusgf96/NP_RNNs

² Data can be obtained from <https://www.ncei.noaa.gov/data/local-climatological-data/>.



understanding of the dataset used in our study. In [Figure 6B](#) we present the performance during training over the train and test set for the different methods for 48-h ahead prediction. Additionally, these figures present the final performance, computed as the mean performance over the last 50 epochs, facilitating a quantitative comparison between methods. In [Supplementary materials 5](#), we also include the results for 1-hour ahead and 24-hours ahead predictions.

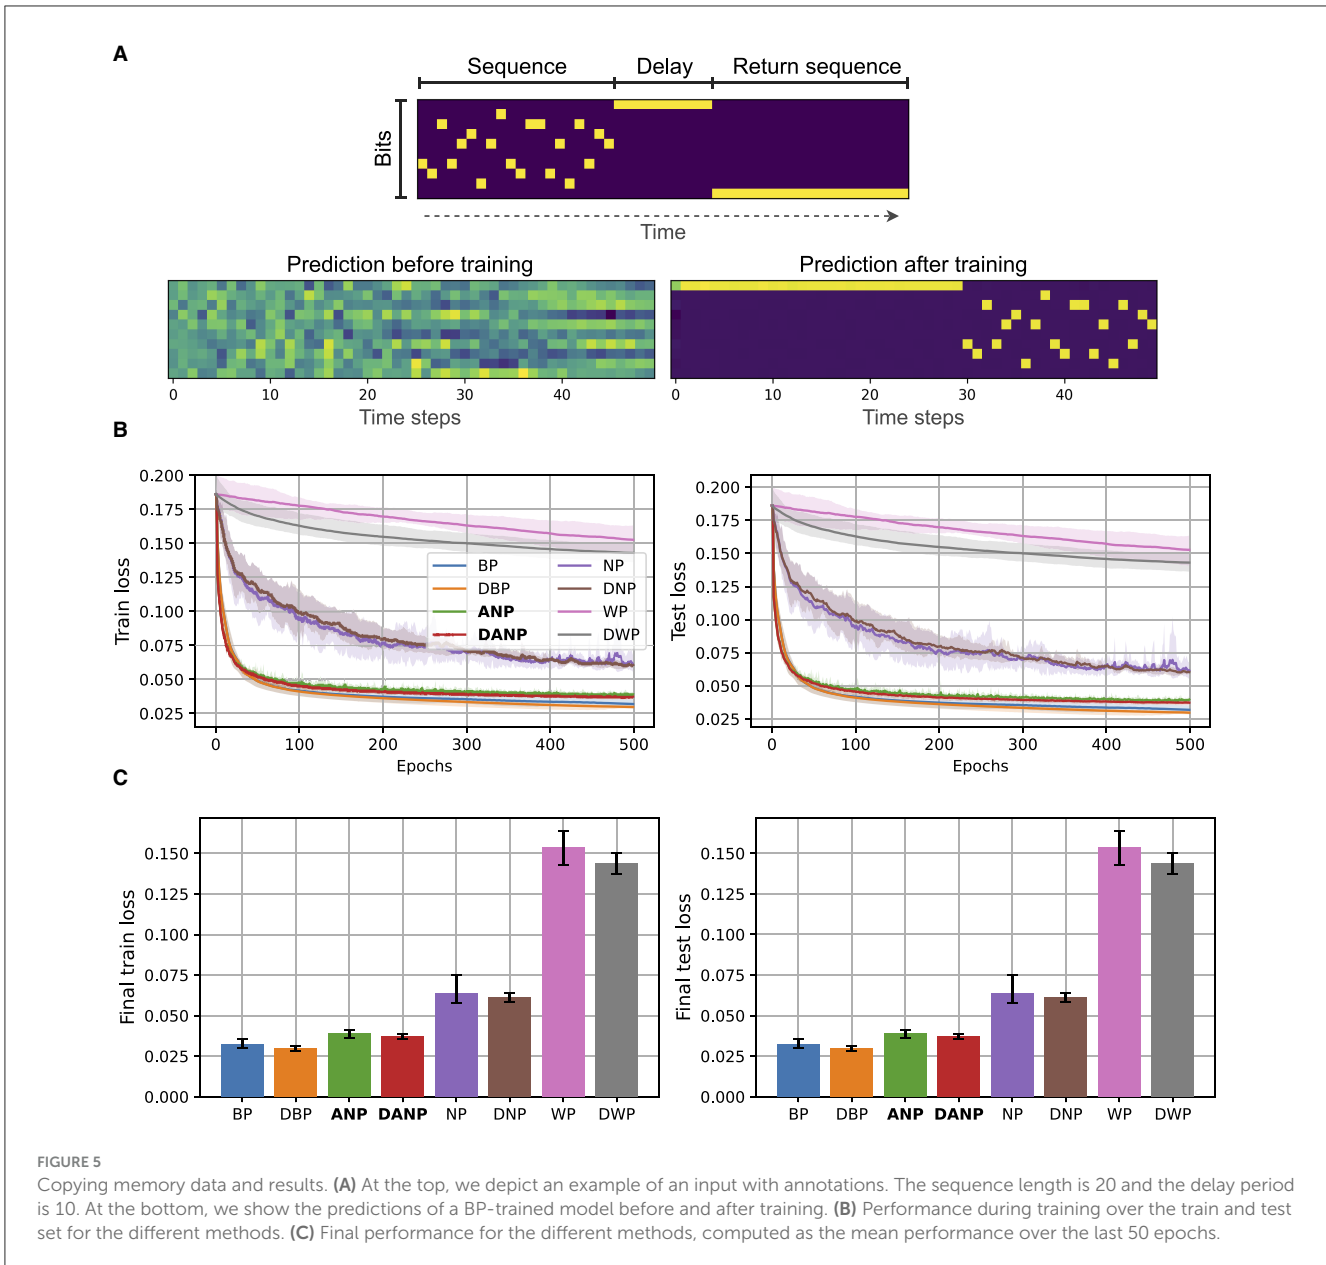
The outcomes of this experiment reveal that standard perturbation-based methods NP and WP, adapted to operate in the time domain, exhibit significantly inferior performance compared to the BP baseline. The convergence time and final performance of ANP closely approach those of the BP baseline, especially when augmented with the decorrelation mechanism, making them comparable in terms of generalization. On the contrary, introducing the decorrelation mechanism to BP results in faster convergence but compromises generalization performance.

3.4 Scaling performance

Here, we investigate the scalability of the described methods in RNNs with a single hidden layer and an increasingly larger number of units. We accomplish this by analyzing the final performance on the 1-hour ahead weather prediction task for networks with differing numbers of hidden units, trained using various methods. We chose this dataset as it is the most challenging among the considered datasets in this paper. The number of hidden units ranges from 100 to 3000.

[Figure 7](#) shows the final performance over the train and test sets of networks with different configurations. Each configuration was repeated using five random seeds, and the results were computed as the mean of the successful executions within the set of five runs.

The results of this experiment reveal two key findings. Firstly, NP and WP exhibit inadequate scalability, leading to unstable runs or inferior performance when employed in the training of



large networks. Secondly, ANP (and its variation incorporating the decorrelation mechanism, DANP) is the only perturbation-based method capable of effective scaling to larger networks, showing performance on par with BP and enhanced generalization when augmented with the decorrelation mechanism.

3.5 Decorrelation results

Here, we investigate the functionality of the decorrelation mechanism within the networks. Using the parameters that yield optimal task-dependent performance, we compare the degree of correlation in the neural outputs between a network incorporating the decorrelation mechanism and one that does not. To achieve this, we visualize the mean squared correlation across epochs during training on the 1-hour ahead weather prediction task.

Figure 8 illustrates the degree of correlation within the hidden units. This represents the decorrelation loss, calculated as the mean squared off-diagonal values of the lower triangular covariance matrix, computed with the recurrent inputs, x_t^* , to the hidden units at each timestep during input presentation, across epochs.

The outcomes of these experiments reveal a low degree of correlation among the hidden units in the network featuring the decorrelation mechanism compared to the network that does not incorporate it. However, as depicted in Figure 8, the parameters that yield optimal task-dependent performance do not lead to completely decorrelated neural outputs. These results suggest that a small level of correlation is needed to achieve optimal task-dependent performance, which may be due to the need for the weight updates to be able to keep up with decorrelation parameter updates.

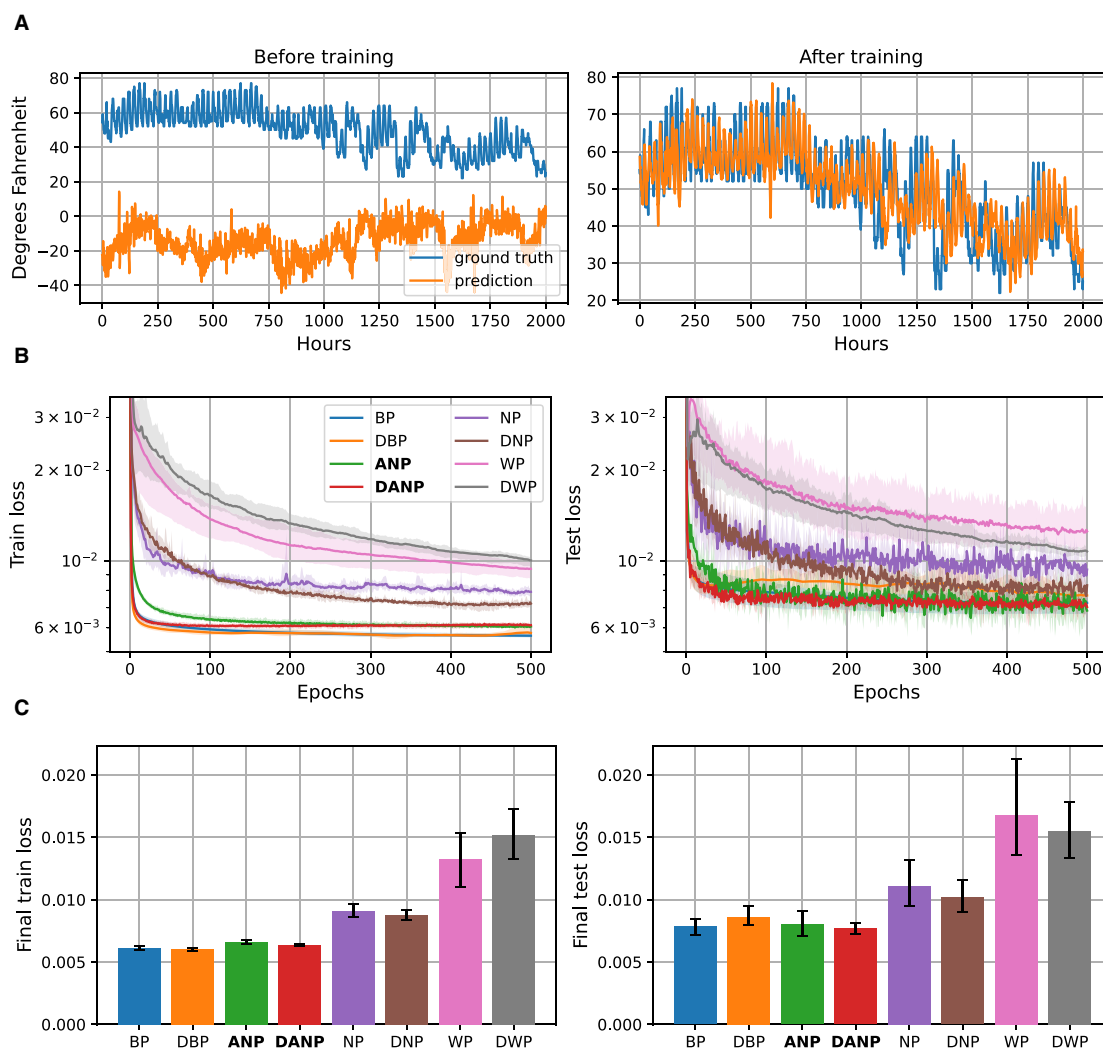


FIGURE 6 48-h ahead weather prediction data and results. **(A)** 2,000 time steps of the target feature, 'Dry bulb', from the weather dataset along with the predictions 48 hours ahead of a BP-trained model before and after training. **(B)** Performance during training over the train and test set for the different methods, represented in a logarithmic scale. **(C)** Final performance for the different methods, computed as the mean performance over the last 50 epochs.

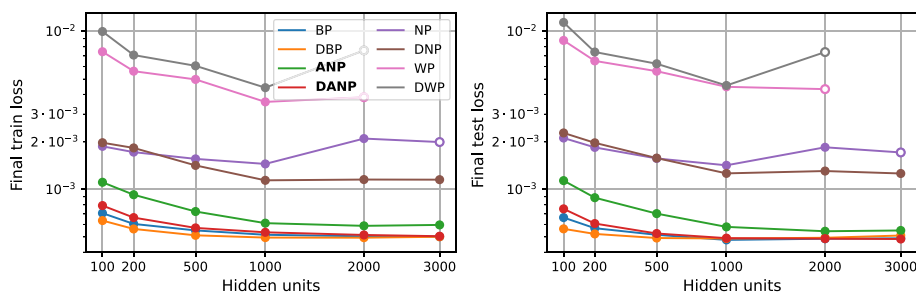


FIGURE 7 Scalability. Panels show final train and test loss. Each data point represents the mean of five runs. Solid circles denote stability across all runs, with the mean calculated from the five stable runs. Empty circles indicate instability in some runs, and the mean is computed solely from the remaining stable executions. The absence of data points indicates that all the runs were unstable.

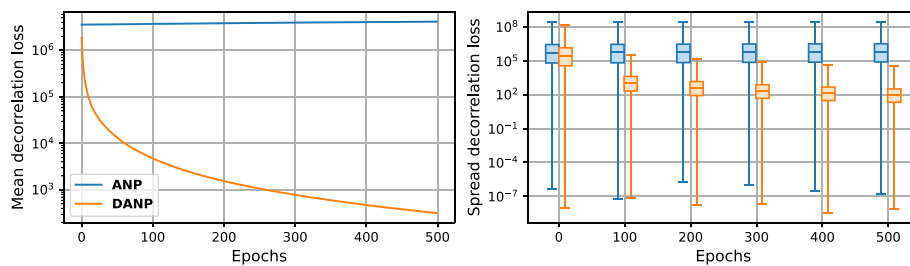


FIGURE 8

Decorrelation loss during training. The lower triangular correlation matrix is computed with the hidden units during the input presentation and subsequently squared and averaged across epochs. On the left, the mean decorrelation loss is depicted. On the right, a boxplot showing the spread of the decorrelation loss values is shown. Statistics were generated during training on the 1-h ahead weather prediction task.

4 Discussion

Backpropagation through time is the default algorithm to train RNNs. Like backpropagation, it relies on the computation and propagation of gradients for weight updates. However, the need to unroll the RNN over time makes this algorithm computationally demanding and memory-intensive, especially with long input sequences. Furthermore, the non-local nature of its updates and the requirement to backtrack through time once the input sequence concludes can pose some challenges in implementing this training method. In this paper, we introduced a viable alternative for training RNNs. Instead of using explicit gradients, our method approximates the gradients stochastically via perturbation-based learning. To this end, we extended the decorrelated activity-based node perturbation approach (Dalm et al., 2023), which approximates stochastic gradient descent more accurately than other perturbation-based methods, to operate effectively in the time domain using RNNs.

The results of our extensive validation show similar performance, convergence time and scalability of ANP and DANP when compared to BP. Remarkably, our approach exhibits superior performance, convergence time and scalability compared to standard perturbation-based methods, NP and WP. These promising findings suggest that perturbation-based learning holds the potential to rival BP's performance while retaining inherent advantages over gradient-based learning. Notably, the computational simplicity of our method stands out, especially compared to BP, which requires a specific second phase for error computation across networks and a dedicated circuit for error propagation. This simplicity is particularly advantageous for RNNs, eliminating the need for time-unrolling that significantly increases computational load. Additional experiments, detailed in [Supplementary materials 3](#), empirically validate our claims that perturbation-based methods are computationally simpler and less memory-intensive than the standard gradient-based backpropagation method. Note that other gradient-free learning methods for RNNs have been proposed that do not require separate learning passes, such as the random feedback local online (RFLO) algorithm (Murray, 2019). In [Supplementary materials 4](#) we show that our method still compares favorably in terms of learning efficiency and computational cost. In follow-up work, further validation on very large neural networks is necessary. This step is usually a common challenge in assessing the efficacy of alternative learning algorithms (Bartunov et al., 2018). Other avenues for

further research include exploring performance in more complex recurrent structures like LSTM and GRU units, as well as further validation on other challenging real-world datasets.

Our extensions differ from typical approaches in how the reinforcement signal is computed. While conventional extensions considers a global loss or aggregate local losses over time into a single loss per sequence and use it as the reinforcement signal, our methods directly incorporate local losses in time into the updates as the reinforcement signal. This allows for more informed and effective updates without increasing the computational load, considering peaks of high or low performance at specific time steps. Furthermore, both the clean and the noisy forward pass can be parallelized in this approach by using two identical copies of the model, making it compatible with online learning setups. This enables our synaptic rules to compute and implement updates online at every time step, reducing memory requirements as updates are applied immediately without storage. In this sense, our approach resembles the RTRL algorithm for computing gradients in RNNs in a forward manner (Williams and Zipser, 1989; Zenke and Nefcici, 2020). A limitation of our approach using local losses is its limited compatibility with delayed or sparse rewards over time. However, as demonstrated by our own results, as well as those of Zenke and Nefcici (2020) and Bellec et al. (2020), local methods can still work well in such delayed settings.

Our approach has several limitations regarding biological fidelity. Our neural network models are significantly simplified, lacking spikes, structural plasticity, and detailed temporal dynamics and neural structure. However, our results are an important step in developing more biologically realistic learning approaches, while performing similarly to backpropagation, which is challenging to implement in biological circuits (Whittington and Bogacz, 2019; Lillicrap et al., 2020). Numerous methods have been proposed for learning synaptic weights in artificial neural networks as alternatives, such as feedback alignment (Lillicrap et al., 2016; Nøkland, 2016), target propagation (Bengio, 2014, 2020; Lee et al., 2015; Ahmad et al., 2020), dendritic error propagation (Guerguiev et al., 2017; Sacramento et al., 2018), and spiking implementations combining global and local errors (Bellec et al., 2020). However, these approaches generally still require computing parameter or neuron-specific errors, necessitating specific and complex circuits for error propagation. In contrast, our approach uses global errors akin to neuromodulatory signals in biological neural networks (Schultz, 1998; Doya, 2002; Marder, 2012; Brzosko et al., 2019). These global errors are uniformly spread across the network,

eliminating the need for pathways to compute and deliver specific errors to each neuron. A key distinction of perturbation-based learning approaches, like ours, is their use of intrinsic brain noise for synaptic plasticity (Faisal et al., 2008), viewing noise as a beneficial feature rather than an obstacle. This mirrors the biological principle of using noise as a mechanism for learning, crucial for adapting to dynamic and unpredictable environments. Finally, our approach's active decorrelation method is similar to decorrelation in the brain, thought to be implemented through neural inhibition (Ecker et al., 2010; Chini et al., 2022). This process reduces correlations among neurons, promotes efficient information coding, enhances learning efficiency, and mirrors the adaptive processes found in biological neural networks.

While not fully biologically realistic, our models have sufficient fidelity to be highly relevant for neuromorphic computation. The local nature of the required computations, combined with a global learning signal, facilitates the deployment of these methods on neuromorphic hardware (Sandamirskaya et al., 2022; Paredes-Vallés et al., 2024), devices that rely on distributed, localized processing units that mimic biological neurons. Embracing noise as a mechanism for learning can be highly suitable in settings where the computational substrate shows a high degree of noise (Gokmen, 2021). This even holds when the noise cannot be measured since ANP still functions when we compare two noisy passes rather than a clean and a noisy pass (Dalm et al., 2023). This resilience to noise underscores the method's suitability for real-world neuromorphic devices where noise is ubiquitous and difficult to control. Additionally, the gradient-free nature of the approach becomes valuable in settings where the computational graph contains non-differentiable components as in spiking recurrent neural networks; a type of network where effective training methods are still under exploration (Neftci et al., 2019; Tavanaei et al., 2019; Wang et al., 2020).

Concluding, our present findings open the door to efficient gradient-free training of RNNs, offering exciting prospects for future research and applications in artificial intelligence, neuroscience and neuromorphic computing.

Data availability statement

The original contributions presented in the study are included in the article/Supplementary material, further inquiries can be directed to the corresponding author.

References

- Ahmad, N., Schrader, E., and van Gerven, M. (2022). Constrained parameter inference as a principle for learning. *arXiv [preprint]* arXiv:2203.13203. doi: 10.48550/arxiv.2203.13203
- Ahmad, N., van Gerven, M. A., and Ambrogioni, L. (2020). Gait-prop: A biologically plausible learning rule derived from backpropagation of error. *Adv. Neural Inf. Process. Syst.* 33, 10913–10923. doi: 10.48550/arxiv.2006.0643
- Arjovsky, M., Shah, A., and Bengio, Y. (2016). "Unitary evolution recurrent neural networks," in *International Conference on Machine Learning* (New York: PMLR), 1120–1128.
- Bartunov, S., Santoro, A., Richards, B., Marris, L., Hinton, G. E., and Lillicrap, T. (2018). "Assessing the scalability of biologically-motivated deep learning algorithms and architectures," in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 31.
- Baydin, A. G., Pearlmutter, B. A., Syme, D., Wood, F., and Torr, P. (2022). Gradients without backpropagation. *arXiv [preprint]* arXiv:2202.08587. doi: 10.48550/arxiv.2202.08587
- Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., et al. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nat. Commun.* 11, 1–15. doi: 10.1038/s41467-020-17236-y
- Bengio, Y. (2014). How auto-encoders could provide credit assignment in deep networks via target propagation. *arXiv [preprint]* arXiv:1407.7906. doi: 10.48550/arxiv.1407.7906

Author contributions

JG: Methodology, Software, Writing – original draft. SK: Supervision, Writing – review & editing. MvG: Conceptualization, Supervision, Writing – review & editing.

Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This publication is part of the DBI2 project (024.005.022, Gravitation), which is financed by the Dutch Ministry of Education (OCW) via the Dutch Research Council (NWO).

Acknowledgments

We would like to thank Nasir Ahmad for helpful discussions.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2024.1439155/full#supplementary-material>

- Bengio, Y. (2020). Deriving differential target propagation from iterating approximate inverses. *arXiv [preprint]* arXiv:2007.15139. doi: 10.48550/arxiv.2007.15139
- Bengio, Y., Simard, P., and Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Netw.* 5, 157–166. doi: 10.1109/72.279181
- Brzosko, Z., Mierau, S. B., and Paulsen, O. (2019). Neuromodulation of spike-timing-dependent plasticity: past, present, and future. *Neuron* 103, 563–581. doi: 10.1016/j.neuron.2019.05.041
- Cauwenberghs, G. (1992). “A fast stochastic error-descent algorithm for supervised learning and optimization,” in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 5.
- Cayco-Gajic, N. A., Clopath, C., and Silver, R. A. (2017). Sparse synaptic connectivity is required for decorrelation and pattern separation in feedforward networks. *Nat. Commun.* 8:1116. doi: 10.1038/s41467-017-01109-y
- Chini, M., Pfeffer, T., and Hanganu-Opatz, I. (2022). An increase of inhibition drives the developmental decorrelation of neural activity. *Elife* 11:e78811. doi: 10.7554/eLife.78811
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., et al. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv [preprint]* arXiv:1406.1078. doi: 10.3115/v1/D14-1179
- Chung, S., and Siegelmann, H. (2021). Turing completeness of bounded-precision recurrent neural networks. *Adv. Neural Inf. Process. Syst.* 34, 28431–28441.
- Dalm, S., Offergeld, J., Ahmad, N., and van Gerven, M. (2024). Efficient deep learning with decorrelated backpropagation. *arXiv [preprint]* arXiv:2405.02385. doi: 10.48550/arxiv.2405.02385
- Dalm, S., van Gerven, M., and Ahmad, N. (2023). Effective learning with node perturbation in deep neural networks. *arXiv [preprint]* arXiv:2310.00965.
- Desjardins, G., Simonyan, K., Pascanu, R., and Kavukcuoglu, K. (2015). “Natural neural networks,” in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 28.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Netw.* 15, 495–506. doi: 10.1016/S0893-6080(02)00044-8
- Ecker, A. S., Berens, P., Keliris, G. A., Bethge, M., Logothetis, N. K., and Tolias, A. S. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science* 327, 584–587. doi: 10.1126/science.1179867
- Faisal, A. A., Selen, L. P., and Wolpert, D. M. (2008). Noise in the nervous system. *Nature Rev. Neurosci.* 9, 292–303. doi: 10.1038/nrn2258
- Fiete, I. R., and Seung, H. S. (2006). Gradient learning in spiking neural networks by dynamic perturbation of conductances. *Phys. Rev. Lett.* 97, 048104. doi: 10.1103/PhysRevLett.97.048104
- Gokmen, T. (2021). Enabling training of neural networks on noisy hardware. *Front. Artif. Intellig.* 4:699148. doi: 10.3389/frac.2021.699148
- Guerguiev, J., Lillicrap, T. P., and Richards, B. A. (2017). Towards deep learning with segregated dendrites. *Elife* 6:e22901. doi: 10.7554/eLife.22901
- Hewamalage, H., Bergmeir, C., and Bandara, K. (2021). Recurrent neural networks for time series forecasting: current status and future directions. *Int. J. Forecast.* 37, 388–427. doi: 10.1016/j.ijforecast.2020.06.008
- Hiratani, N., Mehta, Y., Lillicrap, T., and Latham, P. E. (2022). On the stability and scalability of node perturbation learning. *Adv. Neural Inf. Process. Syst.* 35, 31929–31941.
- Huang, L., Yang, D., Lang, B., and Deng, J. (2018). “Decorrelated batch normalization,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 791–800.
- Kaspar, C., Ravoo, B., van der Wiel, W. G., Wegner, S., and Pernice, W. (2021). The rise of intelligent matter. *Nature* 594, 345–355. doi: 10.1038/s41586-021-03453-y
- Kingma, D. P., and Ba, J. (2014). Adam: a method for stochastic optimization. *arXiv [preprint]* arXiv:1412.6980. doi: 10.48550/arXiv.1412.6980
- Lansdell, B. J., Prakash, P. R., and Kording, K. P. (2019). Learning to solve the credit assignment problem. *arXiv [preprint]* arXiv:1906.00889.
- Lee, D.-H., Zhang, S., Fischer, A., and Bengio, Y. (2015). “Difference target propagation,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (Cham: Springer), 498–515.
- Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. (2016). Random synaptic feedback weights support error backpropagation for deep learning. *Nat. Commun.* 7, 1–10. doi: 10.1038/ncomms13276
- Lillicrap, T. P., and Santoro, A. (2019). Backpropagation through time and the brain. *Curr. Opin. Neurobiol.* 55, 82–89. doi: 10.1016/j.conb.2019.01.011
- Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., and Hinton, G. (2020). Backpropagation and the brain. *Nature Rev. Neurosci.* 21, 335–346. doi: 10.1038/s41583-020-0277-3
- Luo, P. (2017). “Learning deep architectures via generalized whitened neural networks,” in *International Conference on Machine Learning* (New York: PMLR), 2238–2246.
- Mackey, M. C., and Glass, L. (1977). Oscillation and chaos in physiological control systems. *Science* 197, 287–289. doi: 10.1126/science.267326
- Marder, E. (2012). Neuromodulation of neuronal circuits: back to the future. *Neuron* 76, 1–11. doi: 10.1016/j.neuron.2012.09.010
- Murray, J. M. (2019). Local online learning in recurrent networks with random feedback. *Elife* 8:e43299. doi: 10.7554/eLife.43299
- Neftci, E. O., Mostafa, H., and Zenke, F. (2019). Surrogate gradient learning in spiking neural networks: bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Process. Mag.* 36, 51–63. doi: 10.1109/MSP.2019.2931595
- Nóklánd, A. (2016). “Direct feedback alignment provides learning in deep neural networks,” in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 29.
- Orvieto, A., Smith, S. L., Gu, A., Fernando, A., Gulcehre, C., Pascanu, R., et al. (2023). “Resurrecting recurrent neural networks for long sequences,” in *International Conference on Machine Learning* (New York: PMLR), 26670–26698.
- Paredes-Vallés, F., Hagenaars, J. J., Dupeyron, J., Stroobants, S., Xu, Y., and Croon, G. C. H. E. D. (2024). Fully neuromorphic vision and control for autonomous drone flight. *Sci. Robot.* 9:adi0591. doi: 10.1126/scirobotics.adi0591
- Pascanu, R., Mikolov, T., and Bengio, Y. (2013). “On the difficulty of training recurrent neural networks,” in *International Conference on Machine Learning* (New York: PMLR), 1310–1318.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). “Pytorch: an imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 32.
- Sacramento, J., Ponte Costa, R., Bengio, Y., and Senn, W. (2018). “Dendritic cortical microcircuits approximate the backpropagation algorithm,” in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 31.
- Sandamirskaya, Y., Kaboli, M., Conrad, J., and Celikel, T. (2022). Neuromorphic computing hardware and neural architectures for robotics. *Sci. Robot.* 7:abl8419. doi: 10.1126/scirobotics.abl8419
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27. doi: 10.1152/jn.1998.80.1.1
- Schuman, C. D., Kulkarni, S. R., Parsa, M., Mitchell, J. P., Kay, B., et al. (2022). Opportunities for neuromorphic computing algorithms and applications. *Nature Comp. Sci.* 2, 10–19. doi: 10.1038/s43588-021-00184-y
- Spall, J. C. (1992). Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Trans. Automat. Contr.* 37, 332–341. doi: 10.1109/9.119632
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). “Sequence to sequence learning with neural networks,” in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 27.
- Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., and Maida, A. (2019). Deep learning in spiking neural networks. *Neural Netw.* 111, 47–63. doi: 10.1016/j.neunet.2018.12.002
- Tetzlaff, T., Helias, M., Einevoll, G. T., and Diesmann, M. (2012). Decorrelation of neural-network activity by inhibitory feedback. *PLoS Comput. Biol.* 8(8). doi: 10.1371/journal.pcbi.1002596
- Van Rossum, G., Drake, F. L., et al. (1995). *Python Reference Manual, Volume 111*. Amsterdam: Centrum voor Wiskunde en Informatica Amsterdam.
- Voelker, A., Kajić, I., and Eliasmith, C. (2019). “Legendre memory units: Continuous-time representation in recurrent neural networks,” in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 32.
- Wang, X., Lin, X., and Dang, X. (2020). Supervised learning in spiking neural networks: a review of algorithms and evaluations. *Neural Netw.* 125, 258–280. doi: 10.1016/j.neunet.2020.02.011
- Werbos, P. J. (1990). Backpropagation through time: what it does and how to do it. *Proc. IEEE* 78, 1550–1560. doi: 10.1109/5.58337
- Werfel, J., Xie, X., and Seung, H. (2003). “Learning curves for stochastic gradient descent in linear feedforward networks,” in *Advances in Neural Information Processing Systems* (Vancouver, BC: Neural Information Processing Systems Foundation), 16.
- Whittington, J. C., and Bogacz, R. (2019). Theories of error backpropagation in the brain. *Trends Cogn. Sci.* 23, 235–250. doi: 10.1016/j.tics.2018.12.005
- Widrow, B., and Lehr, M. A. (1990). 30 years of adaptive neural networks: perceptron, madaline, and backpropagation. *Proc. IEEE* 78, 1415–1442. doi: 10.1109/5.58323

Wiechert, M. T., Judkewitz, B., Riecke, H., and Friedrich, R. W. (2010). Mechanisms of pattern decorrelation by recurrent neuronal circuits. *Nat. Neurosci.* 13, 1003–1010. doi: 10.1038/nn.2591

Williams, R. J., and Zipser, D. (1989). A learning algorithm for continually running fully recurrent neural networks. *Neural Comput.* 1, 270–280. doi: 10.1162/neco.1989.1.2.270

Yao, K., Zweig, G., Hwang, M.-Y., Shi, Y., and Yu, D. (2013). Recurrent neural networks for language understanding. In *Interspeech* 2524–2528. doi: 10.21437/Interspeech.2013-569

Zenke, F., and Nefci, E. O. (2020). Brain-inspired learning on neuromorphic substrates. *arXiv [preprint]* arXiv:2010.11931.

Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., et al. (2021). Informer: beyond efficient transformer for long sequence time-series forecasting. *Proc. AAAI Conf. Artif. Intellig.* 35, 11106–11115. doi: 10.1609/aaai.v35i12.17325

Züge, P., Klos, C., and Memmesheimer, R.-M. (2023). Weight versus node perturbation learning in temporally extended tasks: weight perturbation often performs similarly or better. *Phys. Rev. X* 13:021006. doi: 10.1103/PhysRevX.13.021006