

Provable Acceleration of Nesterov’s Accelerated Gradient Method over Heavy Ball Method in Training Over-Parameterized Neural Networks

Xin Liu¹, Wei Tao^{2*}, Wei Li¹, Dazhi Zhan¹, Jun Wang¹, Zhisong Pan^{1*}

¹Army Engineering University of PLA

²Strategic Assessments and Consultation Institute, Academy of Military Science

Abstract

Due to its simplicity and efficiency, the first-order gradient method has been extensively employed in training neural networks. Although the optimization problem of the neural network is non-convex, recent research has proved that the first-order method is capable of attaining a global minimum during training over-parameterized neural networks, where the number of parameters is significantly larger than that of training instances. Momentum methods, including the heavy ball (HB) method and Nesterov’s accelerated gradient (NAG) method, are the workhorse of first-order gradient methods owing to their accelerated convergence. In practice, NAG often exhibits superior performance than HB. However, current theoretical works fail to distinguish their convergence difference in training neural networks. To fill this gap, we consider the training problem of the two-layer ReLU neural network under over-parameterization and random initialization. Leveraging high-resolution dynamical systems and neural tangent kernel (NTK) theory, our result not only establishes tighter upper bounds of the convergence rate for both HB and NAG, but also provides the first theoretical guarantee for the acceleration of NAG over HB in training neural networks. Finally, we validate our theoretical results on three benchmark datasets.

1 Introduction

Nowadays, the first-order gradient method is widely adopted as the standard choice for training neural networks [Schmidhuber, 2015]. The gradient descent (GD) method is the most famous first-order method, which only accesses the current gradient to update the model. To accelerate its convergence, the momentum method is proposed by utilizing the historical gradients. The idea of incorporating momentum term into GD can be traced back to 1964, when Polyak introduced the heavy ball (HB) method for solving the quadratic convex

problem [Polyak, 1964]. Compared to GD, HB achieves a local accelerated convergence rate in a neighborhood around the global optimal. Subsequently, Nesterov proposed a different type of momentum method known as Nesterov’s accelerated gradient (NAG) method [Nesterov, 1983]. In the convex setting, NAG achieves the optimal global convergence rate [Nesterov, 2013], which improves the local convergence of HB.

Nowadays, neural networks have gained remarkable performance in a various tasks owing to their over-parameterization, such as GPT-3 with 175 billion parameters [Brown *et al.*, 2020]. As a result, the training process of neural networks suffers from significant computation costs. To save the expensive computational resources, momentum methods are widely adopted in the deep learning community due to their acceleration over GD. Besides, NAG converges faster than HB in practical neural network training [Schmidt *et al.*, 2021; Nado *et al.*, 2021]. Unfortunately, this phenomenon lacks sufficient understandings mainly due to the non-convexity faced by the optimization with neural networks. Recently, some advances have been achieved in providing explanations for the convergence of GD in training over-parameterized neural networks. Later, some works extended these findings to momentum methods [Wang *et al.*, 2021; Bu *et al.*, 2021; Liu *et al.*, 2022b]. However, their analysis only establishes convergence guarantees, but stops short of distinguishing the convergence difference between HB and NAG. Specifically, existing works can be roughly categorized into two parts: discrete-time and continuous-time. For the discrete-time, [Wang *et al.*, 2021] proved that HB can reach a global minimum when training two-layer neural networks. Likewise, [Liu *et al.*, 2022b] established the global convergence guarantee for NAG. These two researches only provide convergence results with specific hyper-parameters, but it is still unclear whether NAG theoretically converges faster than HB in training neural networks. From the continuous-time, [Bu *et al.*, 2021] exploited the connection between the limiting ordinary differential equations (ODEs) and the momentum methods. However, their result relies on a low-resolution ODE, which fails to differentiate the convergence behavior between HB and NAG. As a result, there still exists a gap between practice and theory.

Motivated by the fact that the high-resolution ODEs provide more accurate characterizations than low-resolution

*Corresponding Authors.

Full version is available at <https://arxiv.org/abs/2208.03941>.

ODEs, we study the training dynamics of momentum methods from a high-resolution perspective and establish an explicit convergence difference between HB and NAG in training neural networks. For the architecture of the neural network, we focus on an over-parameterized two-layer neural network with ReLU activation, which is extensively studied in prior works [Du *et al.*, 2019c; Wang *et al.*, 2021; Liu *et al.*, 2022b]. Our contributions are summarized as follows:

- Firstly, we derive the residual dynamics of HB and NAG by applying high-resolution ODEs and neural tangent kernel (NTK) theory, which reveals that the original non-convex weight dynamics corresponds to a strongly convex residual dynamics under over-parameterized assumption.
- Secondly, with the Lyapunov-based approach, we investigate the residual dynamics via a sharper Young’s inequality instead of Cauchy-Schwarz inequality. Compared to existing works, we improve the convergence rates of HB and NAG with tighter upper bounds. Besides, our result provides the first theoretical guarantee for the acceleration of NAG over HB, which is mainly caused by the extra gradient correction term of NAG.
- Finally, we validate the findings through a convergence comparison between HB and NAG under different settings of hyper-parameters on three benchmark datasets. The maximum distance between the training iterate and its initial value is also tested to support our theory.

2 Related Works

2.1 Momentum Methods

Momentum methods apply the historical information of gradients to accelerate GD. For the quadratic convex problem, Polyak proved HB achieves a local iteration complexity $\mathcal{O}(\sqrt{\kappa} \log(1/\epsilon))$ [Polyak, 1964], where $\kappa > 1$ is the condition number of the quadratic problem. Compared to the iteration complexity $\mathcal{O}(\kappa \log(1/\epsilon))$ of GD, HB achieves an acceleration. For L -smooth and μ -strongly convex problems, [Ghadimi *et al.*, 2015] proved that the average iterates of HB converge to a global optimal with $\mathcal{O}(\frac{L}{\mu} \log(1/\epsilon))$ iteration complexity, similar to the convergence of GD. Nevertheless, [Lessard *et al.*, 2016] discovered a counterexample that HB fails to converge in tackling some smooth and strongly convex problems. In contrast, NAG enjoys a global iteration complexity $\mathcal{O}(\sqrt{\kappa} \log(1/\epsilon))$ in the same setting [Nesterov, 2013].

Meanwhile, researchers investigated the connection between discrete momentum methods and ODEs. With the step size towards infinitesimal, [Wilson *et al.*, 2021] derived HB and NAG share the same ODE formulation. [Su *et al.*, 2014] obtained the limiting ODE for NAG with a time-varying momentum parameter. Subsequently, [Shi *et al.*, 2021] applied dimensional analysis to establish high-resolution ODEs for momentum methods, where HB and NAG exhibit different limiting ODEs. However, in the continuous viewpoint, [Shi *et al.*, 2021] only proved HB and NAG have the same convergence rate $\mathcal{O}(e^{-\sqrt{\mu}t/4})$ for L -smooth and μ -strongly con-

vex problems. By discretizing the Lyapunov function and the high-resolution ODEs, Shi further proved that NAG converges faster than HB in the strongly convex setting. Additionally, [Sun *et al.*, 2020] improved above continuous-time result of NAG with a faster $\mathcal{O}(e^{-3\sqrt{\mu}t/7})$ convergence rate. Recently, [Zhang *et al.*, 2021] further derived a sharper $\mathcal{O}(e^{-\sqrt{\mu}t/2})$ result.

	Obj	HB	NAG
[Shi <i>et al.</i> , 2021]	SC	$\mathcal{O}(e^{-\frac{\sqrt{\mu}t}{4}})$	$\mathcal{O}(e^{-\frac{\sqrt{\mu}t}{4}})$
[Sun <i>et al.</i> , 2020]	SC	-	$\mathcal{O}(e^{-\frac{3\sqrt{\mu}t}{7}})$
[Zhang <i>et al.</i> , 2021]	SC	-	$\mathcal{O}(e^{-\frac{\sqrt{\mu}t}{2}})$
This work	NC	$\mathcal{O}(e^{-(2-\sqrt{2})\sqrt{\mu}t})$	$\mathcal{O}(e^{-\rho_{NAG}\sqrt{\mu}t})$

Table 1: Comparison of existing results for HB and NAG from the high-resolution view, where SC, NC and Obj represent strongly-convex, non-convex and Objective function, respectively. In addition, $2 - \sqrt{2} < \rho_{NAG} \leq \frac{11 - \sqrt{73}}{4}$.

2.2 Neural Tangent Kernel

[Jacot *et al.*, 2018] established the connection between the optimization of infinite wide neural networks and NTK. Later, this result was further extended in providing convergence guarantees of GD for different architectures of neural networks, including convolutional neural network [Arora *et al.*, 2019a], recurrent neural network [Alemohammad *et al.*, 2020], residual network [Huang *et al.*, 2020] and graph neural network [Du *et al.*, 2019a].

With the NTK theory, there is an increasing interest in explaining the acceleration phenomenon of momentum methods in training neural networks. From the discrete-time view, [Wang *et al.*, 2021] established the first convergence result of HB in training two-layer ReLU neural networks and deep linear neural networks. Besides, their result theoretically shows the acceleration guarantee of HB over GD. In addition, [Liu *et al.*, 2022b] presented the convergence result of NAG in training two-layer neural networks. From the continuous-time view, [Bu *et al.*, 2021] exploited the limiting ODEs in analyzing the convergence of momentum methods for training neural networks. However, their result relies on low-resolution ODEs, which fails to capture the convergence difference between HB and NAG. At the same setting, our work differentiates the convergence behaviors between HB and NAG through sharper inequality and analysis. Specifically, our result shows that HB converges at a $\mathcal{O}(e^{-(2-\sqrt{2})\sqrt{\mu}t})$ convergence rate and NAG enjoys a faster $\mathcal{O}(e^{-\rho_{NAG}\sqrt{\mu}t})$ rate, where $2 - \sqrt{2} < \rho_{NAG} \leq (11 - \sqrt{73})/4$. The comparison of existing results for HB and NAG from the high-resolution view can be founded in Table 1.

3 Preliminary

3.1 Problem Setting

In this paper, we focus on the supervised learning problem for training a two-layer ReLU neural network with m hidden

nodes:

$$f(\mathbf{W}, \mathbf{a}, \mathbf{x}) = \frac{1}{\sqrt{m}} \sum_{r=1}^m a_r \sigma(\mathbf{w}_r^\top \mathbf{x}), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ denotes the input feature, $\sigma(z) = \max\{0, z\}$ denotes the ReLU activation function, $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_m) \in \mathbb{R}^{d \times m}$ is the parameter of the hidden layer and $\mathbf{a} \in \mathbb{R}^m$ is the parameter of the output layer. For simplicity, we assume the input feature has unit norm, i.e., $\|\mathbf{x}\|_2 = 1$, which follows [Bu *et al.*, 2021; Du *et al.*, 2019c]. For initialization, $\mathbf{w}_r(0) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $a_r \sim \text{Rademacher}(0.5)$ for any $r \in [m]$. According to [Wang *et al.*, 2021; Bu *et al.*, 2021; Du *et al.*, 2019c], we use the square loss as the loss function. Then, it can find the optimal parameter by minimizing the empirical risk ℓ to :

$$\ell(\mathbf{W}, \mathbf{a}) = \frac{1}{2} \sum_{i=1}^n (f(\mathbf{W}, \mathbf{a}, \mathbf{x}_i) - y_i)^2, \quad (2)$$

where $(\mathbf{x}_i, y_i) \in \mathbb{R}^d \times \mathbb{R}$ denotes the i -th training instance.

Following [Bu *et al.*, 2021; Du *et al.*, 2019c], we only update \mathbf{W} and keep \mathbf{a} fixed during training. Note that the problem in Eq.(2) is both non-convex and non-smooth. Recently, some works provided global convergence guarantees for gradient-based methods in optimizing Eq.(2) [Bu *et al.*, 2021; Du *et al.*, 2019c]. GD is the most widely used gradient-based method due to its simplicity and efficiency, which updates as

$$\mathbf{w}_r(t+1) = \mathbf{w}_r(t) - \eta \frac{\partial \ell(\mathbf{W}(t), \mathbf{a})}{\partial \mathbf{w}_r(t)}, \quad (3)$$

where $\eta > 0$ is the learning rate.

For accelerating GD, momentum methods blend history information of gradients into current updates. HB starts from the initial parameter $\mathbf{w}_r(-1) = \mathbf{w}_r(0)$ and updates as follows

$$\mathbf{w}_r(t+1) = \mathbf{w}_r(t) + \beta(\mathbf{w}_r(t) - \mathbf{w}_r(t-1)) - \eta \frac{\partial \ell(\mathbf{W}(t), \mathbf{a})}{\partial \mathbf{w}_r(t)}, \quad (4)$$

where $\beta \in [0, 1)$ denotes the momentum parameter. According to [Nesterov, 2013], NAG has several variants with different schemes of hyper-parameters. In this paper, we mainly focus on NAG with a constant momentum parameter β . Given the initial parameters $\mathbf{w}_r(-1) = \mathbf{w}_r(0)$, NAG involves the update procedure as

$$\begin{aligned} \mathbf{w}_r(t+1) = & \mathbf{w}_r(t) + \beta(\mathbf{w}_r(t) - \mathbf{w}_r(t-1)) - \eta \frac{\partial \ell(\mathbf{W}(t), \mathbf{a})}{\partial \mathbf{w}_r(t)} \\ & - \beta \eta \left(\frac{\partial \ell(\mathbf{W}(t), \mathbf{a})}{\partial \mathbf{w}_r(t)} - \frac{\partial \ell(\mathbf{W}(t-1), \mathbf{a})}{\partial \mathbf{w}_r(t-1)} \right). \end{aligned} \quad (5)$$

Compared to HB, NAG has an extra term $\beta \eta \left(\frac{\partial \ell(\mathbf{W}(t), \mathbf{a})}{\partial \mathbf{w}_r(t)} - \frac{\partial \ell(\mathbf{W}(t-1), \mathbf{a})}{\partial \mathbf{w}_r(t-1)} \right)$, which is associated with the difference between two consecutive gradients and is referred to as gradient correction [Shi *et al.*, 2021]. This term is crucial for the provable acceleration of NAG over HB in the convex setting, but lacks enough understandings in the non-convex setting.

3.2 Continuous View of Momentum Methods

Continuous view of GD. Before introducing the limiting ODEs of momentum methods, we present existing result of GD in training two-layer ReLU neural networks.

When learning rate tends to infinitesimal, GD corresponds to the gradient flow

$$\frac{d\mathbf{w}_r(t)}{dt} = - \frac{\partial \ell(\mathbf{W}(t), \mathbf{a})}{\partial \mathbf{w}_r(t)}. \quad (6)$$

Based on Eq.(2), the time derivative of the hidden layer has

$$\frac{d\mathbf{w}_r(t)}{dt} = - \frac{\partial \ell(\mathbf{W}(t), \mathbf{a})}{\partial \mathbf{w}_r(t)} = - \left(\frac{\partial \mathbf{f}(t)}{\partial \mathbf{w}_r(t)} \right)^\top (\mathbf{f}(t) - \mathbf{y}), \quad (7)$$

where $\mathbf{f}(t) = (f_1, \dots, f_n)^\top \in \mathbb{R}^n$ with $f_i = f(\mathbf{W}(t), \mathbf{a}; \mathbf{x}_i)$ and $\mathbf{y} = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$. Denotes $\Delta = \mathbf{f} - \mathbf{y}$ as the residual vector, it has

$$\frac{d\mathbf{f}(t)}{dt} = \sum_r \frac{\partial \mathbf{f}(t)}{\partial \mathbf{w}_r(t)} \frac{d\mathbf{w}_r(t)}{dt} = -\mathbf{H}(t)(\mathbf{f}(t) - \mathbf{y}), \quad (8)$$

where $\mathbf{H}(t)$ is a Gram matrix induced by the two-layer ReLU neural network with following formulation

$$\mathbf{H}(t) := \sum_{r=1}^m \frac{\partial \mathbf{f}(t)}{\partial \mathbf{w}_r(t)} \left(\frac{\partial \mathbf{f}(t)}{\partial \mathbf{w}_r(t)} \right)^\top. \quad (9)$$

When m goes to infinity, \mathbf{H} at the initialization has

$$\begin{aligned} \mathbf{H}_{ij}^\infty &:= \lim_{m \rightarrow \infty} \mathbf{H}_{ij}(0) \\ &= \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\mathbf{x}_i^\top \mathbf{x}_j \mathbb{I}\{\mathbf{w}^\top \mathbf{x}_i \geq 0, \mathbf{w}^\top \mathbf{x}_j \geq 0\}], \end{aligned} \quad (10)$$

which holds for any $i, j \in [n]$. Note that \mathbf{H}^∞ is positive definite with following lemma.

Lemma 1 (Theorem 3.1 in [Du *et al.*, 2019c]). *Suppose \mathbf{x}_i is not parallel with \mathbf{x}_j for any $i \neq j$, then $\lambda_0 := \lambda_{\min}(\mathbf{H}^\infty) > 0$.*

When m is large enough, $\mathbf{H}(0)$ is also positive definite.

Lemma 2 (Lemma 3.1 in [Du *et al.*, 2019c]). *If $m = \Omega\left(\frac{n^2}{\lambda_0^2} \log\left(\frac{n^2}{p}\right)\right)$, with probability at least $1 - p$, it has $\|\mathbf{H}(0) - \mathbf{H}^\infty\|_2 \leq \frac{\lambda_0}{4}$ and $\lambda_{\min}(\mathbf{H}(0)) \geq \frac{3\lambda_0}{4}$.*

Next, we introduce a lemma that shows for any t , if $\mathbf{w}_r(t)$ is close to $\mathbf{w}_r(0)$, then $\mathbf{H}(t)$ is identical to $\mathbf{H}(0)$, which ensures the positive definite of $\mathbf{H}(t)$.

Lemma 3 (Lemma 3.2 in [Du *et al.*, 2019c]). *Assume $\mathbf{w}_r(0) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ for $r \in [m]$ and $\|\mathbf{w}_r(0) - \mathbf{w}_r\|_2 \leq \frac{cp\lambda_0}{n^2} =: R$ for some small positive constant c , with probability at least $1 - p$, it has $\|\mathbf{H}(t) - \mathbf{H}(0)\|_2 < \frac{\lambda_0}{4}$, $\lambda_{\min}(\mathbf{H}(t)) > \frac{\lambda_0}{2}$ and $\lambda_{\max}(\mathbf{H}(t)) < \lambda_m := \lambda_{\max}(\mathbf{H}^\infty) + \frac{\lambda_0}{4}$.*

Therefore, under the over-parameterized assumption, $\mathbf{H}(t)$ stays positive with the condition number $\kappa := \lambda_m / (\lambda_0/2) > 1$ and $\mathbf{w}_r(t)$ remains close to its initialization during the training process.

Low-resolution view of momentum methods. From a continuous view, [Bu *et al.*, 2021] derived the convergence guarantees of HB and NAG in optimizing Eq.(2). They consider a non-linear dissipative dynamical system with $b > 0$

$$\ddot{\mathbf{w}}_r(t) + b\dot{\mathbf{w}}_r(t) + \frac{\partial \ell(\mathbf{W}(t), \mathbf{a})}{\partial \mathbf{w}_r(t)} = 0, \quad (11)$$

which corresponds to the limiting ODE for both HB and NAG when the learning rate tends to infinitesimal [Wilson *et al.*,

2021]. Consequently, Eq.(11) fails to capture the convergence difference between HB and NAG.

High-resolution view of momentum methods. In this subsection, we first present the high-resolution formulations of HB and NAG, then derive their corresponding residual dynamics when training the two-layer ReLU neural network.

On the one hand, according to [Shi *et al.*, 2021], HB has the following high-resolution ODE

$$\ddot{\mathbf{W}}(t) + 2\sqrt{\alpha}\dot{\mathbf{W}}(t) + (1 + \sqrt{\alpha s})\frac{\partial\ell(\mathbf{W}(t), \mathbf{a})}{\partial\mathbf{W}(t)} = 0, \quad (12)$$

where $\alpha \geq 0$ and $s > 0$. For each neuron $r \in [m]$, HB has the corresponding ODE

$$\ddot{w}_r(t) + 2\sqrt{\alpha}\dot{w}_r(t) + (1 + \sqrt{\alpha s})\frac{\partial\ell(\mathbf{W}(t), \mathbf{a})}{\partial w_r(t)} = 0. \quad (13)$$

On the other hand, NAG has a different ODE representation

$$\begin{aligned} \ddot{\mathbf{W}}(t) + 2\sqrt{\alpha}\dot{\mathbf{W}}(t) + \sqrt{s}\frac{\partial^2\ell(\mathbf{W}(t), \mathbf{a})}{\partial\mathbf{W}^2(t)}\dot{\mathbf{W}}(t) \\ + (1 + \sqrt{\alpha s})\frac{\partial\ell(\mathbf{W}(t), \mathbf{a})}{\partial\mathbf{W}(t)} = 0. \end{aligned} \quad (14)$$

The third term in Eq.(14) is the main difference between NAG and HB, which is induced by the gradient correction term in Eq.(5) [Shi *et al.*, 2021]. Accordingly, NAG has the following ODE for any w_r

$$\begin{aligned} \ddot{w}_r(t) + 2\sqrt{\alpha}\dot{w}_r(t) + \sqrt{s}\sum_{i \in [m]}\frac{\partial^2\ell(\mathbf{W}(t), \mathbf{a})}{\partial w_r(t)\partial w_i(t)}\dot{w}_i(t) \\ + (1 + \sqrt{\alpha s})\frac{\partial\ell(\mathbf{W}(t), \mathbf{a})}{\partial w_r(t)} = 0. \end{aligned} \quad (15)$$

When $s \rightarrow 0$, Eq.(13) and Eq.(15) degenerate to the same low-resolution Eq.(11). However, ℓ is non-convex while [Shi *et al.*, 2021] derived the convergence rate only in the convex setting. For the μ -strongly convex and smooth problem ℓ , Shi *et al.* sets $\alpha = \mu$ for both Eq.(13) and Eq.(15). However, their result only shows that HB and NAG converge at the same rate $\mathcal{O}(e^{-\sqrt{\mu}t/4})$, which fails to distinguish them in the continuous view.

Instead of the convex setting, we mainly focus on a non-convex ℓ . Following [Bu *et al.*, 2021], it has

$$\frac{\partial f_i}{\partial w_r} = \frac{1}{\sqrt{m}}a_r\mathbf{x}_i\mathbb{I}\{\langle \mathbf{w}_r, \mathbf{x}_i \rangle \geq 0\}, \quad (16)$$

$$\frac{\partial^2 f_i}{\partial w_r \partial w_l} = \mathbf{0}_{d \times d} \text{ for } r \neq l, \quad (17)$$

$$\frac{\partial^2 f_i}{\partial w_r^2} = \frac{1}{\sqrt{m}}a_r\mathbf{x}_i\mathbf{x}_i^\top \delta(\langle \mathbf{w}_r, \mathbf{x}_i \rangle) \stackrel{a.s.}{=} \mathbf{0}_{d \times d}, \quad (18)$$

where $\mathbf{0}_{d \times d}$ represents the zero matrix with dimension $d \times d$, $\frac{\partial\mathbb{I}\{x \geq 0\}}{\partial x} = \delta(x)$ and δ denotes the delta function that

$$\delta(x) = \begin{cases} +\infty, & x = 0 \\ 0, & x \neq 0. \end{cases} \quad (19)$$

Then, it has

$$\dot{f}_i = \sum_r \left\langle \frac{\partial f_i}{\partial w_r}, \dot{w}_r \right\rangle, \quad (20)$$

$$\ddot{f}_i = \sum_{r, l \in [m]} \dot{w}_r^\top \frac{\partial^2 f_i}{\partial w_r \partial w_l} \dot{w}_l + \sum_{r \in [m]} \left\langle \frac{\partial f_i}{\partial w_r}, \ddot{w}_r \right\rangle$$

$$\stackrel{a.s.}{=} \sum_{r \in [m]} \left\langle \frac{\partial f_i}{\partial w_r}, \ddot{w}_r \right\rangle, \quad (21)$$

and

$$\begin{aligned} \frac{\partial^2 \ell}{\partial w_r^2} &= \left(\frac{\partial \mathbf{f}}{\partial w_r} \right)^\top \frac{\partial \mathbf{f}}{\partial w_r} + \frac{\partial^2 \mathbf{f}}{\partial w_r^2} (\mathbf{f} - \mathbf{y}) \\ &\stackrel{a.s.}{=} \left(\frac{\partial \mathbf{f}}{\partial w_r} \right)^\top \frac{\partial \mathbf{f}}{\partial w_r}, \end{aligned} \quad (22)$$

$$\begin{aligned} \frac{\partial^2 \ell}{\partial w_r \partial w_l} &= \left(\frac{\partial \mathbf{f}}{\partial w_r} \right)^\top \frac{\partial \mathbf{f}}{\partial w_l} + \frac{\partial^2 \mathbf{f}}{\partial w_r \partial w_l} (\mathbf{f} - \mathbf{y}) \\ &= \left(\frac{\partial \mathbf{f}}{\partial w_r} \right)^\top \frac{\partial \mathbf{f}}{\partial w_l} \text{ for } r \neq l. \end{aligned} \quad (23)$$

When multiplying Eq.(13) with $\frac{\partial \mathbf{f}}{\partial w_r}$ and summing over all $r \in [m]$, it has the following ODE based on Eq.(9) and Eq.(20)

$$\ddot{\mathbf{f}}(t) + 2\sqrt{\alpha}\dot{\mathbf{f}}(t) + (1 + \sqrt{\alpha s})\mathbf{H}(t)(\mathbf{f}(t) - \mathbf{y}) = 0. \quad (24)$$

Then, the dynamics of the residual vector $\Delta = \mathbf{f} - \mathbf{y}$ has

$$\ddot{\Delta}(t) + 2\sqrt{\alpha}\dot{\Delta}(t) + (1 + \sqrt{\alpha s})\mathbf{H}(t)\Delta(t) = 0. \quad (25)$$

Instead of directly analyzing the non-convex ℓ over \mathbf{W} , we show that optimizing ℓ corresponds to the pseudo-loss $\hat{\ell}(t) := \frac{1}{2}\Delta^\top(t)\mathbf{H}(t)\Delta(t)$, which is $\frac{\lambda_0}{2}$ -strongly convex over $\Delta(t)$ according to Lemma 3. With $\mu = \lambda_0/2$, Eq.(13) can be transformed to

$$\ddot{\Delta}(t) + \sqrt{2\lambda_0}\dot{\Delta}(t) + (1 + \sqrt{\frac{\lambda_0 s}{2}})\frac{\partial \hat{\ell}(t)}{\partial \Delta(t)} = 0, \quad (26)$$

where $\alpha = \lambda_0/2$ according to [Shi *et al.*, 2021]. Similarly, it can derive following residual dynamics of NAG as

$$\ddot{\Delta}(t) + \sqrt{2\lambda_0}\dot{\Delta}(t) + \sqrt{s}\mathbf{H}(t)\dot{\Delta}(t) + (1 + \sqrt{\frac{\lambda_0 s}{2}})\frac{\partial \hat{\ell}(t)}{\partial \Delta(t)} = 0, \quad (27)$$

where the third term in Eq.(27) is derived by

$$\begin{aligned} &\sum_{r \in [m]} \left\langle \frac{\partial \mathbf{f}}{\partial w_r}, \sum_{i \in [m]} \frac{\partial^2 \ell(\mathbf{W}(t), \mathbf{a})}{\partial w_r(t) \partial w_i(t)} \dot{w}_i(t) \right\rangle \\ &= \sum_{r \in [m]} \left\langle \frac{\partial \mathbf{f}}{\partial w_r}, \sum_{i \in [m]} \left(\frac{\partial \mathbf{f}}{\partial w_r} \right)^\top \frac{\partial \mathbf{f}}{\partial w_i} \dot{w}_i(t) \right\rangle \\ &= \left\langle \sum_{r \in [m]} \frac{\partial \mathbf{f}}{\partial w_r} \left(\frac{\partial \mathbf{f}}{\partial w_r} \right)^\top, \sum_{i \in [m]} \frac{\partial \mathbf{f}}{\partial w_i} \dot{w}_i(t) \right\rangle = \mathbf{H}(t)\dot{\mathbf{f}}(t). \end{aligned}$$

Note that the third term in Eq.(27) is derived from the gradient correction term.

4 Main Results

In this section, we provide the convergence analysis and comparison of HB and NAG.

4.1 Convergence Analysis of HB

Theorem 1. Suppose $m = \Omega(\frac{n^6}{p^3\lambda_0^4})$, $w_r(0) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $a_r \sim \text{unif}\{-1, 1\}$ for any $r \in [m]$, with probability at least $1 - p$, it has

$$\ell(t) \leq \frac{6\hat{\ell}(0)}{\lambda_0} e^{-(2-\sqrt{2})\sqrt{\frac{\lambda_0}{2}}t}. \quad (28)$$

Remark 1. The above theorem implies HB can attain a global minimum of Eq.(2) at a linear rate $\mathcal{O}(e^{-(2-\sqrt{2})\sqrt{\frac{\lambda_0}{2}}t})$ in training two-layer ReLU neural networks. It is worth noting that [Shi *et al.*, 2021] studied the same high-resolution ODE of HB and proved HB can attain a global convergence rate $\mathcal{O}(e^{-\frac{1}{4}\sqrt{\mu}t})$ for μ -strongly convex problems. In contrast, our focus is on a non-convex optimization problem. We show that optimizing the non-convex loss ℓ corresponds to an error dynamics, as if we were actually optimizing a strongly convex loss $\hat{\ell}$ over the prediction error. Compared to [Shi *et al.*, 2021], Theorem 1 further provides a tighter upper bound for the convergence rate of HB in the high-resolution continuous-time view.

Remark 2. We are aware of some recent works that improve the requirement of the over-parameterization [Song and Yang, 2019; Nguyen, 2021]. In this work, we mainly focus on providing the theoretical guarantee for the acceleration of NAG over HB in training over-parameterized neural networks. Therefore, we incorporate the first and also the most well-known result about the requirement on m that ensures $\mathbf{H}(t)$ positive [Du *et al.*, 2019c]. We believe incorporating these improved results on m into our analysis can also improve the bound on m in our work.

The proof of Theorem 1 mainly consists of three parts:

1. Firstly, based on the positive definite assumption of \mathbf{H} , we demonstrate the training loss decays to zero at a linear rate and the distance between \mathbf{w} and its initial value has an upper bound R' .
2. Secondly, we show that if R' is smaller than the distance R defined in Lemma 3, our conclusions from the first procedures hold true and the positive definite of \mathbf{H} is guaranteed.
3. Finally, the width m is determined by satisfying $R' < R$, which completes the proof of Theorem 1.

To start with, we present the first procedure of the proof, which involves the smallest eigenvalue $\frac{\lambda_0}{2}$ of \mathbf{H} as introduced in Lemma 3.

Lemma 4. Assume $\lambda_{\min}(\mathbf{H}(i)) \geq \frac{\lambda_0}{2}$ for $0 \leq i \leq t$. With $0 < s \leq 1/\lambda_m$ and $\dot{\mathbf{w}}_r(0) = 0$ for any $r \in [m]$, HB has $\ell(t) \leq \frac{6\hat{\ell}(0)}{\lambda_0} e^{-(2-\sqrt{2})\sqrt{\frac{\lambda_0}{2}}t}$ and $\|\mathbf{w}_r(t) - \mathbf{w}_r(0)\|_2 \leq R' := 10\sqrt{6\hat{\ell}(0)n/(\lambda_0^3 m)}$.

The proof is provided in the supplementary materials. Then, the conclusions in Lemma 4 hold when the distance R' defined in Lemma 4 is less than the distance R given in Lemma 3.

Lemma 5. Assume $R' < R$, it has (1) $\lambda_{\min}(\mathbf{H}(t)) \geq \lambda_0/2$, (2) $\|\mathbf{w}_r(t) - \mathbf{w}_r(0)\|_2 \leq R'$ for all $r \in [m]$ and (3) $\ell(t) \leq \frac{6\hat{\ell}(0)}{\lambda_0} e^{-(2-\sqrt{2})\sqrt{\frac{\lambda_0}{2}}t}$.

The proof can be founded in the supplementary materials. To satisfy the assumption $R' < R$ in Lemma 5, it requires $m = \Omega(\frac{n^6}{p^3 \lambda_0^4})$ based on $\hat{\ell}(0) = \Omega(n\lambda_0/p)$ as proved in [Bu *et al.*, 2021], which completes the proof of Theorem 1.




Figure 1: The value of $\rho_{NAG}^*(\alpha)$ with respect to α .

4.2 Convergence Analysis of NAG

In this subsection, based on the similar analysis framework as the proof of HB, we present the convergence guarantee of NAG and show its acceleration.

Theorem 2. Suppose $m = \Omega(\frac{n^6}{p^3 \lambda_0^4})$, $\mathbf{w}_r(0) \sim \mathcal{N}(0, 1)$ and $a_r \sim \text{unif}\{-1, 1\}$ for any $r \in [m]$, with probability at least $1 - p$, it has

$$\ell(t) \leq \frac{26\hat{\ell}(0)}{3\lambda_0} e^{-\rho_{NAG}^*(\alpha)\sqrt{\frac{\lambda_0}{2}}t}, \quad (29)$$

where

$$\alpha = \sqrt{2\lambda_0 s}/4, \quad \rho_{NAG}^*(\alpha) = \frac{1}{2}(4 + 3\alpha - \sqrt{8 + 16\alpha + 9\alpha^2}).$$

Remark 3. The above theorem indicates that NAG can reach a global minimum at a linear rate $\mathcal{O}(e^{-\rho_{NAG}^*(\alpha)\sqrt{\frac{\lambda_0}{2}}t})$. When $0 < s \leq 1/\lambda_m$, it holds that $0 < \alpha \leq 1/(2\sqrt{\kappa}) \leq 0.5$. From Fig. 1, it can observe that $\rho_{NAG}^*(\alpha)$ is a monotonic increasing function with respect to α when $0 < \alpha \leq 0.5$. Thus, it has

$$2 - \sqrt{2} < \rho_{NAG}^*(\alpha) \leq (11 - \sqrt{73})/4. \quad (30)$$

Compared to the convergence rate $\mathcal{O}(e^{-(2-\sqrt{2})\sqrt{\frac{\lambda_0}{2}}t})$ of HB, our theoretical result demonstrates that NAG converges faster than HB.

In addition, according to Eq.(44) in the supplementary materials, the derivative of the corresponding Lyapunov function V of NAG has

$$\begin{aligned} \dot{V} = & -(1 + \sqrt{\frac{\lambda_0 s}{2}})\sqrt{2\lambda_0}\hat{\ell}(t) - \frac{\sqrt{s}}{2}(1 + \sqrt{\frac{\lambda_0 s}{2}})\|\mathbf{H}(t)\Delta(t)\|^2 \\ & - \frac{\sqrt{2\lambda_0}}{2}\|\dot{\Delta}(t)\|^2 - \frac{\sqrt{s}}{2}\dot{\Delta}^\top(t)\mathbf{H}(t)\dot{\Delta}(t), \end{aligned} \quad (31)$$

where the third term on the right hand side of Eq.(31) is induced from $\sqrt{s}\mathbf{H}(t)\dot{\Delta}(t)$, which corresponds to the gradient correction term as mentioned in Section 3.2. If neglecting this term, based on the upper bound of $V(t)$ (Eq.(43) in the supplementary materials), the corresponding convergence rate has

$$\rho^* = \max_{\phi > 0} \min\left\{\frac{2}{3+2\phi}, \frac{4}{2+1/\phi}, \frac{4}{1+\phi}\right\}\sqrt{\frac{\lambda_0}{2}} = 2 - \sqrt{2},$$

which is the same as the convergence rate of HB. Therefore, we provably show that the extra gradient correction term of NAG in Eq.(27) leads to the acceleration of NAG over HB in training neural networks.

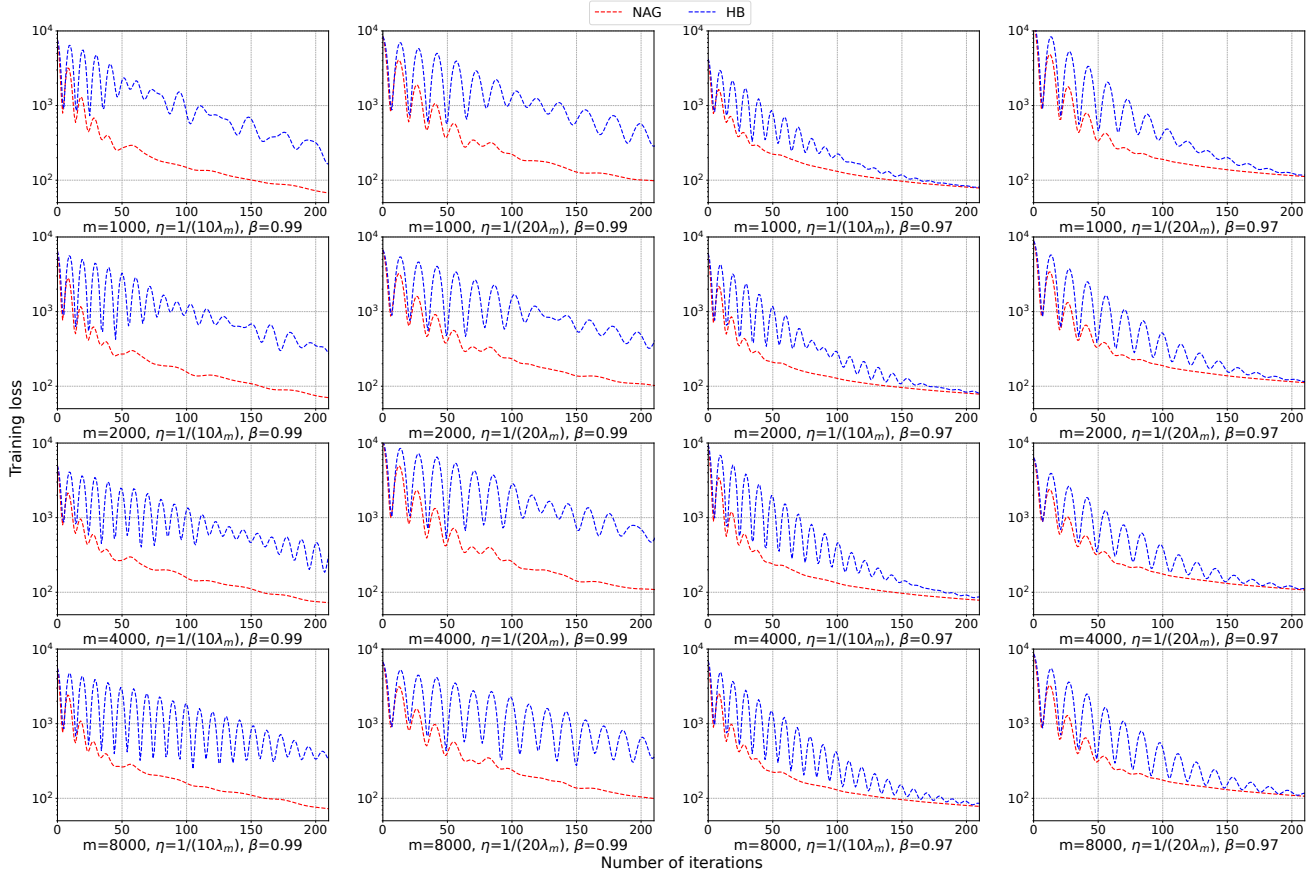


Figure 2: Convergence comparison between NAG and HB in training two-layer ReLU neural networks on the MNIST dataset with different width m , learning rate η and momentum parameter β .

Remark 4. It is noted that [Shi *et al.*, 2021] conducted an analysis on the same high-resolution ODE of NAG for μ -strongly convex problems. Their result demonstrated that the convergence rate of NAG is upper bounded by $\mathcal{O}(e^{-\frac{1}{4}\sqrt{\mu}t})$. Later, [Sun *et al.*, 2020] and [Zhang *et al.*, 2021] improved the convergence rate to $\mathcal{O}(e^{-\frac{3}{2}\sqrt{\mu}t})$ and $\mathcal{O}(e^{-\frac{1}{2}\sqrt{\mu}t})$, respectively. Comparing to [Shi *et al.*, 2021; Sun *et al.*, 2020; Zhang *et al.*, 2021], our study focuses on a non-convex ℓ over the weights. Furthermore, we demonstrate that ℓ corresponds to a $\frac{\lambda_0}{2}$ -strongly convex problem by utilizing the NTK theory. Additionally, we provide a tight convergence rate of NAG with a more refined inequality and analysis.

Remark 5. Noted that directly studying the convergence of discrete algorithms for non-convex ℓ is far too complex and general to offer insightful explanation about their efficiency in deep learning. Instead of the low-resolution approach in previous work [Bu *et al.*, 2021], we study from a high-resolution continuous viewpoint. Our result provides the first acceleration guarantee of NAG over HB with a simple and intuitive proof. Furthermore, Proposition 1 in [Shi *et al.*, 2021] shows that the discrete iterates and the high-resolution ODE trajectories remain close for both HB and NAG when the learning rate is small enough. It would be interesting to consider the discrete setting in the future work.

Lemma 6. Assume $\lambda_{\min}(\mathbf{H}(i)) \geq \frac{\lambda_0}{2}$ for $0 \leq i \leq t$. With $0 < s \leq 1/\lambda_m$ and $\dot{\mathbf{w}}_r(0) = 0$ for any $r \in [m]$, NAG has $\ell(t) \leq \frac{26\hat{\ell}(0)}{3\lambda_0} e^{-\rho_{NAG}^* \sqrt{\lambda_0}/2t}$ and $\|\mathbf{w}_r(t) - \mathbf{w}_r(0)\| \leq 67\sqrt{\hat{\ell}(0)n/(\lambda_0^3 m)}$.

The proof is presented in the supplementary materials. Similarly, it can establish the conclusion of the Theorem 2 following Lemma 5. Moreover, the required width should meet $m = \Omega(\max\{\sqrt{sn}/\sqrt{\lambda_0}, \frac{n^6}{p^3\lambda_0^4}\}) = \Omega(\frac{n^6}{p^3\lambda_0^4})$ using $\lambda_0 \leq 1/2$ as proved in [Bu *et al.*, 2021].

5 Numerical Experiments

5.1 Setup

We use three commonly used datasets: MNIST [LeCun *et al.*, 1998], FMNIST [Xiao *et al.*, 2017] and CIFAR10 [Krizhevsky *et al.*, 2009], where the pre-processing procedure follows the existing work [Arora *et al.*, 2019b]. Specifically, we use the first two classes of instances, assigning the label of the first class to +1 and the second to -1. In addition, each instance is normalized with the unit norm.

Regarding the architecture and initialization of the neural network, we adopt the setting in Section 3.1. Besides, we set the width $m \in \{1000, 2000, 4000, 8000\}$, learning

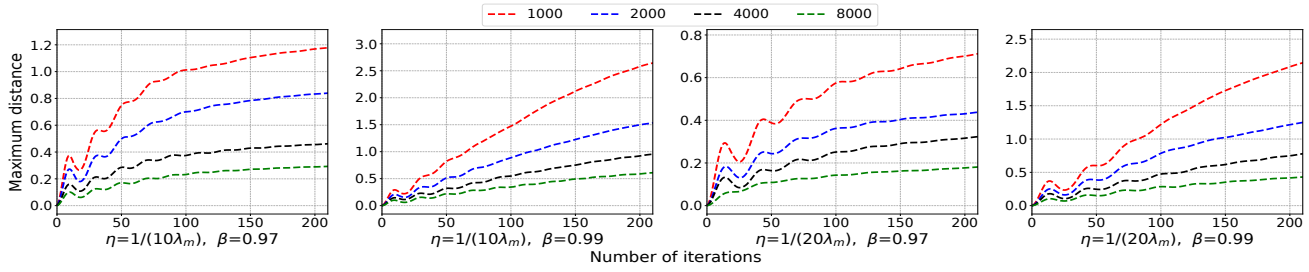


Figure 3: The maximum distance $\max_{r \in [m]} \|\mathbf{w}_r(t) - \mathbf{w}_0(t)\|$ of HB in training two-layer ReLU neural networks on the MNIST dataset during training process.

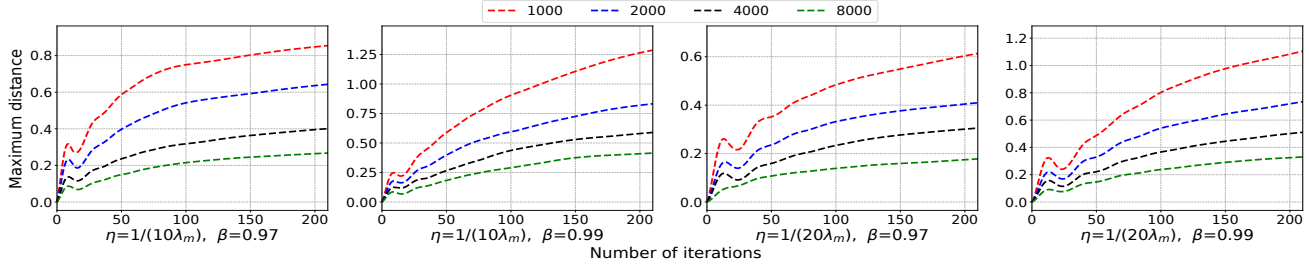


Figure 4: The maximum distance $\max_{r \in [m]} \|\mathbf{w}_r(t) - \mathbf{w}_0(t)\|$ of NAG in training two-layer ReLU neural networks on the MNIST dataset during training process.

rate $\eta \in \{1/(10\lambda_m), 1/(20\lambda_m)\}$ and momentum parameter $\beta \in \{0.97, 0.99\}$, where λ_m is calculated with the eigenvalues of the NTK matrix in Eq.(10). For each setting of hyper-parameters, we conduct 5 independent runs with different random seeds for initialization. The average training loss over iterations is presented by the dashed line. We use JAX [Bradbury *et al.*, 2018] library for code implementation and all experiments are conducted on 4 NVIDIA V100 GPU.

5.2 Results Analysis

Due to the space limitation, only the results on the MNIST dataset are presented here. The remaining results can be founded in the supplementary materials. From Fig. 2, it can observe that the training losses of NAG and HB decay at a almost linear rate, which demonstrates the linear convergence of HB and NAG in training the two-layer ReLU neural networks. Additionally, the loss curve of NAG lies below that of HB for different hyper-parameter settings, which indicates NAG achieves a smaller training loss compared to HB at the same iteration. Thus, NAG achieve an acceleration over HB. These empirical findings validate the conclusion in Section 4 that NAG theoretically obtains a faster global convergence rate compared to HB.

To further validate our findings, we present the maximum distance $\max_{r \in [m]} \|\mathbf{w}_r(t) - \mathbf{w}_0(t)\|$ of HB and NAG during the training process with different widths. According to Lemma 4 and Lemma 6, the theoretical maximum distances of HB and NAG both satisfy $\max_{r \in [m]} \|\mathbf{w}_r(t) - \mathbf{w}_0(t)\| \sim \mathcal{O}(1/\sqrt{m})$, which indicate that the maximum distance decreases when the width of the neural network increases. The corresponding empirical results are shown in Fig. 3 and Fig. 4, which is in accordance with our theoretical findings.

6 Conclusion and Future Work

In this paper, we study the convergence of momentum methods in training an over-parameterized two-layer ReLU neural network. By exploiting high-resolution ODE formulations and the NTK theory, the global convergences of HB and NAG are established. In contrast with previous works, we exploit high-resolution ODEs instead of the low-resolution. In addition, our work yields tighter upper bounds via a more refined inequality and analysis, which provides the first theoretical guarantee for the acceleration of NAG over HB in tackling non-convex optimization problem on neural networks.

There are several directions for future research. Firstly, it is natural to further extend our result to deep neural networks, where the key technical challenge lies in deriving and analyzing the spectral properties of the corresponding \mathbf{H} . We note that there is some works about the eigenvalues of \mathbf{H} , including deep fully-connected neural networks [Du *et al.*, 2019b], graph neural networks [Du *et al.*, 2019a] and convolutional neural networks [Li *et al.*, 2019]. Secondly, stochastic optimization methods are widely used in training neural networks due to relatively low computational costs. Therefore, it would be interesting to extend existing stochastic convex works [Nakakita, 2022; Laborde and Oberman, 2020] to analyze the stochastic optimization problem in neural networks. Finally, the understanding of the adaptive first-order method is still lacking, such as AdaBound [Liu *et al.*, 2022a] and RMSProp [Xu *et al.*, 2021]. These adaptive methods also incorporate the idea of momentum and our work may inspire the development of new tools for understanding their convergence behaviors.

Acknowledgements

This work was supported by National Natural Science Foundation of China (No.62076251 and No.62106281).

References

- [Alemohammad *et al.*, 2020] Sina Alemohammad, Zichao Wang, Randall Balestriero, and Richard Baraniuk. The recurrent neural tangent kernel. In *International Conference on Learning Representations*, 2020.
- [Arora *et al.*, 2019a] Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, pages 8141–8150, 2019.
- [Arora *et al.*, 2019b] Sanjeev Arora, Simon S. Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332, 2019.
- [Bradbury *et al.*, 2018] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: Composable transformations of Python+NumPy programs, 2018.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, pages 1877–1901, 2020.
- [Bu *et al.*, 2021] Zhiqi Bu, Shiyun Xu, and Kan Chen. A dynamical view on optimization algorithms of overparameterized neural networks. In *International Conference on Artificial Intelligence and Statistics*, pages 3187–3195, 2021.
- [Du *et al.*, 2019a] Simon S. Du, Kangcheng Hou, Ruslan Salakhutdinov, Barnabás Póczos, Ruosong Wang, and Keyulu Xu. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. In *Advances in Neural Information Processing Systems*, pages 5724–5734, 2019.
- [Du *et al.*, 2019b] Simon S Du, Jason D Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. *International Conference on Machine Learning*, pages 1675–1685, 2019.
- [Du *et al.*, 2019c] Simon S. Du, Xiyu Zhai, Barnabás Póczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019.
- [Ghadimi *et al.*, 2015] Euhanna Ghadimi, Hamid Reza Feyzmahdavian, and Mikael Johansson. Global convergence of the heavy-ball method for convex optimization. In *European Control Conference*, pages 310–315, 2015.
- [Huang *et al.*, 2020] Kaixuan Huang, Yuqing Wang, Molei Tao, and Tuo Zhao. Why do deep residual networks generalize better than deep feedforward networks - A neural tangent kernel perspective. In *Advances in Neural Information Processing Systems*, pages 2698–2709, 2020.
- [Jacot *et al.*, 2018] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing System*, pages 8571–8580, 2018.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [Laborde and Oberman, 2020] Maxime Laborde and Adam M. Oberman. A Lyapunov analysis for accelerated gradient methods: from deterministic to stochastic case. In *International Conference on Artificial Intelligence and Statistics*, pages 602–612, 2020.
- [LeCun *et al.*, 1998] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [Lessard *et al.*, 2016] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.
- [Li *et al.*, 2019] Zhiyuan Li, Ruosong Wang, Dingli Yu, Simon S Du, Wei Hu, Ruslan Salakhutdinov, and Sanjeev Arora. Enhanced convolutional neural tangent kernels. *arXiv preprint arXiv:1911.00809*, 2019.
- [Liu *et al.*, 2022a] Jinlan Liu, Jun Kong, Dongpo Xu, Miao Qi, and Yinghua Lu. Convergence analysis of adabound with relaxed bound functions for non-convex optimization. *Neural Networks*, 145:300–307, 2022.
- [Liu *et al.*, 2022b] Xin Liu, Zhisong Pan, and Wei Tao. Provable convergence of nesterov’s accelerated gradient method for over-parameterized neural networks. *Knowledge-Based Systems*, 251:109277, 2022.
- [Nado *et al.*, 2021] Zachary Nado, Justin M Gilmer, Christopher J Shallue, Rohan Anil, and George E Dahl. A large batch optimizer reality check: Traditional, generic optimizers suffice across batch sizes. *arXiv preprint arXiv:2102.06356*, 2021.
- [Nakakita, 2022] Shogo Nakakita. Parametric estimation of stochastic differential equations via online gradient descent. *arXiv preprint arXiv:2210.08800*, 2022.
- [Nesterov, 1983] Yurii Nesterov. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. In *Sov. Math. Dokl*, 1983.
- [Nesterov, 2013] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*. Springer Science, 2013.
- [Nguyen, 2021] Quynh Nguyen. On the proof of global convergence of gradient descent for deep relu networks with

- linear widths. In *International Conference on Machine Learning*, pages 8056–8062, 2021.
- [Polyak, 1964] Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- [Schmidhuber, 2015] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.
- [Schmidt *et al.*, 2021] Robin M. Schmidt, Frank Schneider, and Philipp Hennig. Descending through a crowded valley - benchmarking deep learning optimizers. In *International Conference on Machine Learning*, pages 9367–9376, 2021.
- [Shi *et al.*, 2021] Bin Shi, Simon S Du, Michael I Jordan, and Weijie J Su. Understanding the acceleration phenomenon via high-resolution differential equations. *Mathematical Programming*, pages 1–70, 2021.
- [Song and Yang, 2019] Zhao Song and Xin Yang. Quadratic suffices for over-parametrization via matrix chernoff bound. *arXiv preprint arXiv:1906.03593*, 2019.
- [Su *et al.*, 2014] Weijie Su, Stephen Boyd, and Emmanuel Candes. A differential equation for modeling nesterov’s accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems*, pages 2510–2518, 2014.
- [Sun *et al.*, 2020] Boya Sun, Jemin George, and Solmaz S. Kia. High-resolution modeling of the fastest first-order optimization method for strongly convex functions. In *Conference on Decision and Control*, pages 4237–4242, 2020.
- [Wang *et al.*, 2021] Jun-Kun Wang, Chi-Heng Lin, and Jacob D. Abernethy. A modular analysis of provable acceleration via polyak’s momentum: Training a wide relu network and a deep linear network. In *International Conference on Machine Learning*, pages 10816–10827, 2021.
- [Wilson *et al.*, 2021] Ashia C Wilson, Ben Recht, and Michael I Jordan. A lyapunov analysis of accelerated methods in optimization. *Journal of Machine Learning Research*, 22:1–34, 2021.
- [Xiao *et al.*, 2017] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [Xu *et al.*, 2021] Dongpo Xu, Shengdong Zhang, Huisheng Zhang, and Danilo P. Mandic. Convergence of the rmsp-prop deep learning method with penalty for nonconvex optimization. *Neural Networks*, 139:17–23, 2021.
- [Zhang *et al.*, 2021] Peiyuan Zhang, Antonio Orvieto, Hadi Daneshmand, Thomas Hofmann, and Roy S. Smith. Revisiting the role of euler numerical integration on acceleration and stability in convex optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 3979–3987, 2021.