

# How Does Momentum Benefit Deep Neural Networks Architecture Design? A Few Case Studies

Bao Wang · Hedi Xia · Tan Nguyen · Stanley Osher

Received: date / Accepted: date

**Abstract** We present and review an algorithmic and theoretical framework for improving neural network architecture design via momentum. As case studies, we consider how momentum can improve the architecture design for recurrent neural networks (RNNs), neural ordinary differential equations (ODEs), and transformers. We show that integrating momentum into neural network architectures has several remarkable theoretical and empirical benefits, including 1) integrating momentum into RNNs and neural ODEs can overcome the vanishing gradient issues in training RNNs and neural ODEs, resulting in effective learning long-term dependencies. 2) momentum in neural ODEs can reduce the stiffness of the ODE dynamics, which significantly enhances the computational efficiency in training and testing. 3) momentum can improve the efficiency and accuracy of transformers.

## 1 Introduction

Deep learning has radically advanced artificial intelligence [48], which has achieved state-of-the-art performance in various applications, including com-

---

Bao Wang  
Department of Mathematics  
Scientific Computing and Imaging Institute, University of Utah  
E-mail: wangbaonj@gmail.com

Hedi Xia  
Department of Mathematics, UCLA  
E-mail: hedixia@math.ucla.edu

Tan Nguyen  
Department of Mathematics, UCLA  
E-mail: tannguyen89@math.ucla.edu

Stanley Osher  
Department of Mathematics, UCLA  
E-mail: sjo@math.ucla.edu

puter vision [22, 96], natural language processing [21, 8], and control [85]. Nevertheless, deep neural networks (DNNs) designs are mostly heuristic and the resulting architectures have many well-known issues: 1) convolutional neural networks (CNNs) are not robust to unperceptible adversarial attacks [89]; 2) recurrent neural networks (RNNs) cannot learn long-term dependencies effectively due to vanishing gradients [67]; 3) training neural ordinary differential equations (ODEs) can take an excessive number of function evaluations (NFEs) [13]; and 4) training transformers suffers from quadratic computational time and memory costs [98]. See Sections 2, 3, and 4, respectively, for the details of these problems 2)-4).

Addressing the above grand challenges is at the forefront of deep learning research. 1) Adversarial defense [55] has been proposed to train robust neural networks against adversarial attacks; a survey of adversarial defense algorithms is available, see, e.g., [11]. Training ResNets has also been interpreted as solving a control problem of the transport equation [102, 106], resulting in PDE-motivated adversarial defenses [106, 101, 104]. 2) Learning long-term dependencies with improved RNNs has been an active research area for several decades; celebrated works include long short-term memory [36] and gated recurrent unit [17]. 3) Several algorithms and techniques have been proposed to reduce NFEs in training neural ODEs, including input augmentation [24], regularizing solvers and learning dynamics [26, 40, 28, 64], high-order ODE [63], data control [56], and depth-variance [56]. 4) Transformers are the current state-of-the-art machine learning (ML) models for sequential learning [98], which processes the input sequence concurrently and can learn long-term dependencies effectively. However, transformers suffer from quadratic computational time and memory costs with respect to the input sequence length; see Section 4 for details. In response, efficient attention has been proposed leveraging sparse and low-rank approximation of the attention matrix [53, 66, 4, 1, 112, 107, 39, 16], locality-sensitive hashing [44], clustered attention [100], and decomposed near-field and far-field attention [61].

### 1.1 Background: Momentum acceleration for gradient descent

Let us recall heavy-ball momentum, a.k.a. classical momentum [69], for accelerating gradient descent in solving  $\min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x})$ . Starting from  $\mathbf{x}^0$  and  $\mathbf{x}^1$ , the heavy-ball method iterates as follows

$$\mathbf{x}^{k+1} = \mathbf{x}^k - s\nabla F(\mathbf{x}^k) + \beta(\mathbf{x}^k - \mathbf{x}^{k-1}), \quad (1)$$

where  $s > 0$  is the step size and  $0 \leq \beta < 1$  is the momentum parameter. By introducing the momentum state  $\mathbf{m}$ , we can rewrite the HB method as

$$\mathbf{m}^{k+1} = \beta\mathbf{m}^k + \nabla F(\mathbf{x}^k); \quad \mathbf{x}^{k+1} = \mathbf{x}^k - s\mathbf{m}^{k+1}. \quad (2)$$

In contrast, gradient descent updates at each step as follows

$$\mathbf{x}^{k+1} = \mathbf{x}^k - s\nabla F(\mathbf{x}^k). \quad (3)$$

## 1.2 Contribution

This paper aims to present and review an algorithmic and theoretical framework for improving neural network architecture design via momentum, a well-established first-order optimization tool [69,59]. In particular, we focus on leveraging momentum to design new RNNs and neural ODEs to accelerate their training and testing and improve learning long-term dependencies with theoretical guarantees. Moreover, we present a new efficient attention mechanism with momentum augmentation, which significantly improves computational efficiency over transformers [98] and accuracy over linear transformers [39]. Finally, we present some perspectives of how momentum can further improve neural networks design and solve existing grand challenges.

## 1.3 Notations


We denote scalars by lower- or upper-case letters. We also denote vectors and matrices by lower- and upper-case boldface letters, respectively. For a vector  $\mathbf{x} = (x_1, \dots, x_d)^\top \in \mathbb{R}^d$ , where  $(x_1, \dots, x_d)^\top$  denotes the transpose of the vector  $(x_1, \dots, x_d)$ , we use  $\|\mathbf{x}\| = (\sum_{i=1}^d |x_i|^2)^{1/2}$  to denote its  $\ell_2$  norm. We denote the vector whose entries are all 0s as  $\mathbf{0}$ . For a matrix  $\mathbf{A}$ , we use  $\mathbf{A}^\top$ ,  $\mathbf{A}^{-1}$ , and  $\|\mathbf{A}\|$  to denote its transpose, inverse, and spectral norm, respectively. We use  $\mathbf{I}$  to denote the identity matrix, whose dimension can be determined in its context. For a function  $f(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ , we denote its gradient as  $\nabla f(\mathbf{x})$ . Given two sequences  $\{a_n\}$  and  $\{b_n\}$ , we write  $a_n = \mathcal{O}(b_n)$  if there exists a positive constant  $0 < C < +\infty$  such that  $a_n \leq Cb_n$ .

## 1.4 Organization

We organize this paper as follows: In Section 2, we present RNN models and their difficulties in learning long-term dependencies. We also show how to integrate momentum into RNNs to accelerate training RNNs and enhance RNNs' capability in learning long-term dependencies. In Section 3, we show how the ODE limit of momentum can improve neural ODEs in terms of training and test efficiency and learning long-term dependencies. In Section 4, we show how momentum can be integrated into efficient transformers and improve their accuracy. We conclude and present potential new directions in Section 5.

## 2 Recurrent Neural Networks

In this section, we present how to leverage momentum to improve RNN architecture design. The main results have been presented at NeurIPS 2020 [60].



**Fig. 1** Illustration of the recurrent cell (left), Momentum/NAG cell (middle), and Adam/RMSProp cell (right). We draw a connection between the dynamics of hidden states in the recurrent cell and GD. We then introduce momentum to recurrent cell as an analogy of the momentum accelerated GD.

## 2.1 Recap on RNNs and LSTM

Recurrent cells are the building blocks of RNNs. A recurrent cell employs a cyclic connection to update the current hidden state ( $\mathbf{h}_t$ ) using the past hidden state ( $\mathbf{h}_{t-1}$ ) and the current input data ( $\mathbf{x}_t$ ) [25]; the dependence of  $\mathbf{h}_t$  on  $\mathbf{h}_{t-1}$  and  $\mathbf{x}_t$  in a recurrent cell can be written as

$$\mathbf{h}_t = \sigma(\mathbf{U}\mathbf{h}_{t-1} + \mathbf{W}\mathbf{x}_t + \mathbf{b}), \quad \mathbf{x}_t \in \mathbb{R}^d, \quad \text{and } \mathbf{h}_{t-1}, \mathbf{h}_t \in \mathbb{R}^h, \quad t = 1, 2, \dots, T, \quad (4)$$

where  $\mathbf{U} \in \mathbb{R}^{h \times h}$ ,  $\mathbf{W} \in \mathbb{R}^{h \times d}$ , and  $\mathbf{b} \in \mathbb{R}^h$  are trainable parameters;  $\sigma(\cdot)$  is a nonlinear activation function, e.g., sigmoid or hyperbolic tangent; see Fig. 1 (left) for an illustration of the RNN cell. Error backpropagation through time is used to train RNN, but it tends to result in exploding or vanishing gradients [5]. Thus RNNs may fail to learn long term dependencies.

LSTM cells augment the recurrent cell with “gates” [36] and results in

$$\begin{aligned} \mathbf{i}_t &= \sigma(\mathbf{U}_{ih}\mathbf{h}_{t-1} + \mathbf{W}_{ix}\mathbf{x}_t + \mathbf{b}_i), & (\mathbf{i}_t : \text{input gate}) \\ \tilde{\mathbf{c}}_t &= \tanh(\mathbf{U}_{\tilde{c}h}\mathbf{h}_{t-1} + \mathbf{W}_{\tilde{c}x}\mathbf{x}_t + \mathbf{b}_{\tilde{c}}), & (\tilde{\mathbf{c}}_t : \text{cell input}) \\ \mathbf{c}_t &= \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t, & (\mathbf{c}_t : \text{cell state}) \\ \mathbf{o}_t &= \sigma(\mathbf{U}_{oh}\mathbf{h}_{t-1} + \mathbf{W}_{ox}\mathbf{x}_t + \mathbf{b}_o), & (\mathbf{o}_t : \text{output gate}) \\ \mathbf{h}_t &= \mathbf{o}_t \odot \tanh \mathbf{c}_t, & (\mathbf{h}_t : \text{hidden state}) \end{aligned} \quad (5)$$

where  $\mathbf{U}_* \in \mathbb{R}^{h \times h}$ ,  $\mathbf{W}_* \in \mathbb{R}^{h \times d}$ , and  $\mathbf{b}_* \in \mathbb{R}^h$  are learnable parameters, and  $\odot$  denotes the Hadamard product. The input gate decides what new information to be stored in the cell state, and the output gate decides what information to output based on the cell state value. The gating mechanism in LSTMs can lead to the issue of saturation [97, 12].

## 2.2 Gradient descent analogy for RNN and MomentumRNN

Now, we are going to establish a connection between RNN and gradient descent, and further leverage momentum to improve RNNs. Let  $\widetilde{\mathbf{W}} = [\mathbf{W}, \mathbf{b}]$  and  $\widetilde{\mathbf{x}}_t = [\mathbf{x}_t, 1]^\top$  in (4), then we have  $\mathbf{h}_t = \sigma(\mathbf{U}\mathbf{h}_{t-1} + \widetilde{\mathbf{W}}\widetilde{\mathbf{x}}_t)$ . For the ease of notation, without ambiguity we denote  $\mathbf{W} := \widetilde{\mathbf{W}}$  and  $\mathbf{x}_t := \widetilde{\mathbf{x}}_t$ . Then the recurrent cell can be reformulated as

$$\mathbf{h}_t = \sigma(\mathbf{U}\mathbf{h}_{t-1} + \mathbf{W}\mathbf{x}_t). \quad (6)$$

Moreover, let  $\phi(\cdot) := \sigma(\mathbf{U}(\cdot))$  and  $\mathbf{u}_t := \mathbf{U}^{-1}\mathbf{W}\mathbf{x}_t$ , we can rewrite (6) as

$$\mathbf{h}_t = \phi(\mathbf{h}_{t-1} + \mathbf{u}_t). \quad (7)$$

If we regard  $-\mathbf{u}_t$  as the ‘‘gradient’’ at the  $t$ -th iteration, then we can consider (7) as the dynamical system which updates the hidden state by the gradient and then transforms the updated hidden state by the nonlinear activation function  $\phi$ . We propose the following accelerated dynamical system to accelerate the dynamics of (7), which is principled by the accelerated gradient descent theory (see subsection 1.1):

$$\mathbf{p}_t = \mu\mathbf{p}_{t-1} - s\mathbf{u}_t; \quad \mathbf{h}_t = \phi(\mathbf{h}_{t-1} - \mathbf{p}_t), \quad (8)$$

where  $\mu \geq 0, s > 0$  are two hyperparameters, which are the analogies of the momentum coefficient and step size in the momentum-accelerated GD, respectively. Let  $\mathbf{v}_t := -\mathbf{U}\mathbf{p}_t$ , we arrive at the following dynamical system:

$$\mathbf{v}_t = \mu\mathbf{v}_{t-1} + s\mathbf{W}\mathbf{x}_t; \quad \mathbf{h}_t = \sigma(\mathbf{U}\mathbf{h}_{t-1} + \mathbf{v}_t). \quad (9)$$

The architecture of the momentum cell that corresponds to the dynamical system (9) is plotted in Fig. 1 (middle). Compared with the recurrent cell, the momentum cell introduces an auxiliary momentum state in each update and scales the dynamical system with two positive hyperparameters  $\mu$  and  $s$ .


*Remark 1* Different parameterizations of (8) can result in different momentum cell architectures. For instance, if we let  $\mathbf{v}_t = -\mathbf{p}_t$ , we end up with the following dynamical system:

$$\mathbf{v}_t = \mu\mathbf{v}_{t-1} + s\widehat{\mathbf{W}}\mathbf{x}_t; \quad \mathbf{h}_t = \sigma(\mathbf{U}\mathbf{h}_{t-1} + \mathbf{U}\mathbf{v}_t), \quad (10)$$

where  $\widehat{\mathbf{W}} := \mathbf{U}^{-1}\mathbf{W}$  is the trainable weight matrix. Even though (9) and (10) are mathematically equivalent, the training procedure might cause the MomentumRNNs that are derived from different parameterizations to have different performances.

*Remark 2* We put the nonlinear activation in the second equation of (8) to ensure that the value of  $\mathbf{h}_t$  is in the same range as the original recurrent cell.

*Remark 3* The derivation above also applies to the dynamical systems in the LSTM cells, and we can design the MomentumLSTM in the same way as designing the MomentumRNN.



**Fig. 2**  $\ell_2$  norm of the gradients of the loss  $\mathcal{L}$  w.r.t. the state vector  $\mathbf{h}_t$  at each time step  $t$  for RNN (left) and MomentumRNN (right). MomentumRNN does not suffer from vanishing gradients.

### 2.3 Analysis of the vanishing gradient: Momentum cell vs. Recurrent cell

Let  $\mathbf{h}_T$  and  $\mathbf{h}_t$  be the state vectors at the time step  $T$  and  $t$ , respectively, and we suppose  $T \gg t$ . Furthermore, assume that  $\mathcal{L}$  is the objective to minimize, then

$$\frac{\partial \mathcal{L}}{\partial \mathbf{h}_t} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \cdot \frac{\partial \mathbf{h}_T}{\partial \mathbf{h}_t} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \cdot \prod_{k=t}^{T-1} \frac{\partial \mathbf{h}_{k+1}}{\partial \mathbf{h}_k} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \cdot \prod_{k=t}^{T-1} (\mathbf{D}_k \mathbf{U}^\top), \quad (11)$$

where  $\mathbf{U}^\top$  is the transpose of  $\mathbf{U}$  and  $\mathbf{D}_k = \text{diag}(\sigma'(\mathbf{U}\mathbf{h}_k + \mathbf{W}\mathbf{x}_{k+1}))$  is a diagonal matrix with  $\sigma'(\mathbf{U}\mathbf{h}_k + \mathbf{W}\mathbf{x}_{k+1})$  being its diagonal entries.  $\|\prod_{k=t}^{T-1} (\mathbf{D}_k \mathbf{U}^\top)\|_2$  tends to either vanish or explode [5]. We can use regularization or gradient clipping to mitigate the exploding gradient, leaving vanishing gradient as the major obstacle to training RNN to learn long-term dependency [67]. We can rewrite (9) as

$$\mathbf{h}_t = \sigma(\mathbf{U}(\mathbf{h}_{t-1} - \mu\mathbf{h}_{t-2}) + \mu\sigma^{-1}(\mathbf{h}_{t-1}) + s\mathbf{W}\mathbf{x}_t), \quad (12)$$

where  $\sigma^{-1}(\cdot)$  is the inverse function of  $\sigma(\cdot)$ . We compute  $\partial \mathcal{L} / \partial \mathbf{h}_t$  as follows

$$\frac{\partial \mathcal{L}}{\partial \mathbf{h}_t} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \cdot \frac{\partial \mathbf{h}_T}{\partial \mathbf{h}_t} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \cdot \prod_{k=t}^{T-1} \frac{\partial \mathbf{h}_{k+1}}{\partial \mathbf{h}_k} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \cdot \prod_{k=t}^{T-1} \hat{\mathbf{D}}_k [\mathbf{U}^\top + \mu\boldsymbol{\Sigma}_k], \quad (13)$$

where  $\hat{\mathbf{D}}_k = \text{diag}(\sigma'(\mathbf{U}(\mathbf{h}_k - \mu\mathbf{h}_{k-1}) + \mu\sigma^{-1}(\mathbf{h}_k) + s\mathbf{W}\mathbf{x}_{k+1}))$  and  $\boldsymbol{\Sigma} = \text{diag}((\sigma^{-1})'(\mathbf{h}_k))$ . For mostly used  $\sigma$ , e.g., sigmoid and tanh,  $(\sigma^{-1}(\cdot))' > 1$  and  $\mu\boldsymbol{\Sigma}_k$  dominates  $\mathbf{U}^\top$ .<sup>1</sup> Therefore, with an appropriate choice of  $\mu$ , the momentum cell can alleviate vanishing gradient and accelerate training.

We empirically corroborate that momentum cells can alleviate vanishing gradients by training a MomentumRNN and its corresponding RNN on the PMNIST classification task and plot  $\|\partial \mathcal{L} / \partial \mathbf{h}_t\|_2$  for each time step  $t$ . Figure 2 confirms that unlike in RNN, the gradients in MomentumRNN do not vanish.

<sup>1</sup> In the vanishing gradient scenario,  $\|\mathbf{U}\|_2$  is small; also it can be controlled by regularizing the loss function.

## 2.4 Beyond MomentumRNN: NAG and Adam principled RNNs

There are several other advanced formalisms of momentum existing in optimization, which can be leveraged for RNN architecture design. In this subsection, we present two additional variants of MomentumRNN that are derived from the Nesterov accelerated gradient (NAG)-style momentum with restart [59, 103] and Adam [43].

**NAG Principled RNNs.** The momentum-accelerated GD can be further accelerated by replacing the constant momentum coefficient  $\mu$  in (9) with the NAG-style momentum, i.e. setting  $\mu$  to  $(t-1)/(t+2)$  at the  $t$ -th iteration. Furthermore, we can accelerate NAG by resetting the momentum to 0 after every  $F$  iterations, i.e.  $\mu = (t \bmod F)/((t \bmod F) + 3)$ , which is the NAG-style momentum with a scheduled restart of the appropriately selected frequency  $F$  [103]. For convex optimization, NAG has a convergence rate  $O(1/t^2)$ , which is significantly faster than GD or GD with constant momentum whose convergence rate is  $O(1/t)$ . Scheduled restart not only accelerates NAG to a linear convergence rate  $O(\alpha^{-t})$  ( $0 < \alpha < 1$ ) under mild extra assumptions but also stabilizes the NAG iteration [103]. We call the MomentumRNN with the NAG-style momentum and scheduled restart momentum the NAG-based RNN and the scheduled restart RNN (SRRNN), respectively.

**Adam Principled RNNs.** Adam [43] leverages the moving average of historical gradients and entry-wise squared gradients to accelerate the stochastic gradient dynamics. We use Adam to accelerate (7) and end up with the following iteration

$$\begin{aligned} \mathbf{p}_t &= \mu \mathbf{p}_{t-1} + (1 - \mu) \mathbf{u}_t \\ \mathbf{m}_t &= \beta \mathbf{m}_{t-1} + (1 - \beta) \mathbf{u}_t \odot \mathbf{u}_t \\ \mathbf{h}_t &= \phi\left(\mathbf{h}_{t-1} - s \frac{\mathbf{p}_t}{\sqrt{\mathbf{r}_t} + \epsilon}\right), \end{aligned} \quad (14)$$

where  $\mu, s, \beta > 0$  are hyperparameters,  $\epsilon$  is a small constant and chosen to be  $10^{-8}$  by default, and  $\odot/\sqrt{\cdot}$  denotes the entrywise product/square root<sup>2</sup>. Again, let  $\mathbf{v}_t = -\mathbf{U}\mathbf{p}_t$ , we rewrite (14) as follows

$$\begin{aligned} \mathbf{v}_t &= \mu \mathbf{v}_{t-1} + (1 - \mu) \mathbf{W} \mathbf{x}_t \\ \mathbf{m}_t &= \beta \mathbf{m}_{t-1} + (1 - \beta) \mathbf{u}_t \odot \mathbf{u}_t \\ \mathbf{h}_t &= \sigma\left(\mathbf{U} \mathbf{h}_{t-1} + s \frac{\mathbf{v}_t}{\sqrt{\mathbf{m}_t} + \epsilon}\right). \end{aligned}$$

As before, here  $\mathbf{u}_t := \mathbf{U}^{-1} \mathbf{W} \mathbf{x}_t$ . Computing  $\mathbf{U}^{-1}$  is expensive. Our experiments suggest that replacing  $\mathbf{u}_t \odot \mathbf{u}_t$  by  $\mathbf{W} \mathbf{x}_t \odot \mathbf{W} \mathbf{x}_t$  is sufficient and more efficient to compute. In our implementation, we also relax  $\mathbf{v}_t = \mu \mathbf{v}_{t-1} + (1 - \mu) \mathbf{W} \mathbf{x}_t$  to  $\mathbf{v}_t = \mu \mathbf{v}_{t-1} + s \mathbf{W} \mathbf{x}_t$  that follows the momentum in the MomentumRNN (9) for better performance. Therefore, we propose the *AdamRNN* that

<sup>2</sup> In contrast to Adam, we do not normalize  $\mathbf{p}_t$  and  $\mathbf{m}_t$  since they can be absorbed in the weight matrices.

is given by

$$\begin{aligned} \mathbf{v}_t &= \mu \mathbf{v}_{t-1} + s \mathbf{W} \mathbf{x}_t \\ \mathbf{m}_t &= \beta \mathbf{m}_{t-1} + (1 - \beta) (\mathbf{W} \mathbf{x}_t \odot \mathbf{W} \mathbf{x}_t) \\ \mathbf{h}_t &= \sigma(\mathbf{U} \mathbf{h}_{t-1} + \frac{\mathbf{v}_t}{\sqrt{\mathbf{m}_t + \epsilon}}). \end{aligned} \quad (15)$$

In AdamRNN, if  $\mu$  is set to 0, we achieve another new RNN, which obeys the RMSProp gradient update rule [94]; which we call *RMSPropRNN*.

*Remark 4* Both AdamRNN and RMSPropRNN can also be derived by letting  $\mathbf{v}_t = -\mathbf{p}_t$  and  $\widehat{\mathbf{W}} := \mathbf{U}^{-1} \mathbf{W}$  as in Remark 1. This parameterization yields the following formulation for AdamRNN

$$\begin{aligned} \mathbf{v}_t &= \mu \mathbf{v}_{t-1} + s \widehat{\mathbf{W}} \mathbf{x}_t \\ \mathbf{m}_t &= \beta \mathbf{m}_{t-1} + (1 - \beta) (\widehat{\mathbf{W}} \mathbf{x}_t \odot \widehat{\mathbf{W}} \mathbf{x}_t) \\ \mathbf{h}_t &= \sigma(\mathbf{U} \mathbf{h}_{t-1} + \frac{\mathbf{U} \mathbf{v}_t}{\sqrt{\mathbf{m}_t + \epsilon}}). \end{aligned}$$

Here, we simply need to learn  $\widehat{\mathbf{W}}$  and  $\mathbf{U}$  without any relaxation. In contrast, we relaxed  $\mathbf{U}^{-1}$  to an identity matrix in (15). Our experiments suggest that both parameterizations yield similar results.

## 2.5 Experimental results

In this subsection, we evaluate the effectiveness of our momentum approach in designing RNNs in terms of convergence speed and accuracy. We compare the performance of the MomentumLSTM with the baseline LSTM [36] in the following tasks: 1) the object classification task on pixel-permuted MNIST [46], 2) the speech prediction task on the TIMIT dataset [3, 33, 109, 57, 34], 3) the celebrated copying and adding tasks [36, 3], and 4) the language modeling task on the Penn TreeBank (PTB) dataset [58]. These four tasks are among standard benchmarks to measure the performance of RNNs and their ability to handle long-term dependencies. Also, these tasks cover different data modalities – image, speech, and text data – as well as a variety of model sizes, ranging from thousands to millions of parameters with one (MNIST and TIMIT tasks) or multiple (PTB task) recurrent cells in concatenation. Our experimental results confirm that MomentumLSTM converges faster and yields better test accuracy than the baseline LSTM across tasks and settings. We also discuss the AdamLSTM, RMSPropLSTM, and scheduled restart LSTM (SRLSTM) and show their advantage over MomentumLSTM in specific tasks. All of our results are averaged over 5 runs with different seeds. For MNIST and TIMIT experiments, we use the baseline codebase provided by [9]. For PTB experiments, we use the baseline codebase provided by [79].



**Table 1** Best test accuracy at the MNIST and PMNIST tasks (%). We use the baseline results reported in [33], [109], [99]. Our proposed models outperform the baseline LSTM. Among the models using  $N = 256$  hidden units, RMSPropLSTM yields the best results.

MODEL	N	# PARAMS	MNIST	PMNIST
LSTM	128	$\approx 68K$	98.70[33],97.30 [99]	92.00 [33],92.62 [99]
LSTM	256	$\approx 270K$	98.90 [33], 98.50 [109]	92.29 [33], 92.10 [109]
MOMENTUMLSTM	128	$\approx 68K$	<b>99.04 <math>\pm</math> 0.04</b>	<b>93.40 <math>\pm</math> 0.25</b>
MOMENTUMLSTM	256	$\approx 270K$	<b>99.08 <math>\pm</math> 0.05</b>	<b>94.72 <math>\pm</math> 0.16</b>
ADAMLSTM	256	$\approx 270K$	99.09 $\pm$ 0.03	95.05 $\pm$ 0.37
RMSPROPLSTM	256	$\approx 270K$	<b>99.15 <math>\pm</math> 0.06</b>	<b>95.38 <math>\pm</math> 0.19</b>
SRLSTM	256	$\approx 270K$	99.01 $\pm$ 0.07	93.82 $\pm$ 1.85


### 2.5.1 Pixel-by-Pixel MNIST

In this task, we classify image samples of hand-written digits from the MNIST dataset [49] into one of the ten classes. Following the implementation of [46], we flatten the image of original size  $28 \times 28$  pixels and feed it into the model as a sequence of length 784. In the unpermuted task (MNIST), the sequence of pixels is processed row-by-row. In the permuted task (PMNIST), a fixed permutation is selected at the beginning of the experiments and then applied to both training and test sequences. We summarize the results in Table 1. Our experiments show that *MomentumLSTM achieves better test accuracy than the baseline LSTM in both MNIST and PMNIST digit classification tasks* using different numbers of hidden units (i.e.  $N = 128, 256$ ). Especially, the improvement is significant on the PMNIST task, which is designed to test the performance of RNNs in the context of long-term memory. Furthermore, we notice that *MomentumLSTM converges faster than LSTM* in all settings. Figure 3 corroborates this observation when using  $N = 256$  hidden units.

### 2.5.2 TIMIT speech dataset

We study how MomentumLSTM performs on audio data with speech prediction experiments on the TIMIT speech dataset [27], which is a collection of real-world speech recordings. As first proposed by [109], the recordings are downsampled to 8kHz and then transformed into log-magnitudes via a short-time Fourier transform (STFT). The task accounts for predicting the next log-magnitude given the previous ones. We use the standard train/validation/test separation in [109,51,10], thereby having 3640 utterances for the training set with a validation set of size 192 and a test set of size 400.

The results for this TIMIT speech prediction are shown in Table 2. Results are reported on the test set using the model parameters that yield the best validation loss. Again, we see the advantage of MomentumLSTM over the baseline LSTM. In particular, MomentumLSTM yields much better prediction



**Fig. 3** Train and test loss of MomentumLSTM, AdamLSTM, RMSPropLSTM, SRLSTM, and LSTM for MNIST (left two panels) and TIMIT (right two panels) tasks. MomentumLSTM converges faster than LSTM in both tasks. For MNIST, AdamLSTM and RMSPropLSTM converge fastest. For TIMIT, MomentumLSTM and SRLSTM converge fastest.

**Table 2** Test and validation MSEs at the end of the epoch with the lowest validation MSE for the TIMIT task. All of our proposed models outperform the baseline LSTM. Among models using  $N = 158$  hidden units, SRLSTM performs the best.

MODEL	N	# PARAMS	VAL. MSE	TEST MSE
LSTM	84	$\approx 83K$	$14.87 \pm 0.15$	$14.94 \pm 0.15$
LSTM	120	$\approx 135K$	$11.77 \pm 0.14$	$11.83 \pm 0.12$
LSTM	158	$\approx 200K$	$9.33 \pm 0.14$	$9.37 \pm 0.14$
MOMENTUMLSTM	84	$\approx 83K$	<b><math>10.90 \pm 0.19</math></b>	<b><math>10.98 \pm 0.18</math></b>
MOMENTUMLSTM	120	$\approx 135K$	<b><math>8.00 \pm 0.30</math></b>	<b><math>8.04 \pm 0.30</math></b>
MOMENTUMLSTM	158	$\approx 200K$	<b><math>5.86 \pm 0.14</math></b>	<b><math>5.87 \pm 0.15</math></b>
ADAMLSTM	158	$\approx 200K$	$8.66 \pm 0.15$	$8.69 \pm 0.14$
RMSPROPLSTM	158	$\approx 200K$	$9.13 \pm 0.33$	$9.17 \pm 0.33$
SRLSTM	158	$\approx 200K$	<b><math>5.81 \pm 0.10</math></b>	<b><math>5.83 \pm 0.10</math></b>


accuracy and faster convergence speed compared to LSTM. Figure 3 shows the convergence of MomentumLSTM vs. LSTM when using  $N = 158$  hidden units.

### 2.5.3 Copying and adding tasks

Two other important tasks for measuring the ability of a model to learn long-term dependency are the copying and adding tasks [36, 3]. In both copying and adding tasks, avoiding vanishing/exploding gradients becomes more relevant when the input sequence length increases. We compare the performance of MomentumLSTM over LSTM on these tasks. We also examine the performance of AdamLSTM, RMSPropLSTM, and SRLSTM on the same tasks. We summarize our results in Figure 4. In copying task for sequences of length 2K, MomentumLSTM obtains slightly better final training loss than the baseline LSTM (0.009 vs. 0.01). In adding task for sequence of length 750, both models achieve similar training loss of 0.162. However, AdamLSTM and RMSPropLSTM significantly outperform the baseline LSTM.

### 2.5.4 Word-level Penn TreeBank

To study the advantage of MomentumLSTM over LSTM on text data, we perform language modeling on a preprocessed version of the PTB dataset [58], which has been a standard benchmark for evaluating language models. Unlike the baselines used in




**Fig. 4** Train loss vs. iteration for (left) copying task with sequence length 2K and (right) adding task with sequence length 750. AdamLSTM and RMSPropLSTM converge faster and to better final losses than other models.

**Table 3** Model test perplexity at the end of the epoch with the lowest validation perplexity for the Penn Treebank language modeling task (word level).

MODEL	# PARAMS	VAL. PPL	TEST PPL
LSTM	$\approx 24M$	$61.96 \pm 0.83$	$59.71 \pm 0.99$
MOMENTUMLSTM	$\approx 24M$	<b><math>60.71 \pm 0.24</math></b>	<b><math>58.62 \pm 0.22</math></b>
SRLSTM	$\approx 24M$	$61.12 \pm 0.68$	$58.83 \pm 0.62$

the (P)MNIST and TIMIT experiments which contain one LSTM cell, in this PTB experiment, we use a three-layer LSTM model, which contains three concatenated LSTM cells, as the baseline. The size of this model in terms of the number of parameters is also much larger than those in the (P)MNIST and TIMIT experiments. Table 3 shows the test and validation perplexity (PPL) using the model parameters that yield the best validation loss. Again, MomentumLSTM achieves better perplexities and converges faster than the baseline LSTM (see Figure 5).



**Fig. 5** Train (left) and test loss (right) of MomentumLSTM (blue), SRLSTM (cyan), and LSTM (red) for the Penn Treebank language modeling tasks at word level.

### 2.5.5 NAG and Adam principled RNNs

Finally, we evaluate AdamLSTM, RMSPropLSTM and SRLSTM on all tasks. For (P)MNIST and TIMIT tasks, we summarize the test accuracy of the trained models in Tables 1 and 2 and provide the plots of train and test losses in Figure 3. We observe that though AdamLSTM and RMSPropLSTM work better than the MomentumLSTM at (P)MNIST task, they yield worse results at the TIMIT task. Interestingly, SRLSTM shows an opposite behavior - better than MomentumLSTM at TIMIT task but worse at (P)MNIST task. For the copying and adding tasks, Figure 4 shows that AdamLSTM and RMSPropLSTM converge faster and to better final training loss than other models

in both tasks. Finally, for the PTB task, both MomentumLSTM and SRLSTM outperform the baseline LSTM (see Figure 5 and Table 3). However, in this task, AdamLSTM and RMSPropLSTM yields slightly worse performance than the baseline LSTM. In particular, test PPL for AdamLSTM and RMSPropLSTM are  $61.11 \pm 0.31$ , and  $64.53 \pm 0.20$ , respectively, which are higher than the test PPL for LSTM ( $59.71 \pm 0.99$ ). We observe that there is no model that win in all tasks. This is somewhat expected, given the connection between our model and its analogy to optimization algorithm. An optimizer needs to be chosen for each particular task, and so is for our MomentumRNN. All of our models outperform the baseline LSTM.

### 3 Neural ODEs


In this section, we derive the continuous limit of heavy-ball momentum and then present a new class of neural ODEs, named heavy-ball neural ODEs (HBNODEs), which have two properties that imply practical advantages over NODEs: 1) The adjoint state of an HBNODE also satisfies an HBNODE, accelerating both forward and backward ODE solvers, thus significantly reducing the NFEs and improving the utility of trained models. 2) The spectrum of HBNODEs is well structured, enabling effective learning of long-term dependencies from complex sequential data. We verify the advantages of HBNODEs over NODEs on benchmark tasks, including image classification, learning complex dynamics, and sequential modeling. Our method requires remarkably fewer forward and backward NFEs, is more accurate, and learns long-term dependencies more effectively than the other ODE-based neural network models. Part of the results in this section has been accepted for publication at NeurIPS 2021 [110].

#### 3.1 Recap on neural ODEs

Neural ODEs (NODEs) are a family of continuous-depth machine learning models whose forward and backward propagations rely on solving an ODE and its adjoint equation [13]. NODEs model the dynamics of hidden features  $\mathbf{h}(t) \in \mathbb{R}^N$  using an ODE, which is parametrized by a neural network  $f(\mathbf{h}(t), t, \theta) \in \mathbb{R}^N$  with learnable parameters  $\theta$ , i.e.,

$$\frac{d\mathbf{h}(t)}{dt} = f(\mathbf{h}(t), t, \theta). \quad (16)$$

Starting from the input  $\mathbf{h}(t_0)$ , NODEs obtain the output  $\mathbf{h}(T)$  by solving (16) for  $t_0 \leq t \leq T$  with the initial value  $\mathbf{h}(t_0)$ , using a black-box numerical ODE solver. The NFEs that the black-box ODE solver requires in a single forward pass is an analogue for the continuous-depth models [13] to the depth of networks in ResNets [31]. The loss between  $\mathbf{h}(T)$  and the ground truth is



**Fig. 6** Contrasting NODE, ANODE, SONODE, HBNODE, and GHBNODE for CIFAR10 classification in NFEs, training time, and test accuracy. (Tolerance:  $10^{-5}$ , see subsection 3.4.2 for experimental details.)

denoted by  $\mathcal{L}(\mathbf{h}(T))$ ; we update parameters  $\theta$  using the following gradient [70]

$$\frac{d\mathcal{L}(\mathbf{h}(T))}{d\theta} = \int_{t_0}^T \mathbf{a}(t) \frac{\partial f(\mathbf{h}(t), t, \theta)}{\partial \theta} dt, \quad (17)$$

where  $\mathbf{a}(t) := \partial \mathcal{L} / \partial \mathbf{h}(t)$  satisfies the following adjoint equation

$$\frac{d\mathbf{a}(t)}{dt} = -\mathbf{a}(t) \frac{\partial f(\mathbf{h}(t), t, \theta)}{\partial \mathbf{h}}. \quad (18)$$

NODEs are flexible in learning from irregularly-sampled sequential data and particularly suitable for learning complex dynamical systems [13, 78, 113, 63, 23, 41], which can be trained by efficient algorithms [72, 19, 114]. The drawback of NODEs is also prominent. In many ML tasks, NODEs require very high NFEs in both training and inference, especially in high accuracy settings where a lower tolerance is needed. The NFEs increase rapidly with training; high NFEs reduce computational speed and accuracy of NODEs and can lead to blow-ups in the worst-case scenario [29, 24, 56, 63]. As an illustration, we train NODEs for CIFAR10 classification using the same model and experimental settings as in [24], except using a tolerance of  $10^{-5}$ ; Fig. 6 shows both forward and backward NFEs and the training time of different ODE-based models; we see that NFEs and computational times increase very rapidly for NODE, ANODE [24], and SONODE [63]. More results on the large NFE and degrading utility issues for different benchmark experiments are available in Section 3.4. Another issue is that NODEs often fail to effectively learn long-term dependencies in sequential data [47], discussed in subsection 3.3.

## 3.2 Heavy-ball neural ODEs

### 3.2.1 Heavy-ball ODE

We derive the HBODE from the heavy-ball momentum method. For any fixed step size  $s$ , let  $\mathbf{m}^k := (\mathbf{x}^{k+1} - \mathbf{x}^k) / \sqrt{s}$ , and let  $\beta := 1 - \gamma\sqrt{s}$ , where  $\gamma \geq 0$  is another hyperparameter. Then we can rewrite (1) as

$$\mathbf{m}^{k+1} = (1 - \gamma\sqrt{s})\mathbf{m}^k - \sqrt{s}\nabla F(\mathbf{x}^k); \quad \mathbf{x}^{k+1} = \mathbf{x}^k + \sqrt{s}\mathbf{m}^{k+1}. \quad (19)$$

Let  $s \rightarrow 0$  in (19); we obtain the following system of first-order ODEs,

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{m}(t); \quad \frac{d\mathbf{m}(t)}{dt} = -\gamma\mathbf{m}(t) - \nabla F(\mathbf{x}(t)). \quad (20)$$

This can be further rewritten as a second-order heavy-ball ODE (HBODE), which also models a damped oscillator,

$$\frac{d^2\mathbf{x}(t)}{dt^2} + \gamma\frac{d\mathbf{x}(t)}{dt} = -\nabla F(\mathbf{x}(t)). \quad (21)$$

### 3.2.2 Heavy-ball neural ODEs

Similar to NODE, we parameterize  $-\nabla F$  in (21) using a neural network  $f(\mathbf{h}(t), t, \theta)$ , resulting in the following HBNODE with initial position  $\mathbf{h}(t_0)$  and momentum  $\mathbf{m}(t_0) := d\mathbf{h}/dt(t_0)$ ,

$$\frac{d^2\mathbf{h}(t)}{dt^2} + \gamma\frac{d\mathbf{h}(t)}{dt} = f(\mathbf{h}(t), t, \theta), \quad (22)$$

where  $\gamma \geq 0$  is the damping parameter, which can be set as a tunable or a learnable hyperparameter with positivity constraint. In the trainable case, we use  $\gamma = \epsilon \cdot \text{sigmoid}(\omega)$  for a trainable  $\omega \in \mathbb{R}$  and a fixed tunable upper bound  $\epsilon$  (we set  $\epsilon = 1$  below). According to (20), HBNODE (22) is equivalent to

$$\frac{d\mathbf{h}(t)}{dt} = \mathbf{m}(t); \quad \frac{d\mathbf{m}(t)}{dt} = -\gamma\mathbf{m}(t) + f(\mathbf{h}(t), t, \theta). \quad (23)$$

Equation (22) (or equivalently, the system (23)) defines the forward ODE for the HBNODE, and we can use either the first-order (Prop. 2) or the second-order (Prop. 1) adjoint sensitivity method to update the parameter  $\theta$  [63].

**Proposition 1 (Adjoint equation for HBNODE)** *The adjoint state  $\mathbf{a}(t) := \partial\mathcal{L}/\partial\mathbf{h}(t)$  for the HBNODE (22) satisfies the following HBODE with the same damping parameter  $\gamma$  as that in (22),*

$$\frac{d^2\mathbf{a}(t)}{dt^2} - \gamma\frac{d\mathbf{a}(t)}{dt} = \mathbf{a}(t)\frac{\partial f}{\partial\mathbf{h}}(\mathbf{h}(t), t, \theta). \quad (24)$$

*Proof* Consider the following coupled first-order ODE system

$$\frac{\partial}{\partial t} \begin{bmatrix} \mathbf{h} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ f(\mathbf{h}(t), \mathbf{v}(t), t, \theta) \end{bmatrix}, \quad \begin{bmatrix} \mathbf{h} \\ \mathbf{v} \end{bmatrix} (t_0) = \begin{bmatrix} \mathbf{h}_{t_0} \\ \mathbf{v}_{t_0} \end{bmatrix}. \quad (25)$$

Denote  $\mathbf{z} = \begin{bmatrix} \mathbf{h} \\ \mathbf{v} \end{bmatrix}$  and final state as

$$\begin{bmatrix} \mathbf{h}(T) \\ \mathbf{v}(T) \end{bmatrix} = \begin{bmatrix} \mathbf{h}_T \\ \mathbf{v}_T \end{bmatrix} = \mathbf{z}_T. \quad (26)$$

Then the adjoint equation is given by

$$\frac{\partial \mathbf{A}(t)}{\partial t} = -\mathbf{A}(t) \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \frac{\partial f}{\partial \mathbf{h}} & \frac{\partial f}{\partial \mathbf{v}} \end{bmatrix}, \quad \mathbf{A}(T) = -\mathbf{I}, \quad \mathbf{a}(t) = -\frac{d\mathcal{L}}{dz_T} \mathbf{A}(t). \quad (27)$$

By rewriting  $\mathbf{A} = [\mathbf{A}_h \ \mathbf{A}_v]$ , we have the following differential equations

$$\frac{\partial \mathbf{A}_h(t)}{\partial t} = -\mathbf{A}_v(t) \frac{\partial f}{\partial \mathbf{h}}, \quad \frac{\partial \mathbf{A}_v(t)}{\partial t} = -\mathbf{A}_h(t) - \mathbf{A}_v(t) \frac{\partial f}{\partial \mathbf{v}}, \quad (28)$$

with initial conditions

$$\mathbf{A}_h(T) = -\begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{A}_v(T) = -\begin{bmatrix} \mathbf{0} \\ \mathbf{I} \end{bmatrix}, \quad (29)$$

and adjoint states

$$\mathbf{a}_h(t) = \frac{d\mathcal{L}}{dz_T} \mathbf{A}_h(t), \quad \mathbf{a}_v(t) = \frac{d\mathcal{L}}{dz_T} \mathbf{A}_v(t). \quad (30)$$

The gradient equations becomes

$$\frac{d\mathcal{L}}{d\theta} = \int_{t_0}^T \mathbf{a} \begin{bmatrix} \mathbf{0} \\ \frac{\partial f}{\partial \theta} \end{bmatrix} dt = \int_{t_0}^T \mathbf{a}_v \frac{\partial f}{\partial \theta} dt, \quad \frac{d\mathcal{L}}{d\mathbf{h}_{t_0}} = \mathbf{a}_h(t_0), \quad \frac{d\mathcal{L}}{d\mathbf{v}_{t_0}} = \mathbf{a}_v(t_0). \quad (31)$$

Note  $\mathbf{h}_{t_0}$  is fixed, and thus  $\mathbf{a}_h$  disappears in gradient computation. Therefore, we are only interested in  $\mathbf{a}_v$ . Thus the adjoint  $\mathbf{A}_v$  satisfies the following second-order ODE

$$\frac{\partial^2 \mathbf{A}_v(t)}{\partial t^2} = \mathbf{A}_v(t) \frac{\partial f}{\partial \mathbf{h}} - \frac{\partial(\mathbf{A}_v(t) \frac{\partial f}{\partial \mathbf{v}})}{\partial t}, \quad (32)$$

and thus

$$\frac{\partial^2 \mathbf{a}_v(t)}{\partial t^2} = \mathbf{a}_v(t) \frac{\partial f}{\partial \mathbf{h}} - \frac{\partial(\mathbf{a}_v(t) \frac{\partial f}{\partial \mathbf{v}})}{\partial t}, \quad (33)$$

with initial conditions

$$\mathbf{a}_v(T) = -\frac{d\mathcal{L}}{dz} \mathbf{A}_v(T) = \frac{d\mathcal{L}}{dv_T}, \quad \frac{\partial \mathbf{a}_v(T)}{\partial t} = -\frac{d\mathcal{L}}{d\mathbf{h}_T} - \mathbf{a}_v(T) \frac{\partial f}{\partial \mathbf{v}}(T). \quad (34)$$

As HBNODE takes the form

$$\frac{d^2 \mathbf{h}(t)}{dt^2} + \gamma \frac{d\mathbf{h}(t)}{dt} = f(\mathbf{h}(t), t, \theta), \quad (35)$$

which can also be viewed as a SONODE. By applying the adjoint equation (33), we arrive at

$$\frac{\partial^2 \mathbf{a}(t)}{\partial t^2} = \mathbf{a}(t) \frac{\partial f}{\partial \mathbf{h}} + \gamma \frac{\partial \mathbf{a}(t)}{\partial t}. \quad (36)$$

As HBNODE only carries its state  $\mathbf{h}$  to the loss  $\mathcal{L}$ , we have  $\frac{d\mathcal{L}}{dv_T} = 0$ , and thus the initial conditions in equation (34) becomes

$$\mathbf{a}(T) = \mathbf{0}, \quad \frac{\partial \mathbf{a}(T)}{\partial t} = -\frac{d\mathcal{L}}{d\mathbf{h}_T}, \quad (37)$$

which concludes the proof of Proposition 1.

*Remark 5* Note that we solve the adjoint equation (24) from time  $t = T$  to  $t = t_0$  in the backward propagation. By letting  $\tau = T - t$  and  $\mathbf{b}(\tau) = \mathbf{a}(T - \tau)$ , we can rewrite (24) as follows,

$$\frac{d^2 \mathbf{b}(\tau)}{d\tau^2} + \gamma \frac{d\mathbf{b}(\tau)}{d\tau} = \mathbf{b}(\tau) \frac{\partial f}{\partial \mathbf{h}}(\mathbf{h}(T - \tau), T - \tau, \theta). \quad (38)$$

Therefore, the adjoint of the HBNODE is also a HBNODE and they have the same damping parameter.

**Proposition 2 (Adjoint equations for the first-order HBNODE system)** *The adjoint states  $\mathbf{a}_h(t) := \partial \mathcal{L} / \partial \mathbf{h}(t)$  and  $\mathbf{a}_m(t) := \partial \mathcal{L} / \partial \mathbf{m}(t)$  for the first-order HBNODE system (23) satisfy*

$$\frac{d\mathbf{a}_h(t)}{dt} = -\mathbf{a}_m(t) \frac{\partial f}{\partial \mathbf{h}}(\mathbf{h}(t), t, \theta); \quad \frac{d\mathbf{a}_m(t)}{dt} = -\mathbf{a}_h(t) + \gamma \mathbf{a}_m(t). \quad (39)$$

*Proof* The coupled form of HBNODE is a coupled first-order ODE system of the form

$$\frac{\partial}{\partial t} \begin{bmatrix} \mathbf{h} \\ \mathbf{m} \end{bmatrix} = \begin{bmatrix} \mathbf{m} \\ -\gamma \mathbf{m} + f(\mathbf{h}(t), t, \theta) \end{bmatrix}, \quad \begin{bmatrix} \mathbf{h} \\ \mathbf{m} \end{bmatrix}(t_0) = \begin{bmatrix} \mathbf{h}_{t_0} \\ \mathbf{m}_{t_0} \end{bmatrix}. \quad (40)$$

Denote the final state as

$$\begin{bmatrix} \mathbf{h}(T) \\ \mathbf{m}(T) \end{bmatrix} = \begin{bmatrix} \mathbf{h}_T \\ \mathbf{m}_T \end{bmatrix} = z. \quad (41)$$

Using the conclusions from the proof of Proposition 1, we have the adjoint equation

$$\frac{\partial \mathbf{A}(t)}{\partial t} = -\mathbf{A}(t) \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \frac{\partial f}{\partial \mathbf{h}} & -\gamma \mathbf{I} \end{bmatrix}, \quad \mathbf{A}(T) = -\mathbf{I}, \quad \mathbf{a}(t) = -\frac{d\mathcal{L}}{dz} \mathbf{A}(t). \quad (42)$$

Let  $[\mathbf{a}_h \ \mathbf{a}_m] = \mathbf{a}$ , by linearity we have

$$\frac{\partial [\mathbf{a}_h \ \mathbf{a}_m]}{\partial t} = -[\mathbf{a}_h \ \mathbf{a}_m] \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \frac{\partial f}{\partial \mathbf{h}} & -\gamma \mathbf{I} \end{bmatrix}, \quad [\mathbf{a}_h(T) \ \mathbf{a}_m(T)] = \left[ \frac{d\mathcal{L}}{d\mathbf{h}_T} \ \frac{d\mathcal{L}}{d\mathbf{m}_T} \right], \quad (43)$$

which gives us the initial conditions at  $t = T$ , and the simplified first-order ODE system

$$\frac{\partial \mathbf{a}_h}{\partial t} = -\mathbf{a}_m \frac{\partial f}{\partial \mathbf{h}}, \quad \frac{\partial \mathbf{a}_m}{\partial t} = -\mathbf{a}_h + \gamma \mathbf{a}_m, \quad (44)$$

concluding the proof of Proposition 2.

*Remark 6* Let  $\tilde{\mathbf{a}}_m(t) = d\mathbf{a}_m(t)/dt$ , then  $\mathbf{a}_m(t)$  and  $\tilde{\mathbf{a}}_m(t)$  satisfies the following first-order heavy-ball ODE system

$$\frac{d\mathbf{a}_m(t)}{dt} = \tilde{\mathbf{a}}_m(t); \quad \frac{d\tilde{\mathbf{a}}_m(t)}{dt} = \mathbf{a}_m(t) \frac{\partial f}{\partial \mathbf{h}}(\mathbf{h}(t), t, \theta) + \gamma \tilde{\mathbf{a}}_m(t). \quad (45)$$

Note that we solve this system backward in time in back-propagation. Moreover, we have  $\mathbf{a}_h(t) = \gamma \mathbf{a}_m(t) - \tilde{\mathbf{a}}_m(t)$ .



Similar to [63], we use the coupled first-order HBNODE system (23) and its adjoint first-order HBNODE system (39) for practical implementation, since the entangled representation permits faster computation [63] of the gradients of the coupled ODE systems.


### 3.2.3 Generalized heavy-ball neural ODEs

In this part, we propose a generalized version of HBNODE (GHBNODE), see (46), to mitigate the potential blow-up issue in training ODE-based models. We observe that  $\mathbf{h}(t)$  of ANODEs [24], SONODEs [63], and HBNODEs (23) usually grows much faster than that of NODEs. The fast growth of  $\mathbf{h}(t)$  can lead to finite-time blow up. As an illustration, we compare the performance of NODE, ANODE, SONODE, HBNODE, and GHBNODE on the Silverbox task as in [63]. The goal of the task is to learn the voltage of an electronic circuit that resembles a Duffing oscillator, where the input voltage  $V_1(t)$  is used to predict the output  $V_2(t)$ . Similar to the setting in [63], we first augment ANODE by 1 dimension with 0-augmentation and augment SONODE, HBNODE, and GHBNODE with a dense network. We use a simple dense layer to parameterize  $f$  for all five models, with an extra input term for  $V_1(t)$ <sup>3</sup>. For both HBNODE and GHBNODE, we set the damping parameter  $\gamma$  to be  $\text{sigmoid}(-3)$ . For GHBNODE (46) below, we set  $\sigma(\cdot)$  to be the `hardtanh` function with bound  $[-5, 5]$  and  $\xi = \ln(2)$ . As shown in Fig. 7, compared to the vanilla NODE, the  $\ell_2$  norm of  $\mathbf{h}(t)$  grows much faster when a higher order NODE is used, which leads to blow-up during training. Similar issues arise in the time series experiments (see Section 3.4.4), where SONODE blows up during long term integration in time, and HBNODE suffers from the same issue with some initialization.

To alleviate the problem above, we propose the following GHBNODE

$$\begin{aligned} \frac{d\mathbf{h}(t)}{dt} &= \sigma(\mathbf{m}(t)), \\ \frac{d\mathbf{m}(t)}{dt} &= -\gamma\mathbf{m}(t) + f(\mathbf{h}(t), t, \theta) - \xi\mathbf{h}(t), \end{aligned} \quad (46)$$

where  $\sigma(\cdot)$  is a nonlinear activation, which is set as `tanh` in our experiments. The positive hyperparameters  $\gamma, \xi > 0$  are tunable or learnable. In the trainable case, we let  $\gamma = \epsilon \cdot \text{sigmoid}(\omega)$  as in HBNODE, and  $\xi = \text{softplus}(\chi)$  to ensure that  $\gamma, \xi \geq 0$ . Here, we integrate two main ideas into the design of GHBNODE: (i) We incorporate the gating mechanism used in LSTM [36] and



**Fig. 7** Contrasting  $\mathbf{h}(t)$  for different models.  $\mathbf{h}(t)$  in ANODE, SONODE, and HBNODE grows much faster than that in NODE. GHBNODE controls the growth of  $\mathbf{h}(t)$  effectively when  $t$  is large.

<sup>3</sup> Here, we exclude an  $\mathbf{h}^3$  term that appeared in the original Duffing oscillator model because including it would result in finite-time explosion.

GRU [15], which can suppress the aggregation of  $\mathbf{m}(t)$ ; (ii) Following the idea of skip connection [32], we add the term  $\xi\mathbf{h}(t)$  into the governing equation of  $\mathbf{m}(t)$ , which benefits training and generalization of GHBNODEs. Fig. 7 shows that GHBNODE can indeed control the growth of  $\mathbf{h}(t)$  effectively.

**Proposition 3 (Adjoint equations for GHBNODEs)** *The adjoint states  $\mathbf{a}_h(t) := \partial\mathcal{L}/\partial\mathbf{h}(t)$ ,  $\mathbf{a}_m(t) := \partial\mathcal{L}/\partial\mathbf{m}(t)$  for the GHBNODE (46) satisfy the following first-order ODE system*

$$\begin{aligned}\frac{\partial\mathbf{a}_h(t)}{\partial t} &= -\mathbf{a}_m(t)\left(\frac{\partial f}{\partial\mathbf{h}}(\mathbf{h}(t), t, \theta) - \xi\mathbf{I}\right) \\ \frac{\partial\mathbf{a}_m(t)}{\partial t} &= -\mathbf{a}_h(t)\sigma'(\mathbf{m}(t)) + \gamma\mathbf{a}_m(t).\end{aligned}\quad (47)$$

*Proof* GHBNODE can be written as the following first-order ODE system

$$\frac{\partial}{\partial t} \begin{bmatrix} \mathbf{h} \\ \mathbf{m} \end{bmatrix} = \begin{bmatrix} \sigma(\mathbf{m}) \\ -\gamma\mathbf{m} + f(\mathbf{h}(t), t, \theta) - \xi\mathbf{h}(t) \end{bmatrix}, \quad \begin{bmatrix} \mathbf{h} \\ \mathbf{m} \end{bmatrix}(t_0) = \begin{bmatrix} \mathbf{h}_{t_0} \\ \mathbf{m}_{t_0} \end{bmatrix}. \quad (48)$$

Denote the final state as  $\mathbf{z}_T := [\mathbf{h}_T \mathbf{m}_T]$ . We have the adjoint equation

$$\frac{\partial\mathbf{A}(t)}{\partial t} = -\mathbf{A}(t) \begin{bmatrix} \mathbf{0} & \sigma'(\mathbf{m}) \\ \frac{\partial f}{\partial\mathbf{h}} - \xi\mathbf{I} & -\gamma\mathbf{I} \end{bmatrix}, \quad \mathbf{A}(T) = -\mathbf{I}, \quad \mathbf{a}(t) = -\frac{d\mathcal{L}}{d\mathbf{z}_T}\mathbf{A}(t). \quad (49)$$

Let  $[\mathbf{a}_h \ \mathbf{a}_m] = \mathbf{a}$ , by linearity we have

$$\begin{aligned}\frac{\partial [\mathbf{a}_h \ \mathbf{a}_m]}{\partial t} &= -[\mathbf{a}_h \ \mathbf{a}_m] \begin{bmatrix} \mathbf{0} & \sigma'(\mathbf{m}) \\ \frac{\partial f}{\partial\mathbf{h}} - \xi\mathbf{I} & -\gamma\mathbf{I} \end{bmatrix}, \\ [\mathbf{a}_h(T) \ \mathbf{a}_m(T)] &= \left[ \frac{d\mathcal{L}}{d\mathbf{h}_T} \ \frac{d\mathcal{L}}{d\mathbf{m}_T} \right],\end{aligned}\quad (50)$$

which gives us the initial conditions at  $t = T$ , and the simplified first-order ODE system

$$\frac{\partial\mathbf{a}_h}{\partial t} = -\mathbf{a}_m\left(\frac{\partial f}{\partial\mathbf{h}} - \xi\mathbf{I}\right), \quad \frac{\partial\mathbf{a}_m}{\partial t} = -\mathbf{a}_h\sigma'(\mathbf{m}) + \gamma\mathbf{a}_m, \quad (51)$$

concluding the proof of Proposition 3.

Though the adjoint state of the GHBNODE (47) does not satisfy the exact heavy-ball ODE, based on our empirical study, it also significantly reduces the backward NFEs.

### 3.3 Learning long-term dependencies – Vanishing gradient

As mentioned in Section 2, the vanishing gradient is the main bottleneck for training RNNs with long-term dependencies. As the continuous analogue of RNN, NODEs as well as their hybrid ODE-RNN models, may also suffer from vanishing in the adjoint state  $\mathbf{a}(t) := \partial\mathcal{L}/\partial\mathbf{h}(t)$  [47]. When the vanishing gradient issue happens,  $\mathbf{a}(t)$  goes to  $\mathbf{0}$  quickly as  $T - t$  increases, then  $d\mathcal{L}/d\theta$  in (17) will be independent of these  $\mathbf{a}(t)$ . We have the following expressions for the adjoint states of the NODE and HBNODE:

– For NODE, we have

$$\frac{\partial \mathcal{L}}{\partial \mathbf{h}_t} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \frac{\partial \mathbf{h}_T}{\partial \mathbf{h}_t} = \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \exp \left\{ - \int_T^t \frac{\partial f}{\partial \mathbf{h}}(\mathbf{h}(s), s, \theta) ds \right\}. \quad (52)$$

– For GHBNODE<sup>4</sup>, from (39) we can derive

$$\left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}} \frac{\partial \mathcal{L}}{\partial \mathbf{m}_t} \right] = \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \underbrace{\left[ \begin{array}{c} \frac{\partial \mathbf{h}_T}{\partial \mathbf{h}_t} \frac{\partial \mathbf{h}_T}{\partial \mathbf{m}_t} \\ \frac{\partial \mathbf{m}_T}{\partial \mathbf{h}_t} \frac{\partial \mathbf{m}_T}{\partial \mathbf{m}_t} \end{array} \right]}_{:=M} = \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \exp \left\{ - \int_T^t \underbrace{\left[ \begin{array}{c} \mathbf{0} \\ (\frac{\partial f}{\partial \mathbf{h}} - \xi \mathbf{I}) - \gamma \mathbf{I} \end{array} \right]}_{:=M} ds \right\}. \quad (53)$$

Note that the matrix exponential is directly related to its eigenvalues. By Schur decomposition, there exists an orthogonal matrix  $\mathbf{Q}$  and an upper triangular matrix  $\mathbf{U}$ , where the diagonal entries of  $\mathbf{U}$  are eigenvalues of  $\mathbf{Q}$  ordered by their real parts, such that

$$-M = \mathbf{Q} \mathbf{U} \mathbf{Q}^\top \implies \exp\{-M\} = \mathbf{Q} \exp\{\mathbf{U}\} \mathbf{Q}^\top. \quad (54)$$

Let  $\mathbf{v}^\top := \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \mathbf{Q}$ , then (53) can be rewritten as

$$\begin{aligned} \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_t} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_t} \right] &= \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \exp\{-M\} \\ &= \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \mathbf{Q} \exp\{\mathbf{U}\} \mathbf{Q}^\top = \mathbf{v}^\top \exp\{\mathbf{U}\} \mathbf{Q}^\top. \end{aligned} \quad (55)$$

Taking the  $\ell_2$  norm in (55) and dividing both sides by  $\left\| \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \right\|_2$ , we have

$$\frac{\left\| \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_t} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_t} \right] \right\|_2}{\left\| \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \right\|_2} = \frac{\left\| \mathbf{v}^\top \exp\{\mathbf{U}\} \mathbf{Q}^\top \right\|_2}{\left\| \mathbf{v}^\top \mathbf{Q}^\top \right\|_2} = \frac{\left\| \mathbf{v}^\top \exp\{\mathbf{U}\} \right\|_2}{\left\| \mathbf{v} \right\|_2} = \left\| \mathbf{e}^\top \exp\{\mathbf{U}\} \right\|_2, \quad (56)$$

i.e.,  $\left\| \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_t} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_t} \right] \right\|_2 = \left\| \mathbf{e}^\top \exp\{\mathbf{U}\} \right\|_2 \left\| \left[ \frac{\partial \mathcal{L}}{\partial \mathbf{h}_T} \quad \frac{\partial \mathcal{L}}{\partial \mathbf{m}_T} \right] \right\|_2$  where  $\mathbf{e} = \mathbf{v} / \|\mathbf{v}\|_2$ .

**Proposition 4** *The eigenvalues of  $-M$  can be paired so that the sum of each pair equals  $(t - T)\gamma$ .*

*Proof* Let  $\mathbf{F} = \frac{1}{t-T} \int_T^t \frac{\partial f}{\partial \mathbf{h}}(\mathbf{h}(s), s, \theta) ds - \xi \mathbf{I}$ ,  $\mathbf{J} = \frac{1}{t-T} \int_T^t \frac{\partial \sigma}{\partial \mathbf{m}}(\mathbf{m}(s)) ds$ , and  $\mathbf{H} = \frac{1}{t-T} \mathbf{M}$ , then we have the following equation

$$\mathbf{H} = \frac{1}{t-T} \mathbf{M} = \begin{bmatrix} \mathbf{0} & \mathbf{J} \\ \mathbf{F} & -\gamma \mathbf{I} \end{bmatrix}. \quad (57)$$

As  $(\lambda + \gamma)\mathbf{I}$  commutes with any matrix  $\mathbf{F}$ , the characteristics polynomials of  $\mathbf{H}$  and  $\mathbf{J}\mathbf{F}$  satisfy the relation

$$\begin{aligned} ch_{\mathbf{H}}(\lambda) &= \det(\lambda \mathbf{I} - \mathbf{H}) = \det \begin{bmatrix} \lambda \mathbf{I} & -\mathbf{J} \\ -\mathbf{F} & (\lambda + \gamma) \mathbf{I} \end{bmatrix} \\ &= \det(\lambda(\lambda + \gamma) \mathbf{I} - \mathbf{J}\mathbf{F}) = -ch_{\mathbf{J}\mathbf{F}}(\lambda(\lambda + \gamma)). \end{aligned} \quad (58)$$

<sup>4</sup> HBNODE can be seen as a special GHBNODE with  $\xi = 0$  and  $\sigma$  be the identity map.

Since the characteristics polynomial of  $\mathbf{JF}$  splits in the field  $\mathbb{C}$  of complex numbers, i.e.  $ch_{\mathbf{JF}}(x) = \prod_{i=1}^n (x - \lambda_{\mathbf{JF},i})$ , we have

$$ch_{\mathbf{H}}(\lambda) = -ch_{\mathbf{JF}}(\lambda(\lambda + \gamma)) = -\prod_{i=1}^n (\lambda(\lambda + \gamma) - \lambda_{\mathbf{JF},i}). \quad (59)$$

Therefore, the eigenvalues of  $\mathbf{H}$  appear in  $n$  pairs with each pair satisfying the quadratic equation

$$\lambda(\lambda + \gamma) - \lambda_{\mathbf{JF},i} = 0. \quad (60)$$

By Vieta's formulas, the sum of these pairs are all  $-\gamma$ . Therefore, the eigenvalues of  $\mathbf{M}$  comes in  $n$  pairs and the sum of each pair is  $-(t - T)\gamma$ , which finishes the proof of Proposition 4.

For a given constant  $a > 0$ , we can group the upper triangular matrix  $\exp\{\mathbf{U}\}$  as follows


$$\exp\{\mathbf{U}\} := \begin{bmatrix} \exp\{\mathbf{U}_L\} & \mathbf{P} \\ \mathbf{0} & \exp\{\mathbf{U}_V\} \end{bmatrix}, \quad (61)$$

where the diagonal of  $\mathbf{U}_L$  ( $\mathbf{U}_V$ ) contains eigenvalues of  $-\mathbf{M}$  that are no less (greater) than  $(t - T)a$ . Then, we have  $\|\mathbf{e}^\top \exp\{\mathbf{U}\}\|_2 \geq \|\mathbf{e}_L^\top \exp\{\mathbf{U}_L\}\|_2$  where the vector  $\mathbf{e}_L$  denotes the first  $m$  columns of  $\mathbf{e}$  with  $m$  be the number of columns of  $\mathbf{U}_L$ . By choosing  $0 \leq \gamma \leq 2a$ , for every pair of eigenvalues of  $-\mathbf{M}$  there is at least one eigenvalue whose real part is no less than  $(t - T)a$ . Therefore,  $\exp\{\mathbf{U}_L\}$  decays at a rate at most  $(t - T)a$ , and the dimension of  $\mathbf{U}_L$  is at least  $N \times N$ . We avoid exploding gradients by clipping the  $\ell_2$  norm of the adjoint states similar to that used for training RNNs.

In contrast, all eigenvalues of the matrix  $\int_T^t \partial f / \partial \mathbf{h} ds$  in (52) for NODE can be very positive or negative, resulting in exploding or vanishing gradients. As an illustration, we consider the benchmark Walker2D kinematic simulation task that requires learning long-term dependencies effectively [47, 7]. We train ODE-RNN [78] and (G)HBNODE-RNN on this benchmark dataset, and the detailed experimental settings are provided in Section 3.4.4. Figure 8 plots  $\|\partial \mathcal{L} / \partial \mathbf{h}_t\|_2$  for ODE-RNN and  $\|[\partial \mathcal{L} / \partial \mathbf{h}_t \quad \partial \mathcal{L} / \partial \mathbf{m}_t]\|_2$  for (G)HBNODE-RNN, showing that the adjoint state of ODE-RNN vanishes quickly, while that of (G)HBNODE-RNN does not vanish even when the gap between  $T$  and  $t$  is very large.

### 3.4 Experimental Results


In this section, we compare the performance of the proposed HBNODE and GHBNODE with existing ODE-based models, including NODE [13], ANODE [24], and SONODE [63] on the benchmark point cloud separation, image classification, learning dynamical systems, and kinematic simulation. For all the experiments, we use Adam [43] as the benchmark optimization solver (the learning rate and batch size for each experiment are listed in Table 4). For HBNODE and GHBNODE, we set  $\gamma = \text{sigmoid}(\theta)$ , where  $\theta$  is a trainable weight initialized as  $\theta = -3$ .



**Fig. 8** Plot of the the  $\ell_2$ -norm of the adjoint states for ODE-RNN and (G)HBNODE-RNN back-propagated from the last time stamp. The adjoint state of ODE-RNN vanishes quickly when the gap between the final time  $T$  and intermediate time  $t$  becomes larger, while the adjoint states of (G)HBNODE-RNN decays much more slowly. This implies that (G)HBNODE-RNN is more effective in learning long-term dependency than ODE-RNN.

**Table 4** The batch size and learning rate for different datasets.


Dataset	Point Cloud	MNIST	CIFAR10	Plane Vibration	Walker2D
Batch Size	50	64	64	64	256
Learning Rate	0.01	0.001	0.001	0.0001	0.003



**Fig. 9** Comparison between NODE, ANODE, SONODE, HBNODE, and GHBNODE for two-dimensional point cloud separation. HBNODE and GHBNODE converge better and require less NFEs in both forward and backward propagation than the other benchmark models.

### 3.4.1 Point cloud separation

In this subsection, we consider the two-dimensional point cloud separation benchmark. A total of 120 points are sampled, in which 40 points are drawn uniformly from the circle  $\|\mathbf{r}\| < 0.5$ , and 80 points are drawn uniformly from the annulus  $0.85 < \|\mathbf{r}\| < 1.0$ . This experiment aims to learn effective features to classify these two point clouds. Following [24], we use a three-layer neural network to parameterize the right-hand side of each ODE-based model, integrate the ODE-based model from  $t_0 = 0$  to  $T = 1$ , and pass the integration results to a dense layer to generate the classification results. We set the size of hidden layers so that the models have similar sizes, and the number of pa-



**Fig. 10** Contrasting NODE [13], ANODE [24], SONODE [63], HBNODE, and GHBNODE for MNIST classification in NFE, training time, and test accuracy. (Tolerance:  $10^{-5}$ ).

rameters of NODE, ANODE, SONODE, HBNODE, and GHBNODE are 525, 567, 528, 568, and 568, respectively. To avoid the effects of numerical error of the black-box ODE solver we set tolerance of ODE solver to be  $10^{-7}$ . Figure 9 plots a randomly selected evolution of the point cloud separation for each model; we also compare the forward and backward NFEs and the training loss of these models (100 independent runs). HBNODE and GHBNODE improve training as the training loss consistently goes to zero over different runs, while ANODE and SONODE often get stuck at local minima, and NODE cannot separate the point cloud since it preserves the topology [24].

### 3.4.2 Image classification

We compare the performance of HBNODE and GHBNODE with the existing ODE-based models on MNIST and CIFAR10 classification tasks using the same setting as in [24]. We parameterize  $f(\mathbf{h}(t), t, \theta)$  using a 3-layer convolutional network for each ODE-based model, and the total number of parameters for each model is listed in Table 5. For a given input image of the size  $c \times h \times w$ , we first augment the number of channel from  $c$  to  $c + p$  with the augmentation dimension  $p$  dependent on each method<sup>5</sup>. Moreover, for SONODE, HBNODE and GHBNODE, we further include velocity or momentum with the same shape as the augmented state.

**Table 5** The number of parameters for each models for image classification.


Model	NODE	ANODE	SONODE	HBNODE	GHBNODE
#Params (MNIST)	85,315	85,462	86,179	85,931	85,235
#Params (CIFAR10)	173,611	172,452	171,635	172,916	172,916

*NFEs*. As shown in Figs. 6 and 10, the NFEs grow rapidly with training of the NODE, resulting in an increasingly complex model with reduced performance and the possibility of blow up. Input augmentation has been verified to effectively reduce the NFEs, as both ANODE and SONODE require fewer forward NFEs than NODE for the MNIST and CIFAR10 classification. However, input augmentation is less effective in controlling their backward NFEs. HBNODE

<sup>5</sup> We set  $p = 0, 5, 4, 4, 5/0, 10, 9, 9, 9$  on MNIST/CIFAR10 for NODE, ANODE, SONODE, HBNODE, and GHBNODE, respectively.

and GHBNODE require much fewer NFEs than the existing benchmarks, especially for backward NFEs. In practice, reducing NFEs implies reducing both training and inference time, as shown in Figs. 6 and 10.

*Accuracy.* We also compare the accuracy of different ODE-based models for MNIST and CIFAR10 classification. As shown in Figs. 6 and 10, HBNODE and GHBNODE have slightly better classification accuracy than the other three models; this resonates with the fact that less NFEs lead to simpler models which generalize better [24, 63].



**Fig. 11** NFE vs. tolerance (shown in the colorbar) for training ODE-based models for CIFAR10 classification. Both forward and backward NFEs of HBNODE and GHBNODE grow much more slowly than that of NODE, ANODE, and SONODE; especially the backward NFEs. As the tolerance decreases, the advantage of HBNODE and GHBNODE in reducing NFEs becomes more significant.


*NFEs vs. tolerance.* We further study the NFEs for different ODE-based models under different tolerances of the ODE solver using the same approach as in [13]. Figure 11 depicts the forward and backward NFEs for different models under different tolerances. We see that (i) both forward and backward NFEs grow quickly when tolerance is decreased, and HBNODE and GHBNODE require much fewer NFEs than other models; (ii) under different tolerances, the backward NFEs of NODE, ANODE, and SONODE are much larger than the forward NFEs, and the difference becomes larger when the tolerance decreases. In contrast, the forward and backward NFEs of HBNODE and GHBNODE scale almost linearly with each other. This reflects that the advantage in NFEs of (G)HBNODE over the benchmarks become more significant when a smaller tolerance is used.

### 3.4.3 Learning dynamical systems from irregularly-sampled time series

In this subsection, we learn dynamical systems from experimental measurements. In particular, we use the ODE-RNN framework [13, 78], with the recognition model being set to different ODE-based models, to study the vibration of an airplane dataset [62]. The dataset was acquired, from time 0 to 73627, by attaching a shaker underneath the right wing to provide input signals, and 5 attributes are recorded per time stamp; these attributes include voltage of input signal, force applied to aircraft, and acceleration at 3 different spots of the airplane. We randomly take out 10% of the data to make the time series irregularly-sampled. We use the first 50% of data as our train set, the next 25% as validation set, and the rest as test set. We divide each set into

non-overlapping segments of consecutive 65 time stamps of the irregularly-sampled time series, with each input instance consisting of 64 time stamps of the irregularly-sampled time series, and we aim to forecast 8 consecutive time stamps starting from the last time stamp of the segment. The input is fed through the the hybrid methods in a recurrent fashion; by changing the time duration of the last step of the ODE integration, we can forecast the output in the different time stamps. The output of the hybrid method is passed to a single dense layer to generate the output time series. In our experiments, we compare different ODE-based models hybrid with RNNs. The ODE of each model is parametrized by a 3-layer network whereas the RNN is parametrized by a simple dense network; the total number of parameters for ODE-RNN, ANODE-RNN, SONODE-RNN, HBNODE-RNN, and GHBNODE-RNN with 16, 22, 14, 15, 15 augmented dimensions are 15,986, 16,730, 16,649, 16,127, and 16,127, respectively. To avoid potential error due to the ODE solver, we use a tolerance of  $10^{-7}$ .

In training those hybrid models, we regularize the models by penalizing the L2 distance between the RNN output and the values of the next time stamp. Due to the second-order natural of the underlying dynamics [63], ODE-RNN and ANODE-RNN learn the dynamics very poorly with much larger training and test losses than the other models even they take smaller NFEs. HBNODE-RNN and GHBNODE-RNN give better prediction than SONODE-RNN using less backward NFEs.




**Fig. 12** Contrasting ODE-RNN, ANODE-RNN, SONODE-RNN, HBNODE-RNN, and GHBNODE-RNN for learning a vibrational dynamical system. Left most: The learned curves of each model vs. the ground truth (Time: <66 for training, 66-75 for testing).

#### 3.4.4 Walker2D kinematic simulation

We evaluate the performance of HBNODE-RNN and GHBNODE-RNN on the Walker2D kinematic simulation task, which requires learning long-term dependency effectively [47]. The dataset [7] consists of a dynamical system from kinematic simulation of a person walking from a pre-trained policy, aiming to learn the kinematic simulation of the MuJoCo physics engine [95]. The dataset is irregularly-sampled with 10% of the data removed from the simulation. Each input consists of 64 time stamps fed through the the hybrid methods in a recurrent fashion, and the output is passed to a single dense layer to generate the output time series. The goal is to provide an auto-regressive forecast so that the output time series is as close as the input sequence shifted one time stamp



to the right. We compare ODE-RNN (with 7 augmentation), ANODE-RNN (with 7 ANODE style augmentation), HBNODE-RNN (with 7 augmentation), and GHBNODE-RNN (with 7 augmentation). The RNN is parametrized by a 3-layer network whereas the ODE is parametrized by a simple dense network. The number of parameters of the above four models are 8,729, 8,815, 8,899, and 8,899, respectively. In Fig. 13, we compare the performance of the above four models on the Walker2D benchmark; HBNODE-RNN and GHBNODE-RNN not only require significantly less NFEs in both training (forward and backward) and in testing than ODE-RNN and ANODE-RNN, but also have much smaller training and test losses.



**Fig. 13** Contrasting ODE-RNN, ANODE-RNN, SONODE-RNN, HBNODE-RNN, and GHBNODE-RNN for the Walker-2D kinematic simulation.

## 4 Transformers

We further show that momentum can be integrated into transformers, which can significantly reduce the computational and memory costs of the standard transformer [98] and enhance the performance of linear transformers [39].

The self-attention mechanism is a fundamental building block of transformers [98, 42]. Given an input sequence  $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]^\top \in \mathbb{R}^{N \times D_x}$  of  $N$  feature vectors, the self-attention transformers it into another sequence  $\hat{\mathbf{V}} = [\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \dots, \hat{\mathbf{v}}_N]^\top \in \mathbb{R}^{N \times D_v}$  as follows

$$\hat{\mathbf{v}}_i = \sum_{j=1}^N \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right) \mathbf{v}_j, \text{ for } i = 1, \dots, N, \quad (62)$$

where the scalar  $\text{softmax}((\mathbf{q}_i^\top \mathbf{k}_j)/\sqrt{D})$  can be understood as the attention  $\hat{\mathbf{v}}_i$  pays to the input feature  $\mathbf{x}_j$ . The vectors  $\mathbf{q}_i$ ,  $\mathbf{k}_j$ , and  $\mathbf{v}_j$  are called the query, key, and value vectors, respectively; these vectors are computed as follows

$$\begin{aligned} [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_N]^\top &:= \mathbf{Q} = \mathbf{X} \mathbf{W}_Q^\top \in \mathbb{R}^{N \times D}, \\ [\mathbf{k}_1, \mathbf{k}_2, \dots, \mathbf{k}_N]^\top &:= \mathbf{K} = \mathbf{X} \mathbf{W}_K^\top \in \mathbb{R}^{N \times D}, \\ [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N]^\top &:= \mathbf{V} = \mathbf{X} \mathbf{W}_V^\top \in \mathbb{R}^{N \times D_v}, \end{aligned} \quad (63)$$

where  $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{D \times D_x}$ , and  $\mathbf{W}_V \in \mathbb{R}^{D_v \times D_x}$  are the weight matrices. We can further write (62) into the following compact form

$$\hat{\mathbf{V}} = \text{softmax}\left(\frac{\mathbf{Q} \mathbf{K}^\top}{\sqrt{D}}\right) \mathbf{V}, \quad (64)$$

where the softmax function is applied to each row of  $(\mathbf{Q}\mathbf{K}^\top)/\sqrt{D}$ . Equation (64) is also called the “scaled dot-product attention” or “softmax attention”. Each transformer layer  $T_\ell(\cdot)$  is defined via the following residual connection,

$$T_\ell(\mathbf{X}) = f_\ell(\hat{\mathbf{V}} + \mathbf{X}), \quad (65)$$

where  $f_\ell$  is a function that transforms each feature vector independently and usually chosen to be a feedforward network. We call a transformer built with softmax attention standard transformer or transformer. It is easy to see that both memory and computational complexity of (64) are  $\mathcal{O}(N^2)$  with  $N$  being the length of the input sequence. We can further introduce causal masking into (64) for autoregressive applications [98].

Transformers have become the state-of-the-art model for solving many challenging problems in natural language processing [98, 2, 18, 108, 21, 8, 37, 74] and computer vision [20, 86, 22, 96]. Nevertheless, the quadratic memory and computational cost of computing the softmax attention (64) is a major bottleneck for applying transformers to large-scale applications that involve very long sequences, such as those in [54, 38, 66]. Thus, much recent research on transformers has been focusing on developing efficient transformers, aiming to reduce the memory and computational complexities of the model [71, 14, 35, 66, 4, 1, 107, 90, 91, 44, 77, 100, 112, 107, 39, 16, 84, 82, 6, 76, 87, 68, 111]. A thorough survey of recent advances in efficient transformers is available at [93]. These efficient transformers have better memory and/or computational efficiency at the cost of a significant reduction in accuracy.

#### 4.1 Motivation

In [39], the authors have established a connection between transformers and RNNs through the kernel trick. They proposed the linear transformer, which can be considered a rank-one approximation of the softmax transformer. Linear transformers have computational advantages in training, test, and inference: the RNN formulation (see (69) below) enjoys fast inference, especially for autoregressive tasks, and the unrolled RNN formulation (see (67) below) is efficient for fast training. See subsection 4.2 for a detailed review of the linear transformer and its advantages. [60] proposes integrating momentum into RNNs to accelerate training RNNs and improve learning long-term dependencies. We notice that MomentumRNN also enjoys a closed unrolling form, which is quite unique among existing techniques for improving RNNs, enabling fast training, test, and inference; see section 4.3 for details. As such, in this section we study *how momentum improves linear transformers?*

#### 4.2 Linear transformer

Transformers learn long-term dependencies in sequences effectively and concurrently through the self-attention mechanism. Note we can write (62) as

$\hat{\mathbf{v}}_i = (\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j) \mathbf{v}_j) / (\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j))$ , where  $k(\mathbf{q}_i, \mathbf{k}_j) := \exp(\mathbf{q}_i^\top \mathbf{k}_j / \sqrt{D})$ . In linear transformers [107, 39, 16, 84], the feature map  $k(\mathbf{q}_i, \mathbf{k}_j)$  is linearized as the product of feature maps  $\phi(\cdot)$  on the vectors  $\mathbf{q}_i$  and  $\mathbf{k}_j$ , i.e.,  $k(\mathbf{q}_i, \mathbf{k}_j) = \phi(\mathbf{q}_i)^\top \phi(\mathbf{k}_j)$ . The associative property of matrix multiplication is then utilized to derive the following efficient computation of the attention map

$$\hat{\mathbf{v}}_i = \frac{\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j) \mathbf{v}_j}{\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j)} = \frac{\sum_{j=1}^N \phi(\mathbf{q}_i)^\top \phi(\mathbf{k}_j) \mathbf{v}_j}{\sum_{j=1}^N \phi(\mathbf{q}_i)^\top \phi(\mathbf{k}_j)} = \frac{\phi(\mathbf{q}_i)^\top \sum_{j=1}^N \phi(\mathbf{k}_j) \mathbf{v}_j}{\phi(\mathbf{q}_i)^\top \sum_{j=1}^N \phi(\mathbf{k}_j)}. \quad (66)$$

In the matrix-product form, we can further write (66) as follows

$$\hat{\mathbf{V}} = \frac{\phi(\mathbf{Q})(\phi(\mathbf{K})^\top \mathbf{V})}{\phi(\mathbf{Q})\phi(\mathbf{K})^\top}. \quad (67)$$

Replacing  $(\phi(\mathbf{Q})\phi(\mathbf{K}^\top))\mathbf{V}$  with  $\phi(\mathbf{Q})(\phi(\mathbf{K}^\top)\mathbf{V})$  reduces the memory and computational cost of computing the attention map from  $\mathcal{O}(N^2)$  to  $\mathcal{O}(N)$ , making linear transformers scalable to very long sequences.

Causal masking can be easily implemented in the linearized attention by truncating the summation term in the last equation of (66), resulting in

$$\hat{\mathbf{v}}_i = \frac{\phi(\mathbf{q}_i)^\top \sum_{j=1}^i \phi(\mathbf{k}_j) \mathbf{v}_j}{\phi(\mathbf{q}_i)^\top \sum_{j=1}^i \phi(\mathbf{k}_j)} := \frac{\phi(\mathbf{q}_i)^\top \mathbf{s}_i}{\phi(\mathbf{q}_i)^\top \mathbf{z}_i}, \quad (68)$$

where  $\mathbf{s}_i = \sum_{j=1}^i \phi(\mathbf{k}_j) \mathbf{v}_j$  and  $\mathbf{z}_i = \sum_{j=1}^i \phi(\mathbf{k}_j)$ . The states  $\mathbf{s}_i$  and  $\mathbf{z}_i$  can be computed recurrently.

*Efficient inference via the RNN formulation.* Self-attention processes tokens of a sequence concurrently, enabling fast training of transformers. However, during inference, the output for timestep  $i$  is the input for timestep  $i + 1$ . As a result, the inference in standard transformers cannot be parallelized and is thus inefficient. Linear transformers provide an elegant approach to fixing this issue by leveraging their RNN formulation. In particular, we can further write the linear attention with causal masking in (68) into the following RNN form<sup>6</sup>

$$\begin{aligned} \mathbf{s}_i &= \mathbf{s}_{i-1} + \phi(\mathbf{k}_i) \mathbf{v}_i^\top; \\ \mathbf{z}_i &= \mathbf{z}_{i-1} + \phi(\mathbf{k}_i); \\ \hat{\mathbf{v}}_i &= \frac{\phi(\mathbf{q}_i)^\top \mathbf{s}_i}{\phi(\mathbf{q}_i)^\top \mathbf{z}_i}, \end{aligned} \quad (69)$$

where  $\mathbf{s}_0 = \mathbf{0}$  and  $\mathbf{z}_0 = \mathbf{0}$ . Note that this RNN formulation of linear transformers with causal masking contains two memory states  $\mathbf{s}_i$  and  $\mathbf{z}_i$ .

<sup>6</sup> We omit the nonlinearity (a two-layer feedforward network) compared to [39].

### 4.3 Momentum transformer

In this section, we present the *momentum transformer*. We start by integrating the heavy-ball momentum into the RNN formulation of causal linear attention in (69), resulting in the causal momentum attention. Next, we generalize the causal momentum attention to momentum attention that can efficiently train the model. Moreover, we propose the *momentum connection* to replace residual connections between the attention  $\hat{\mathbf{V}}$  and the input  $\mathbf{X}$  in (65) to boost the model's performance. Finally, we derive the adaptive momentum attention from the theory of optimal choice of momentum for the heavy-ball method.

#### 4.3.1 Momentum transformer

*Integrating momentum into causal linear attention.* Now we consider integrating momentum into causal linear attention. We integrate momentum into the state  $\mathbf{s}_i$  in (69) only since the denominator in causal linear attention is simply a normalizing scalar. If we regard  $-\phi(\mathbf{k}_i)\mathbf{v}_i^\top$  as the gradient vector in (3), then we can add momentum into the state  $\mathbf{s}_i$  by following the heavy-ball method in (2), resulting in the following RNN formulation of causal momentum attention,

$$\begin{aligned}\mathbf{m}_i &= \beta\mathbf{m}_{i-1} - \phi(\mathbf{k}_i)\mathbf{v}_i^\top; \\ \mathbf{s}_i &= \mathbf{s}_{i-1} - \gamma\mathbf{m}_i; \\ \mathbf{z}_i &= \mathbf{z}_{i-1} + \phi(\mathbf{k}_i); \\ \hat{\mathbf{v}}_i &= \frac{\phi(\mathbf{q}_i)^\top \mathbf{s}_i}{\phi(\mathbf{q}_i)^\top \mathbf{z}_i},\end{aligned}\tag{70}$$

where  $\mathbf{m}_0 = \mathbf{0}$ , and  $\gamma > 0$  and  $0 \leq \beta < 1$  are two hyperparameters. The RNN formulation of causal momentum attention in (70) is efficient for autoregressive inference. For training, we need to rewrite (70) into a form that is similar to (68). To this end, we need to eliminate  $\mathbf{m}_i$ ,  $\mathbf{s}_i$ , and  $\mathbf{z}_i$  from (70). Note that

$$\mathbf{s}_i = \mathbf{s}_{i-1} - \underbrace{\gamma\mathbf{m}_i}_{:=\mathbf{p}_i} = \underbrace{\mathbf{s}_0}_{=\mathbf{0}} - (\mathbf{p}_i + \mathbf{p}_{i-1} + \cdots + \mathbf{p}_1),$$

since  $\mathbf{m}_i = \beta\mathbf{m}_{i-1} - \phi(\mathbf{k}_i)\mathbf{v}_i^\top$ , we have  $\mathbf{p}_i = \beta\mathbf{p}_{i-1} - \gamma\phi(\mathbf{k}_i)\mathbf{v}_i^\top$ . Therefore,

$$\begin{aligned}\mathbf{s}_i &= -(\mathbf{p}_i + \mathbf{p}_{i-1} + \cdots + \mathbf{p}_1) = \gamma\phi(\mathbf{k}_i)\mathbf{v}_i^\top - \left( (1+\beta)\mathbf{p}_{i-1} + \mathbf{p}_{i-2} + \cdots + \mathbf{p}_1 \right) \\ &= \gamma\phi(\mathbf{k}_i)\mathbf{v}_i^\top + \gamma(1+\beta)\phi(\mathbf{k}_i)\mathbf{v}_i^\top - \left( (1+\beta)^2\mathbf{p}_{i-2} + \cdots + \mathbf{p}_1 \right) \\ &= \cdots \\ &= \gamma \sum_{j=1}^i \frac{1-\beta^{i-j+1}}{1-\beta} \phi(\mathbf{k}_j)\mathbf{v}_j^\top \text{ for } i \geq 1.\end{aligned}$$

We can then formulate the causal momentum attention as follows

$$\hat{\mathbf{v}}_i = \frac{\gamma\phi(\mathbf{q}_i)^\top \sum_{j=1}^i \left( \frac{1-\beta^{i-j+1}}{1-\beta} \phi(\mathbf{k}_j)\mathbf{v}_j^\top \right)}{\phi(\mathbf{q}_i)^\top \mathbf{z}_i}.\tag{71}$$

Note that (71) is mathematically equivalent to (70), but it can be trained much more efficiently in a concurrent fashion via layerwise parallelism.

*Remark 7* Comparing (71) with (68), we see that momentum plays a role in reweighting the terms  $\{\phi(\mathbf{k}_j)\mathbf{v}_j^\top\}_{j=1}^i$ . It is interesting to note that this reweighting is opposite to that used for reweighting the local attention [18]. It has also been noticed that low-rank attention can complement local attention, resulting in improved performance [61].

*Integrating momentum into linear attention.* To obtain momentum attention without causal masking, we can simply take the sum from 1 to  $N$  instead of summing from 1 to  $i$ . Therefore, we obtain the following momentum attention

$$\hat{\mathbf{v}}_i = \frac{\gamma\phi(\mathbf{q}_i)^\top \sum_{j=1}^N \left( \frac{1-\beta^{N-j+1}}{1-\beta} \phi(\mathbf{k}_j)\mathbf{v}_j^\top \right)}{\phi(\mathbf{q}_i)^\top \sum_{j=1}^N \phi(\mathbf{k}_j)}. \quad (72)$$


*Memory and computational complexity.* Training momentum transformers have the same memory and computational complexities of  $\mathcal{O}(N)$  as the training of linear transformers. For test and inference, momentum transformers also have the same memory and computational complexities as linear transformers. However, in the RNN form, momentum transformers require slightly more memory than linear transformers to store the extra momentum state  $\mathbf{m}_i$ .

#### 4.3.2 Momentum connection

Each transformer layer has a residual connection between the self-attention output and the input as shown in (65). We further integrate momentum into (65) and derive the momentum connection as follows

$$T_\ell(\mathbf{X}) = f_\ell(\hat{\mathbf{V}} + \mathbf{X} + \tilde{\beta}(\mathbf{X} - T_{\ell-1}(\mathbf{X}))), \quad 0 \leq \tilde{\beta} < 1. \quad (73)$$

*Adaptive momentum.* Our momentum transformer introduces additional hyperparameters  $\gamma$  and  $\beta$ , as well as  $\tilde{\beta}$ , compared to the linear transformer. Often  $\gamma$  can be simply set to 1. However, tuning  $\beta$  and  $\tilde{\beta}$  can introduce extra computational cost for training transformers. Moreover, using a constant momentum may not give us optimal performance. In this part, we will introduce an adaptive momentum formula for computing the momentum hyperparameter in momentum connection and thus eliminating the computational overhead for tuning  $\tilde{\beta}$ . Here, the adaptive momentum does not apply to  $\beta$  since it will break the closed unrolling form in (71). Adaptive momentum has been used in optimization, see, e.g., [105], [88]; here, we use the later one for its simplicity.



**Fig. 14** Convergence comparison of adaptive momentum, momentum, reformer, linear, and softmax transformer on the sequence copy task. Momentum and adaptive momentum transformers converge faster and achieve better training loss than both linear transformer and reformer. Softmax transformer converges the fastest but suffers from quadratic memory and computational complexity. Adaptive momentum transformer performs as well as momentum transformer without intensively searching for momentum values.


#### 4.4 Experimental results

We evaluate the benefits of our momentum transformers in terms of convergence speed, efficiency, and accuracy. We compare the performance of momentum and adaptive momentum transformers with the baseline standard softmax transformer and several other efficient transformers in the following tasks: 1) the synthetic copy task, 2) the MNIST and CIFAR image generation task, 3) Long-Range Arena [92], and 4) the non-autoregressive machine translation task. These tasks are among standard benchmarks for measuring the performance of transformers and their efficiency. The tasks we choose also cover different data modalities - text and image - and a variety of model sizes. Our experimental results confirm that momentum and adaptive momentum transformers outperform many existing efficient transformers, including linear transformers and reformers, in accuracy and converge faster. Furthermore, adaptive momentum transformer improves over momentum transformer without the need of searching for momentum hyperparameter.

##### 4.4.1 Copy task

We train momentum transformers and baseline models on a synthetic copy task to analyze their convergence speed. In this task, the model has to duplicate a sequence of symbols. Each training and test sample has the form  $0w0w$  where  $w$  is a sequence of symbols collected from the set  $\{1, \dots, N\}$ .

In our experiments, we follow the same experimental setting as that used in [39]. In particular, we use a sequence of maximum length 128 with 10 different symbols separated by a separator symbol. The baseline architecture for all methods is a 4-layer transformer with 8 attention heads and  $D = 32$ . The models are trained with the RAdam optimizer using a batch size of 64 and a learning rate of  $10^{-3}$  which is reduced to  $10^{-4}$  after 3000 iterations. Figure 14 shows the training loss and the test accuracy over epochs and over GPU time. Both the momentum and the adaptive momentum transformers converge much faster and achieve better training loss than the linear transformer. Notice



**Fig. 15** Momentum transformers outperform linear transformers on the MNIST image generation task. Adaptive momentum transformer achieves the best test bits/dim.

Method	Bits/dim	Images/sec
Standard softmax transformer	0.84	0.45 (1×)
Linear transformer	0.85	142.8 (317×)
Momentum transformer	0.84	139.7 (310×)
Momentum transformer + momentum connection	0.82	135.5 (301×)
Adaptive momentum transformer	0.80	134.9 (300×)

**Table 6** Momentum transformers achieve better test bits/dim than both softmax and linear transformers on MNIST generation.

that while the standard transformer converges the fastest, it has quadratic complexity. Adaptive momentum transformer has similar performance as the momentum transformer without the need of tuning for the momentum value.

#### 4.4.2 Image generation

Transformers have shown great promise in autoregressive generation applications [73, 14], such as autoregressive image generation [75]. However, the training and sampling procedure using transformers are quite slow for these tasks due to the quadratic computational time complexity and the memory scaling with respect to the sequence length. In this section, we train our momentum-based transformers and the baselines with causal masking to predict images pixel by pixel and compare their performance. In particular, we demonstrate that, like linear transformers, both momentum and adaptive momentum transformers are able to generate images much faster than the standard softmax transformer. Furthermore, we show that momentum-based transformers converge much faster than linear transformers while achieving better bits per dimension (bits/dim). Momentum and adaptive momentum transformers also generate images with constant memory per image like linear transformers.

*MNIST.* We first examine our momentum-based transformers on the MNIST image generation task. For all methods, we train a 8-layer transformer with 8 attention heads and the embedding size of 256, which corresponds to 32 dimensions per head. The feedforward dimensions are 4 times larger than the

Method	Bits/dim	Images/sec
Standard softmax transformer	3.20	0.004 (1×)
Linear transformer	3.44	17.85 (4462×)
Momentum transformer	3.43	17.52 (4380×)
Momentum transformer + momentum connection	3.41	17.11 (4277×)
Adaptive momentum transformer	3.38	17.07 (4267×)

**Table 7** Momentum-based transformers achieve better test bits/dim than linear transformer on CIFAR10 image generation task.

Model	ListOps (2K)	Text (4K)	Retrieval (4K)	Image (1K)	Pathfinder (1K)	Avg
Softmax [98]	37.10 (37.10)	64.17 (65.02)	80.71 (79.35)	39.06 (38.20)	72.48 (74.16)	<b>58.70 (58.77)</b>
Linear [39]	18.30	64.22	81.37	38.29	71.17	54.67
Performer [16]	18.80	63.81	78.62	37.07	69.87	53.63
Reformer [44]	19.05	64.88	78.64	43.29	69.36	55.04
Linformer [107]	<b>37.25</b>	55.91	79.37	37.84	67.60	55.59
Momentum transformer	19.56	64.35	81.95	39.40	73.12	55.68
Adaptive momentum transformer	20.16	<b>64.45</b>	<b>82.07</b>	<b>39.53</b>	<b>74.00</b>	56.04

**Table 8** Results on the LRA tasks. We report the test classification accuracy for each task and average accuracy across all tasks. The momentum-based transformers, in particular, the adaptive momentum transformer, outperforms all other transformers except on the ListOps. The numbers in the parenthesis are from the paper [111]. Unit: %.

embedding size. A mixture of 10 logistics is used to model the output as in [80]. For training, we use the RAdam optimizer with a learning rate of  $10^{-4}$  and train all models for 250 epochs except for the adaptive momentum transformer.

We report the bits/dim and image generation throughput in Table 6. Compared to the linear transformer, all momentum-based transformers not only attain better bits/dim but also have comparable image generation throughput, justifying the linear complexity of our models. In addition, we demonstrate that the adaptive momentum transformer converges much faster than the baseline models in Figure 15. Momentum-based transformers even outperform softmax transformers in this task.

*CIFAR10.* Next, we investigate the advantages of our momentum-based transformers when the sequence length and the number of layers in the model increase. We consider the CIFAR-10 image generation task, in which we train 16-layer transformers to generate CIFAR-10 images. The configuration for each layer is the same as in the MNIST experiment. For the linear transformer and our momentum-based transformer, we use a batch size of 4 while using a batch size of 1 for the standard softmax transformer due to the memory limit of the largest GPU available to us, i.e., NVIDIA V100. This is similar to the setting in [39]. Like in the MNIST image generation task, our momentum-based transformers outperform the linear transformer in terms of bits/dim while maintaining comparable image generation throughput. This is a very expensive task, limiting us to perform a thorough hyperparameter search; we believe better results can be obtained with a more thorough hyperparameter search.



Method	BLEU Score	Speed (tokens/s)
Standard softmax transformer	24.34	5104
Linear transformer	21.37	1382
Momentum transformer	22.11	1398
Momentum transformer + momentum connection	22.14	1403
Adaptive momentum transformer	22.20	1410

**Table 9** BLEU scores and tokens per second from machine translation models trained on IWSLT show the advantages of our momentum-based transformers. The number of trainable parameters is almost the same for all models, up to the small difference introduced by the momentum mechanism in our models. Momentum-based transformers outperform the linear transformer in generation quality in terms of BLEU score and obtain comparable generation efficiency in terms of tokens per second.

#### 4.4.3 Long-Range Arena

In this experiment, we evaluate our model on tasks that involve longer sequence lengths in the Long Range Arena (LRA) benchmark [92]. We show that the momentum-based transformer outperforms the baseline linear transformer and standard softmax transformer [98], justifying the advantage of our momentum-based transformers in capturing long-term dependency.

**Datasets and metrics.** We consider all five tasks in the LRA benchmark [92], including Listops, byte-level IMDb reviews text classification, byte-level document retrieval, CIFAR-10 classification on sequences of pixels, and Pathfinder. These tasks involve long sequences of length  $2K$ ,  $4K$ ,  $4K$ ,  $1K$ , and  $1K$ , respectively. We follow the setup/evaluation protocol in [92] and report test accuracy for each task and the average result across all tasks.

**Models and training.** All models have 2 layers, 64 embedding dimension, 128 hidden dimension, 2 attention heads. Mean pooling is applied in all models. Also, we use the nonlinear activation  $elu(x) + 1$  for the linear transformer. Our implementation uses the public code in [111] as a starting point, and we follow their training procedures. The training setting and additional baseline model details are provided in the configuration file used in [111].

**Results.** We summarize our results in Table 8. Both momentum-based transformers outperform linear transformers in all tasks and yield better accuracy than the standard softmax transformer in most tasks except the Listops. The adaptive momentum transformer performs the best on every task except the LipsOps, far behind the softmax transformer and Linformer.

#### 4.4.4 Non-autoregressive machine translation

All of the above experiments are for auto-regressive tasks. In this last experiment, we demonstrate that the benefits of our momentum-based transformers also hold for a non-autoregressive task. We consider a machine translation task on the popular IWSLT’16 En-De dataset. We follow the setting in [50]. In particular, we tokenize each sentence using a script from Moses [45] and segment each word into subword units using BPE [83]. We also use  $40K$  tokens from both source and target. Our baseline model is the small transformer-

based network in [50]. This model has 5 layers, and each layer has 2 attention heads. We replace the softmax attention in this network with the linear and momentum-based attention to obtain the linear transformer baseline and the momentum-based transformer models, respectively.

Table 9 reports the results in terms of generation quality, measured by the BLEU score [65], and generation efficiency, measured by the number of generated tokens per second. Consistent with other experiments above, our momentum-based transformers obtain better BLEU scores than the linear transformer in this non-autoregressive setting. Furthermore, in terms of generation efficiency, momentum-based models are comparable with the linear transformer and much more efficient than the standard softmax transformer.

## 5 Conclusion and Future Work

In this paper, we reviewed how to integrate momentum into neural networks to enhance their theoretical and practical performances. In particular, we showed that momentum improves learning long-term dependencies of RNNs and neural ODEs and significantly reduces their computational costs. Moreover, we showed that momentum can also be used to improve the efficiency and accuracy of transformers. There are numerous directions for future work: 1) Can we leverage the momentum-augmented neural network component to aid the neural architecture search? 2) Can we further improve the momentum-integrated architectures by using the numerical ODE insights [30]? 3) Momentum has also been used in designing CNNs [52, 81]; it is also worth further studying the benefits of momentum for CNNs.

## Acknowledgement

This material is based on research sponsored by NSF grants DMS-1924935 and DMS-1952339, DOE grant DE-SC0021142, and ONR grant N00014-18-1-2527 and the ONR MURI grant N00014-20-1-2787.

## References

1. Joshua Ainslie, Santiago Ontanon, Chris Alberti, Vaclav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. ETC: Encoding long and structured inputs in transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 268–284, November 2020.
2. Rami Al-Rfou, DK Choe, Noah Constant, Mandy Guo, and Llion Jones. Character-level language modeling with deeper self-attention. In *Thirty-Third AAAI Conference on Artificial Intelligence*, 2019.
3. Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *International Conference on Machine Learning*, pages 1120–1128, 2016.
4. Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.

5. Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994.
6. Guy Blanc and Steffen Rendle. Adaptive sampled softmax with kernel based sampling. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 590–599. PMLR, 10–15 Jul 2018.
7. Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym, 2016. cite arxiv:1606.01540.
8. Tom Brown and et al. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901, 2020.
9. Mario Lezcano Casado. Optimization with orthogonal constraints and on general manifolds. <https://github.com/Lezcano/exprNN>, 2019.
10. Mario Lezcano Casado. Trivializations for gradient-based optimization on manifolds. In *Advances in Neural Information Processing Systems*, pages 9154–9164, 2019.
11. Anirban Chakraborty, Manaar Alam, Vishal Dey, Anupam Chattopadhyay, and Debdeep Mukhopadhyay. Adversarial attacks and defences: A survey. *arXiv preprint arXiv:1810.00069*, 2018.
12. Sarath Chandar, Chinnadhurai Sankar, Eugene Vorontsov, Samira Ebrahimi Kahou, and Yoshua Bengio. Towards non-saturating recurrent units for modelling long-term dependencies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3280–3287, 2019.
13. Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 6572–6583, 2018.
14. Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
15. Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
16. Krzysztof Marcin Choromanski and et al. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021.
17. Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. Gated feed-back recurrent neural networks. In *International conference on machine learning*, pages 2067–2075, 2015.
18. Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.
19. Talgat Daulbaev, Alexandr Katrutsa, Larisa Markeeva, Julia Gusak, Andrzej Cichocki, and Ivan Oseledets. Interpolation technique to speed up gradients propagation in neural odes. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16689–16700. Curran Associates, Inc., 2020.
20. Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018.
21. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
22. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
23. Jianzhun Du, Joseph Futoma, and Finale Doshi-Velez. Model-based reinforcement learning for semi-markov decision processes with neural odes. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 19805–19816. Curran Associates, Inc., 2020.

24. Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
25. Jeffrey L Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.
26. Chris Finlay, Joern-Henrik Jacobsen, Levon Nurbekyan, and Adam Oberman. How to train your neural ODE: the world of Jacobian and kinetic regularization. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3154–3164. PMLR, 13–18 Jul 2020.
27. John S Garofolo. Timit acoustic phonetic continuous speech corpus. *Linguistic Data Consortium, 1993*, 1993.
28. Arnab Ghosh, Harkirat Behl, Emilien Dupont, Philip Torr, and Vinay Nambodiri. Steer : Simple temporal regularization for neural ode. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 14831–14843. Curran Associates, Inc., 2020.
29. Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, and David Duvenaud. Scalable reversible generative models with free-form continuous dynamics. In *International Conference on Learning Representations*, 2019.
30. Eldad Haber and Lars Ruthotto. Stable architectures for deep neural networks. *Inverse Problems*, 34(1):014004, 2017.
31. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
32. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer, 2016.
33. Kyle Helfrich, Devin Willmott, and Qiang Ye. Orthogonal recurrent neural networks with scaled Cayley transform. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1969–1978, Stockholm, Sweden, 10–15 Jul 2018. PMLR.
34. Mikael Henaff, Arthur Szlam, and Yann LeCun. Recurrent orthogonal networks and long-memory tasks. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2034–2042, New York, New York, USA, 20–22 Jun 2016. PMLR.
35. Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.
36. Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
37. Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia, July 2018. Association for Computational Linguistics.
38. Cheng-Zhi Anna Huang, Ashish Vaswani, Jakob Uszkoreit, Ian Simon, Curtis Hawthorne, Noam Shazeer, Andrew M Dai, Matthew D Hoffman, Monica Dinulescu, and Douglas Eck. Music transformer: Generating music with long-term structure. In *International Conference on Learning Representations*, 2018.
39. Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pages 5156–5165. PMLR, 2020.
40. Jacob Kelly, Jesse Bettencourt, Matthew J Johnson, and David K Duvenaud. Learning differential equations that are easy to solve. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4370–4380. Curran Associates, Inc., 2020.
41. Patrick Kidger, James Morrill, James Foster, and Terry J. Lyons. Neural controlled differential equations for irregular time series. In *NeurIPS*, 2020.
42. Yoon Kim, Carl Denton, Luong Hoang, and Alexander M Rush. Structured attention networks. *arXiv preprint arXiv:1702.00887*, 2017.

43. Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
44. Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020.
45. Philipp Koehn and et al. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, June 2007.
46. Quoc V Le, Navdeep Jaitly, and Geoffrey E Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015.
47. Mathias Lechner and Ramin Hasani. Learning long-term dependencies in irregularly-sampled time series, 2020.
48. Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
49. Yann LeCun, Corinna Cortes, and CJ Burges. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
50. Jason Lee, Elman Mansimov, and Kyunghyun Cho. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182, October–November 2018.
51. Mario Lezcano-Casado and David Martínez-Rubio. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. In *International Conference on Machine Learning (ICML)*, pages 3794–3803, 2019.
52. Huan Li, Yibo Yang, Dongmin Chen, and Zhouchen Lin. Optimization algorithm inspired deep neural network structure design. In *Asian Conference on Machine Learning*, pages 614–629. PMLR, 2018.
53. Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. Generating wikipedia by summarizing long sequences. In *International Conference on Learning Representations*, 2018.
54. Peter J Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198*, 2018.
55. A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
56. Stefano Massaroli, Michael Poli, Jinkyoo Park, Atsushi Yamashita, and Hajime Asma. Dissecting neural odes. In *34th Conference on Neural Information Processing Systems, NeurIPS 2020*. The Neural Information Processing Systems, 2020.
57. Zakaria Mhammedi, Andrew Hellicar, Ashfaqur Rahman, and James Bailey. Efficient orthogonal parametrisation of recurrent neural networks using householder reflections. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2401–2409. JMLR. org, 2017.
58. Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
59. Yurii E Nesterov. A method for solving the convex programming problem with convergence rate  $o(1/k^2)$ . In *Dokl. Akad. Nauk Sssr*, volume 269, pages 543–547, 1983.
60. Tan Nguyen, Richard Baraniuk, Andrea Bertozzi, Stanley Osher, and Bao Wang. MomentumRNN: Integrating Momentum into Recurrent Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS 2020)*, 2020.
61. Tan M Nguyen, Vai Suliafu, Stanley J Osher, Long Chen, and Bao Wang. Fmmformer: Efficient and flexible transformer via decomposed near-field and far-field attention. *arXiv preprint arXiv:2108.02347*, 2021.
62. Jean-Philippe Noël and M Schoukens. F-16 aircraft benchmark based on ground vibration test data. In *2017 Workshop on Nonlinear System Identification Benchmarks*, pages 19–23, 2017.
63. Alexander Norcliffe, Cristian Bodnar, Ben Day, Nikola Simidjievski, and Pietro Liò. On second order behaviour in augmented neural odes. In *Advances in Neural Information Processing Systems*, 2020.

64. Avik Pal, Yingbo Ma, Viral Shah, and Christopher V Rackauckas. Opening the black-box: Accelerating neural differential equations by regularizing internal solver heuristics. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8325–8335. PMLR, 18–24 Jul 2021.
65. Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, page 311–318, USA, 2002. Association for Computational Linguistics.
66. Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4055–4064. PMLR, 10–15 Jul 2018.
67. Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pages 1310–1318, 2013.
68. Hao Peng, Nikolaos Pappas, Dani Yogatama, Roy Schwartz, Noah Smith, and Lingpeng Kong. Random feature attention. In *International Conference on Learning Representations*, 2021.
69. Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
70. Lev Semenovich Pontryagin. *Mathematical theory of optimal processes*. Routledge, 2018.
71. Jiezhong Qiu, Hao Ma, Omer Levy, Scott Wen-tau Yih, Sinong Wang, and Jie Tang. Blockwise self-attention for long document understanding. *arXiv preprint arXiv:1911.02972*, 2019.
72. Alessio Quaglino, Marco Gallieri, Jonathan Masci, and Jan Koutník. Snode: Spectral discretization of neural odes for system identification. In *International Conference on Learning Representations*, 2020.
73. Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
74. Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas, November 2016. Association for Computational Linguistics.
75. Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Mark Chen, Rewon Child, Vedant Misra, Pamela Mishkin, Gretchen Krueger, Sandhini Agarwal, and Ilya Sutskever. Dall-e: Creating images from text. *OpenAI blog*, 2020.
76. Ankit Singh Rawat, Jiecao Chen, Felix Xinnan X Yu, Ananda Theertha Suresh, and Sanjiv Kumar. Sampled softmax with random fourier features. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
77. Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics*, 9:53–68, 2021.
78. Yulia Rubanova, Ricky T. Q. Chen, and David K Duvenaud. Latent ordinary differential equations for irregularly-sampled time series. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
79. Salesforce. Lstm and qrn language model toolkit for pytorch. <https://github.com/salesforce/awd-lstm-lm>, 2017.
80. Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. In *International Conference on Learning Representations*, 2017.
81. Michael E. Sander, Pierre Ablin, Mathieu Blondel, and Gabriel Peyré. Momentum residual neural networks. *arXiv preprint arXiv:2102.07870*, 2021.

82. Imanol Schlag, Kazuki Irie, and Jürgen Schmidhuber. Linear transformers are secretly fast weight memory systems. *CoRR*, abs/2102.11174, 2021.
83. Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, August 2016.
84. Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Efficient attention: Attention with linear complexities. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3531–3539, 2021.
85. David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
86. David R So, Chen Liang, and Quoc V Le. The evolved transformer. *arXiv preprint arXiv:1901.11117*, 2019.
87. Kyungwoo Song, Yohan Jung, Dongjun Kim, and Il-Chul Moon. Implicit kernel attention. *arXiv preprint arXiv:2006.06147*, 2021.
88. Tao Sun and Huaming Ling and Zuoqiang Shi and Dongsheng Li and Bao Wang. Training Deep Neural Networks with Adaptive Momentum Inspired by the Quadratic Optimization. *arXiv preprint arXiv:2110.09057*, 2021.
89. C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
90. Yi Tay, Dara Bahri, Donald Metzler, Da-Cheng Juan, Zhe Zhao, and Che Zheng. Synthesizer: Rethinking self-attention in transformer models. *arXiv preprint arXiv:2005.00743*, 2020.
91. Yi Tay, Dara Bahri, Liu Yang, Donald Metzler, and Da-Cheng Juan. Sparse Sinkhorn attention. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9438–9447. PMLR, 13–18 Jul 2020.
92. Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena : A benchmark for efficient transformers. In *International Conference on Learning Representations*, 2021.
93. Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*, 2020.
94. T. Tieleman and G. Hinton. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning, 2012.
95. Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.
96. Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877*, 2020.
97. Jos Van Der Westhuizen and Joan Lasenby. The unreasonable effectiveness of the forget gate. *arXiv preprint arXiv:1804.04849*, 2018.
98. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
99. Eugene Vorontsov, Chiheb Trabelsi, Samuel Kadoury, and Chris Pal. On orthogonality and learning recurrent networks with long term dependencies. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3570–3578. JMLR. org, 2017.
100. Apoorv Vyas, Angelos Katharopoulos, and François Fleuret. Fast transformers with clustered attention. *Advances in Neural Information Processing Systems*, 33, 2020.
101. Bao Wang, Alex Lin, Penghang Yin, Wei Zhu, Andrea L. Bertozzi, and Stanley J. Osher. Adversarial defense via the data-dependent activation, total variation minimization, and adversarial training. *Inverse Problems & Imaging*, 15(1):129–145, 2021.

102. Bao Wang, Xiyang Luo, Zhen Li, Wei Zhu, Zuoqiang Shi, and Stanley Osher. Deep neural nets with interpolating function as output activation. *Advances in Neural Information Processing Systems*, 2018.
103. Bao Wang, Tan M Nguyen, Andrea L Bertozzi, Richard G Baraniuk, and Stanley J Osher. Scheduled restart momentum for accelerated stochastic gradient descent. *arXiv preprint arXiv:2002.10583*, 2020.
104. Bao Wang and Stan J. Osher. Graph interpolating activation improves both natural and robust accuracies in data-efficient deep learning. *European Journal of Applied Mathematics*, 32(3):540–569, 2021.
105. Bao Wang and Qiang Ye. Stochastic gradient descent with nonlinear conjugate gradient-style adaptive momentum. *arXiv preprint arXiv:2012.02188*, 2020.
106. Bao Wang, Binjie Yuan, Zzuoqiang Shi, and Stanley Osher. Resnets ensemble via the feynman-kac formalism to improve natural and robust accuracies. *Advances in Neural Information Processing Systems*, 2019.
107. Sinong Wang, Belinda Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
108. Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, June 2018.
109. Scott Wisdom, Thomas Powers, John Hershey, Jonathan Le Roux, and Les Atlas. Full-capacity unitary recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 4880–4888, 2016.
110. Hedi Xia, Vai Suliafu, Hangjie Ji, Tan M. Nguyen, Andrea L. Bertozzi, Stanley J. Osher, and Bao Wang. Heavy ball neural ordinary differential equations. *arXiv preprint arXiv:2010.04840*, 2021.
111. Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nyströmformer: A nyström-based algorithm for approximating self-attention. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
112. Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big bird: Transformers for longer sequences. *arXiv preprint arXiv:2007.14062*, 2021.
113. Tianjun Zhang, Zhewei Yao, Amir Gholami, Joseph E Gonzalez, Kurt Keutzer, Michael W Mahoney, and George Biros. ANODEV2: A Coupled Neural ODE Framework. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
114. Juntang Zhuang, Nicha C Dvornek, sekhar tatikonda, and James s Duncan. {MALI}: A memory efficient and reverse accurate integrator for neural {ode}s. In *International Conference on Learning Representations*, 2021.