

Effects of Momentum in Implicit Bias of Gradient Flow for Diagonal Linear Networks

Bochen Lyu^{1, 2}, He Wang³, Zheng Wang⁴, Zhanxing Zhu^{*2}

¹DataCanvas Lab, DataCanvas, Beijing, China

²University of Southampton, UK

³UCL Centre for Artificial Intelligence, Department of Computer Science, UK

⁴University of Leeds, UK

bochen.lv@gmail.com, he_wang@ucl.ac.uk, z.wang5@leeds.ac.uk, z.zhu@soton.ac.uk

Abstract

This paper targets on the regularization effect of momentum-based methods in regression settings and analyzes the popular diagonal linear networks to precisely characterize the implicit bias of continuous versions of heavy-ball (HB) and Nesterov’s method of accelerated gradients (NAG). We show that, HB and NAG exhibit different implicit bias compared to GD for diagonal linear networks, which is different from the one for classic linear regression problem where momentum-based methods share the same implicit bias with GD. Specifically, the role of momentum in the implicit bias of GD is twofold: (a) HB and NAG induce extra initialization mitigation effects similar to SGD that are beneficial for generalization of sparse regression; (b) the implicit regularization effects of HB and NAG also depend on the initialization of gradients explicitly, which may not be benign for generalization. As a result, whether HB and NAG have better generalization properties than GD jointly depends on the aforementioned twofold effects determined by various parameters such as learning rate, momentum factor, and integral of gradients. Our findings highlight the potential beneficial role of momentum and can help understand its advantages in practice such as when it will lead to better generalization performance.

1 Introduction

Extensive deep learning tasks aim to solve the optimization problem

$$\arg \min_{\beta} L(\beta) \quad (1)$$

where L is the loss function and β is the parameter. Gradient descent (GD) and its variants underpin such optimization of parameters for deep learning, thus understanding these simple yet highly effective algorithms is crucial to unveil the thrilling generalization performance of deep neural networks (DNNs). Recently, Soudry et al. (2018); Ji and Telgarsky (2019); Lyu and Li (2020); Pesme, Pillaud-Vivien, and Flammarion (2021); Azulay et al. (2021); Nacson et al. (2019) have made significant efforts in this direction to understand GD-based methods through the lens of *implicit bias*, which states that *GD and its variants are implicitly biased towards selecting particular solutions among all global minimum*.

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In particular, Soudry et al. (2018) pioneered the study of implicit bias of GD and showed that GD selects the max-margin solution for logistic regression on separable dataset. For regression problems, the simplest setting is the linear regression problem, where GD and its stochastic variant, SGD, are biased towards the interpolation solution that is closest to the initialization measured by the Euclidean distance (Ali, Dobriban, and Tibshirani 2020). In order to investigate the implicit bias for DNNs, diagonal linear network, a simplified version of deep learning models, has been proposed. For this model, Woodworth et al. (2020); Azulay et al. (2021); Yun, Krishnan, and Mobahi (2021) showed that the solution selected by GD is equivalent to that of a constrained norm minimization problem interpolating between ℓ_1 and ℓ_2 norms up to the magnitude of the initialization scale. Pesme, Pillaud-Vivien, and Flammarion (2021) further characterized that adding stochastic noise to GD additionally induces a regularization effect equivalent to reducing the initialization magnitude.

Besides these fruitful progresses, Gunasekar et al. (2018); Wang et al. (2022) studied the implicit bias of momentum-based methods for one-layer linear models and showed that they have the same implicit bias as GD. Jelassi and Li (2022), on the other hand, revealed that momentum-based methods have better generalization performance than GD for a special linear CNN in classification problems. Ghosh et al. (2023) conducted a model-agnostic analysis of $\mathcal{O}(\eta^2)$ approximate continuous version of HB from the perspective of IGR (Barrett and Dherin 2021). Momentum methods adopt a two-step manner and can induce different dynamics compared to vanilla GD: from the perspective of their $\mathcal{O}(\eta)$ -continuous approximation modelling, the approximation for GD is a first-order ODE (gradient flow, GF): $d\beta/dt = -\nabla L(\beta)$, while, as a comparison, the approximation for momentum-based methods can be seen as a damped second-order Hamiltonian dynamic with potential $L(\beta)$:

$$m \frac{d^2 \beta}{dt^2} + \lambda \frac{d\beta}{dt} + \nabla L(\beta) = 0,$$

which was first observed by Polyak (1964). Due to the clear difference in their dynamics, it is natural and intriguing to ask from the theoretical point of view:

(Q): Will adding momentum to GD change its implicit bias for DNNs?

For the least squares problem (single layer linear network), Gunasekar et al. (2018) argued that momentum does not change the implicit bias of GD, while the case for deep learning models is more complex. From the empirical point of view, momentum typically leads to a better generalization performance and is crucial in the training of modern DNNs. This suggests that they might enjoy a different implicit bias compared to GD. Therefore, its theoretical characterization is meaningful and necessary.

Hence, our goal in this work is to precisely characterize the implicit bias of momentum-based methods to take a step towards answering the above fundamental question. To explore the case for deep neural networks, we consider the popular deep linear models: *diagonal linear networks*. Although the structures of diagonal linear networks are simple, they already capture many insightful properties of DNNs, including the dependence on the initialization, the over-parameterization of the parameters, the non-convexity, and the transition from lazy regime to rich regime (Woodworth et al. 2020; Pesme, Pillaud-Vivien, and Flammarion 2021; Azulay et al. 2021) which is an intriguing phenomenon observed in many complex neural networks.

Heavy-Ball algorithms (Polyak 1964) (HB) and Nesterov’s method of accelerated gradients (Nesterov 1983) (NAG) are the most widely adopted momentum-based methods. These algorithms are generally implemented with a fixed momentum factor in deep learning libraries such as PyTorch (Paszke et al. 2017). To be consistent with such practice, we focus on HB and NAG with a fixed momentum factor that is independent of learning rate or iteration count. For the purpose of conducting a tractable theoretical analysis, we rely on the tools of continuous time approximations of momentum-based methods (with fixed momentum factor), HB and NAG flow, which were recently interpreted by Kovachki and Stuart (2021) as *modified equations* in the numerical analysis literature and by Shi et al. (2018) as high resolution ODE approximation.

Our findings are summarized as follows. We show that, *unlike* the case for single layer linear networks where momentum-based methods HB and NAG share similar implicit bias with GF, they exhibit different implicit bias for diagonal linear networks compared to GF in two main aspects:

1. Compared to GF, although HB and NAG flow also converge to solutions that minimize a norm interpolating between ℓ_2 -norm and ℓ_1 -norm up to initialization scales of parameters, they induce an *extra effect that is equivalent to mitigating the influence of initialization of the model parameters*, which is beneficial for generalization properties of sparse regression (Theorem 3). Note that stochastic gradient flow (SGF) could also yield an initialization mitigation effects (Pesme, Pillaud-Vivien, and Flammarion 2021), although momentum-based methods and SGF modify GF differently.
2. The solutions of HB and NAG flow also *depend on the initialization of both parameters and gradients explicitly and simultaneously*, which may not be benign for the generalization performances for sparse regression.

Therefore, HB and NAG are not always better than GD from the perspective of generalization for sparse regression. Whether HB and NAG have the advantages of generalization over GD is up to the overall effects of the above two distinct effects determined by various hyper-parameters. In particular, when mitigation effects of initialization of parameters brought by HB and NAG outperform their dependence on the initialization of gradients, HB and NAG will have better generalization performances than GD, e.g., when the initialization is highly biased (Fig. 2(a)), otherwise, they will not show such advantages over GD (Fig. 1(a)).

Organization. This paper is organized as follows. In Section 2, we summarize notations, setup, and continuous time approximation modelling details for HB and NAG. Section 3 concentrates on our main results of the implicit bias of momentum-based methods for diagonal linear networks with corresponding numerical experiments to support our theoretical findings. We conclude this work in Section 4. All the proof details and additional experiments are presented in Appendix.

1.1 Related works

To characterize the properties of momentum-based methods, continuous time approximations of them are introduced in several recent works. Su, Boyd, and Candes (2014) provided a second-order ODE to precisely describe the NAG with momentum factor depending on the iteration count. Wilson, Recht, and Jordan (2016) derived a limiting equation for both HB and NAG when the momentum factor depends on learning rate or iteration count. Shi et al. (2018) further developed high-resolution limiting equations for HB and NAG, and Wibisono, Wilson, and Jordan (2016) designed a general framework from the perspective of Bregman Lagrangian. When the momentum is fixed and does not depend on learning rate or iteration count, Kovachki and Stuart (2021) developed the continuous time approximation, the modified equation in the numerical analysis literature, for HB and NAG.

Compared to these works, we develop the continuous time approximations of HB and NAG for deep learning models and regression problems, and we focus on the implicit bias of HB and NAG flow rather than GF. Furthermore, we also take into account the effects of other sources of implicit bias such as initialization and model architecture, which is different from the model-agnostic analysis in Ghosh et al. (2023).

The recent work Papazov, Pesme, and Flammarion (2024) also studied HB for diagonal linear networks using continuous approximation where the initialization of speed of parameters are all zero. In particular, they characterized a crucial quantity $\eta(1 - \mu)^{-2}$ that can induce acceleration of the optimization of HB and can make the solution of HB generalize better when $\eta(1 - \mu)^{-2}$ is sufficiently small. As a comparison, we do not impose restriction to initialization and characterize the implicit bias for both HB and NAG flow. In addition, we reveal the role of initialization of gradients and show when HB and NAG flow can generalize better/worse than GF.

2 Preliminaries

Notations. We let $\{1, \dots, L\}$ be all integers between 1 and L . The dataset with n samples is denoted by $\{(x_i, y_i)\}_{i=1}^n$, where $x_i \in \mathbb{R}^d$ is the d -dimensional input and $y_i \in \mathbb{R}$ is the scalar output. The data matrix is represented by $X \in \mathbb{R}^{n \times d}$ where each row is a feature x_i and $y = (y_1, \dots, y_n)^T \in \mathbb{R}^n$ is the collection of y_i . For a vector $a \in \mathbb{R}^d$, a_j denotes its j -th component and its ℓ_p -norm is $\|a\|_p$. For a vector $a(t)$ depending on time, we use $\dot{a}(t) = da/dt$ to denote the first time derivative and $\ddot{a}(t) = d^2a/dt^2$ for the second time derivative. The element-wise product is denoted by \odot such that $(a \odot b)_j = a_j b_j$. We let $\mathbf{e}_d = (1, \dots, 1)^T \in \mathbb{R}^d$. For a square matrix $W \in \mathbb{R}^{d \times d}$, we use $\text{diag}(W)$ to denote the corresponding vector $(W_{11}, \dots, W_{dd})^T \in \mathbb{R}^d$.

Heavy-Ball and Nesterov’s method of accelerated gradients. Heavy-Ball (HB) and Nesterov’s method of accelerated gradients (NAG) are perhaps the most widely adopted momentum-based methods. Different from GD, HB and NAG apply a two-step scheme (Sutskever et al. 2013). In particular, for Eq. (1) let k be the iteration number, μ be the momentum factor, η be the learning rate, and $p \in \mathbb{R}^d$ be the momentum of parameter β , then HB updates β as follows:

$$p_{k+1} = \mu p_k - \eta \nabla L(\beta_k), \quad \beta_{k+1} = \beta_k + p_{k+1} \quad (2)$$

where $p_0 = 0$. Similarly, NAG can also be written as a two-step manner

$$p_{k+1} = \mu p_k - \eta \nabla L(\beta_k + \mu p_k), \quad \beta_{k+1} = \beta_k + p_{k+1} \quad (3)$$

with $p_0 = 0$. Note that although previous works (Su, Boyd, and Candes 2014; Nesterov 2014; Wilson, Recht, and Jordan 2016; Shi et al. 2018) considered HB and NAG with momentum factor depending on the learning rate η or iteration count k , HB and NAG are generally implemented with constant momentum factor such as in PyTorch (Paszke et al. 2017). Therefore a constant momentum factor μ is assumed in this work as in Kovachki and Stuart (2021) to be consistent with such practice.

HB and NAG flow: continuous time approximations of HB and NAG. In this work, we analyze the implicit bias of HB and NAG through their continuous time approximations summarized as follows, which provide insights to the corresponding discrete algorithms and enable us to take the great advantages of the convenience of theoretical analysis at the same time. We start with the definition of order of convergence of continuous approximation for discrete HB and NAG.

Definition 1 (Order of convergence of HB and NAG flow). *An ODE whose solution is $\beta(t)$ is the order $\mathcal{O}(\eta^\gamma)$ continuous approximate version of the discrete HB Eq. (2) and NAG Eq. (3) if for $k = 0, 1, 2, \dots$, let $\bar{\beta}_k$ be the sequence given by Eq. (2) or Eq. (3) and let $\beta_k = \beta(t = k\eta)$, then for any $T \geq 0$, there exists a constant $C > 0$ such that $\sup_{0 \leq k\eta \leq T} |\beta_k - \bar{\beta}_k| \leq C\eta^\gamma$.*

Proposition 1 (HB and NAG flow: $\mathcal{O}(\eta)$ continuous approximate version of HB and NAG). *For the model $f(x; \beta)$ with*

empirical loss function $L(\beta)$, let $\mu \in (0, 1)$ be the fixed momentum factor and η be the learning rate, the $\mathcal{O}(\eta)$ continuous approximate versions of the discrete HB (Eq. (2)) and NAG (Eq. (3)) are of the form

$$\alpha \ddot{\beta} + \dot{\beta} + \frac{\nabla L(\beta)}{1 - \mu} = 0, \quad (4)$$

where $\alpha = \frac{\eta(1+\mu)}{2(1-\mu)}$ for HB, and $\alpha = \frac{\eta(1-\mu+2\mu^2)}{2(1-\mu)}$ for NAG.

Eq. (4) follows from Kovachki and Stuart (2021) and we present an alternative proof in Appendix. Note that since the learning rate η is small, Proposition 1 indicates that, for the model parameter β , modifying GD with fixed momentum is equivalent to perturb the re-scaled gradient flow equation $d\beta/dt = \nabla L(\beta)/(1 - \mu)$ by a small term proportional to η . More importantly, this modification term offers us considerably more qualitative understanding regarding the dynamics of momentum-based methods than the re-scaled gradient flow, which will become more significant for large learning rate—a preferable choice in practice.

Over-parameterized regression. We consider the regression problem for the n -sample dataset $\{(x_i, y_i)\}_{i=1}^n$ where $n < d$ and assume the existence of the perfect solution, i.e., there exist interpolation solutions $\beta^* \in \mathbb{R}^d$ such that $\forall i \in \{1, \dots, n\} : x_i^T \beta^* = y_i$. For the parametric model $f(x; \beta) = \beta^T x$, we use the quadratic loss $\ell_i = (f(x_i; \beta) - y_i)^2$ and the empirical loss $L(\beta)$ is

$$L(\beta) = \frac{1}{2n} \sum_{i=1}^n \ell_i(\beta) = \frac{1}{2n} \sum_{i=1}^n (f(x_i; \beta) - y_i)^2. \quad (5)$$

Diagonal linear networks. The diagonal linear network (Woodworth et al. 2020) is a popular proxy model for DNNs. It corresponds to an equivalent linear predictor $f(x; \beta) = \theta^T x$, where $\theta = \theta(\beta)$ is parameterized by the model parameters β . For the diagonal linear networks considered in this paper, we study the 2-layer diagonal linear network, which corresponds to the parameterization¹ of $\theta = u \odot u - v \odot v$ in the sense that

$$f(x; \beta) := f(x; u, v) = (u \odot u - v \odot v)^T x, \quad (6)$$

and the model parameters are $\beta = (u, v)$, where $u, v \in \mathbb{R}^d$. We slightly abuse the notation of $L(\theta) = L(\beta)$. Our goal in this paper is to characterize the implicit bias of HB and NAG by precisely capturing the property of the limit point of θ and its dependence on various parameters such as the learning rate and the initialization of parameters for diagonal linear networks $f(x; \beta)$ trained with HB and NAG.

The use of $\mathcal{O}(\eta)$ order of convergence. The main reason why we use the $\mathcal{O}(\eta)$ continuous approximation for HB and NAG is that we aim to precisely characterize the role of momentum in the implicit bias of the widely-studied GF, which is the $\mathcal{O}(\eta)$ approximate continuous version of GD, for diagonal linear networks. The same order of approximations ($\mathcal{O}(\eta)$ in our case) for both momentum-based methods and GD should be used to make a “fair” comparison on their implicit bias.

¹A standard diagonal linear network is $\theta = u \odot v$, which is shown in (Woodworth et al. 2020) to be equivalent to our parameterization here. We further discuss this in Appendix.

3 Implicit Bias of HB and NAG Flow for Diagonal Linear Nets

To clearly reveal the difference between the implicit bias of (S)GD and momentum-based methods, we start with discussing existing results under the unbiased initialization assumption, and our main result is summarized in Theorem 3. We then discuss the dynamics of θ for diagonal linear networks trained with HB and NAG flow, which is necessary for the proof of Theorem 3 and may be of independent interest.

For convenience, given a diagonal linear network Eq. (6), let $\xi = (\xi_1, \dots, \xi^d) \in \mathbb{R}^d$ where $\forall i \in \{1, \dots, d\} : \xi_j = |u_j(0)| |v_j(0)|$ measures the scale of the initialization, we first present the definition of the unbiased initialization assumed frequently in previous works (Woodworth et al. 2020; Azulay et al. 2021; Pesme, Pillaud-Vivien, and Flammarion 2021).

Definition 2 (Unbiased initialization for diagonal linear networks). *The initialization for the diagonal linear network Eq. (6) is unbiased if $u(0) = v(0)$, which implies that $\theta(0) = 0$ and $\xi = u(0) \odot v(0)$.*

Implicit bias of GF. Recently, Woodworth et al. (2020); Azulay et al. (2021) showed that, for diagonal linear network with parameterization Eq. (6), if the initialization is unbiased (Definition 2) and $\theta(t) = u(t) \odot u(t) - v(t) \odot v(t)$ converges to the interpolation solution, i.e., $\forall i \in \{1, \dots, n\} : \theta^T(\infty)x_i = y_i$, then under GF $\theta^{\text{GF}}(\infty)$ implicitly solves the constrained optimization problem: $\theta^{\text{GF}}(\infty) = \arg \min_{\theta} Q_{\xi}^{\text{GF}}(\theta)$, s.t. $X\theta = y$, where $Q_{\xi}^{\text{GF}}(\theta) = \sum_{j=1}^d \left[\theta_j \operatorname{arcsinh}(\theta_j / (2\xi_j)) - \sqrt{4\xi_j^2 + \theta_j^2} + 2\xi_j \right] / 4$. The form of $Q_{\xi}^{\text{GF}}(\theta)$ highlights the transition from kernel regimes to rich regimes of diagonal linear networks under gradient flow up to different scales of the initialization: the initialization $\xi \rightarrow \infty$ corresponds to the *kernel regime* or *lazy regime* where $Q_{\xi}^{\text{GF}}(\theta) \propto \|\theta\|_2^2$ and the parameters only move slowly during training, and $\xi \rightarrow 0$ corresponds to the *rich regime* where $Q_{\xi}^{\text{GF}}(\theta) \rightarrow \|\theta\|_1$ and the corresponding solutions enjoy better generalization properties for sparse regression. For completeness, we characterize the implicit bias of GF without requiring the unbiased initialization (Definition 2) in the following proposition.

Proposition 2 (Implicit bias of GF for diagonal linear net with biased initialization). *For diagonal linear network Eq. (6) with biased initialization ($u(0) \neq v(0)$), if $u(t)$ and $v(t)$ follow the gradient flow dynamics for $t > 0$, i.e., $\dot{u} = -\nabla_u L$ and $\dot{v} = -\nabla_v L$, and if the solution converges to the interpolation solution, then*

$$\theta(\infty) = \arg \min_{\theta} Q_{\xi}^{\text{GF}}(\theta) + \theta^T \mathcal{R}^{\text{GF}}, \text{ s.t. } X\theta = y \quad (7)$$

where $\mathcal{R}^{\text{GF}} = (\mathcal{R}_1^{\text{GF}}, \dots, \mathcal{R}_d^{\text{GF}})^T \in \mathbb{R}^d$, $\forall j \in \{1, \dots, d\} : \mathcal{R}_j^{\text{GF}} = \operatorname{arcsinh}(\theta_j(0) / 2\xi_j) / 4$.

Compared to the unbiased initialization case, besides Q_{ξ}^{GF} , an additional term \mathcal{R}^{GF} that depends on $\theta(0)$ is required to capture the implicit bias when the initialization is biased. Note that \mathcal{R}^{GF} indicates that $\theta(\infty)$ also depends on

the direction of the initialization and serves as a kind of self-regularization.

3.1 Implicit bias of HB and NAG flow

Gunasekar et al. (2018); Wang et al. (2022) argued that momentum does not change the implicit bias of GF for single layer model for both linear regression and classification. For DNNs, will modifying GF with the widely adopted momentum change the implicit bias? If it does, will momentum-based methods lead to solutions that have better generalization properties? In the following, we characterize the implicit bias of HB and NAG flow (Proposition 1) for diagonal linear networks to compare with that of GF and answer these questions. For completeness, we do not require the unbiased initialization $u(0) = v(0)$ condition and let $\exp(a) \in \mathbb{R}^d$ denote the vector $(e^{a_1}, \dots, e^{a_d})^T$ for a vector $a \in \mathbb{R}^d$. We now present our main theorem.

Theorem 3 (Implicit bias of HB and NAG flow for diagonal linear networks). *For diagonal linear network Eq. (6), let $\mathcal{R}^{\text{M}} = (\mathcal{R}_1^{\text{M}}, \dots, \mathcal{R}_d^{\text{M}})^T \in \mathbb{R}^d$, if $u(t)$ and $v(t)$ follow the $\mathcal{O}(\eta)$ approximate continuous version of HB and NAG Eq. (4) for $t \geq 0$ and if the solution $\theta(\infty) = u(\infty) \odot u(\infty) - v(\infty) \odot v(\infty)$ converges to the interpolation solution, then, neglecting all terms of the order $\mathcal{O}(\eta^2)$,*

$$\theta(\infty) = \arg \min_{\theta} Q_{\bar{\xi}(\infty)}^{\text{M}}(\theta) + \theta^T \mathcal{R}^{\text{M}}, \text{ s.t. } X\theta = y \quad (8)$$

where we define $\bar{\xi}(\infty) = \xi \odot \exp(-\alpha\phi(\infty))$, and $\forall j \in \{1, \dots, d\} :$

$$\begin{aligned} \mathcal{R}_j^{\text{M}} &= \frac{1}{4} \operatorname{arcsinh} \left(\frac{\theta_j(0)}{2\xi_j} + \frac{4\alpha\partial_{\theta_j} L(\theta(0))}{1-\mu} \sqrt{1 + \frac{\theta_j^2(0)}{4\xi_j^2}} \right) \\ \phi(\infty) &= \frac{8}{(1-\mu)^2} \int_0^{\infty} \nabla_{\theta} L(\theta(s)) \odot \nabla_{\theta} L(\theta(s)) ds, \\ Q_{\bar{\xi}(\infty)}^{\text{M}}(\theta) &= \frac{1}{4} \sum_{j=1}^d \left[\theta_j \operatorname{arcsinh} \left(\frac{\theta_j}{2\bar{\xi}_j(\infty)} \right) - \sqrt{4\bar{\xi}_j^2(\infty) + \theta_j^2} + 2\bar{\xi}_j(\infty) \right]. \end{aligned} \quad (9)$$

Specifically, α is chosen as $\frac{\eta(1+\mu)}{2(1-\mu)}$ if we run HB and $\alpha = \frac{\eta(1-\mu+2\mu^2)}{2(1-\mu)}$ for NAG.

Remark. The $Q_{\bar{\xi}(\infty)}^{\text{M}}$ part for HB and NAG flow has a formulation similar to Q_{ξ}^{GF} of GF: both of them are the *hyperbolic entropy* (Ghai, Hazan, and Singer 2020). The transition from kernel regime to rich regime by decreasing ξ from ∞ to 0 also exists for HB and NAG (see Appendix). The difference between $Q_{\bar{\xi}(\infty)}^{\text{M}}$ and Q_{ξ}^{GF} lies in that HB and NAG flow induce an extra initialization mitigation effect: given ξ , $Q_{\bar{\xi}(\infty)}^{\text{M}}$ for HB and NAG flow is equivalent to the hyperbolic entropy of GF with a smaller initialization scale since $\bar{\xi}(\infty)$ is strictly smaller than ξ due to the fact that $\phi(\infty)$ is a positive integral and finite. As a result, $Q_{\bar{\xi}(\infty)}^{\text{M}}$ is closer to an ℓ_1 -norm of θ than Q_{ξ}^{GF} . Furthermore, compared to the implicit

bias of GF when the initialization is biased (Proposition 2), an additional term in \mathcal{R}^M that depends on the initialization of gradient explicitly is required to capture the implicit bias of HB and NAG flow. Such dependence is as expected since the first step update of momentum methods simply assigns the initialization of gradient to the momentum, which is crucial for the following updates. Therefore, Theorem 3 takes an important step towards positively answering our fundamental question (Q) in the sense that *momentum changes the implicit bias of GD for diagonal linear networks*.

A natural question following the fact that HB and NAG flow induce different implicit bias compared to GF is: will this difference lead to better generalization properties of HB and NAG? The implicit bias of HB and NAG flow is captured by two distinct parts, the hyperbolic entropy $Q_{\xi(\infty)}^M$ and \mathcal{R}^M , where the effects of momentum on $Q_{\xi(\infty)}^M$ is beneficial for generalization while the effects on \mathcal{R}^M may hinder the generalization performance and is affected by the biased initialization. Thus the answer highly depends on various conditions.

Due to the aforementioned initialization mitigation effects of HB and NAG, GD with a smaller initialization might achieve a similar regularization effect as HB and NAG. The main harm, however, of using GD with a smaller initialization is the saddle point escape issue: very small initialization scales lead to the issue that these scales correspond to the initialization highly close to a saddle point (here $u = v = 0$) that might be difficult to escape. This reveals the benefit of the extra effects brought by momentum, i.e., avoiding the saddle point escape problem by using a relatively large initialization to achieve good generalization.

In the following, we present a detailed analysis with corresponding numerical experimental results to compare the implicit bias of GF and that of HB and NAG flow for the case of both unbiased and biased initialization, respectively.

3.2 Comparison of HB/NAG flow and (S)GF for unbiased initialization

When the initialization is unbiased, it is worth to mention that a recent work (Pesme, Pillaud-Vivien, and Flammarion 2021) studied the stochastic version of gradient flow, the stochastic gradient flow (SGF), and revealed that the existence of sampling noise changes the implicit bias of GF in the sense that $\theta^{\text{SGF}}(\infty) = \arg \min_{\theta} Q_{\xi(\infty)}^{\text{SGF}}(\theta)$ under the constraint $X\theta = y$, where

$$Q_{\xi(\infty)}^{\text{SGF}}(\theta) = \sum_{j=1}^d \frac{1}{4} \left[\theta_j \operatorname{arcsinh} \left(\frac{\theta_j}{2\tilde{\xi}_j(\infty)} \right) - \sqrt{4\tilde{\xi}_j^2(\infty) + \theta_j^2} + 2\tilde{\xi}_j(\infty) \right] \quad (10)$$

with $\tilde{\xi}(\infty)$ being strictly smaller than ξ . The remarkable point appears when we compare $Q_{\xi(\infty)}^M$ with $Q_{\xi(\infty)}^{\text{SGF}}$: although SGF and momentum-based methods modify GF differently, i.e., SGF adds stochastic sampling noise while momentum-based methods add momentum to GF, *both of them induce an effect equivalent to reducing the initialization scale!* The

difference between them lies in the way how they control such initialization mitigation effect. For SGF this is controlled by the integral of loss function, while the effect depends on the integral of gradients for HB and NAG flow.

To show the difference between (S)GF and momentum-based methods HB and NAG flow, we note that \mathcal{R}_j^M in Theorem 3 becomes

$$\mathcal{R}_j^M = \operatorname{arcsinh} \left(\frac{4\alpha(X^T y)_j}{n(1-\mu)} \right),$$


which is also determined by the dataset, and \mathcal{R}^{GF} in Proposition 2 is simply zero. Therefore, as long as the initialization of gradients $X^T y = o(\alpha^{-1}n(1-\mu))$, i.e., \mathcal{R}^M is small compared to $Q_{\xi(\infty)}^M$ such that only $Q_{\xi(\infty)}^M$ matters for characterizing the implicit bias, HB and NAG flow will exhibit better generalization properties for sparse regression due to the initialization mitigation effects of HB and NAG flow that lead $Q_{\xi(\infty)}^M$ to be closer to the ℓ_1 -norm of θ than Q_{ξ}^{GF} . On the other hand, when \mathcal{R}_j^M is not small compared to $Q_{\xi(\infty)}^M$, the initialization mitigation effects of HB and NAG flow may not be significant, thus there may not be generalization benefit for HB and NAG flow.

To summarize, for unbiased initialization, HB and NAG outperform GD regarding the generalization when $\alpha X^T y$ is much smaller than $n(1-\mu)$ and they would have worse generalization performance than GD otherwise. In the following, we conduct numerical experiments to verify this claim.


Numerical Experiments. We consider the over-parameterized sparse regression. For the dataset $\{(x_i, y_i)\}_{i=1}^n$ where $x_i \in \mathbb{R}^d$ and $y_i \in \mathbb{R}$, we set $n = 40, d = 100$ and $x_i \sim \mathcal{N}(0, I)$. y_i is generated by $y_i = x_i^T \theta^*$ where $\theta^* \in \mathbb{R}^d$ is the ground truth solution. We let 5 components of θ^* be non-zero. Our models are 2-layer diagonal linear networks $f(x; \beta) = u \odot u - v \odot v$. We use $\|\xi\|_1$ to measure the scale of initialization. The initialization of parameters is unbiased by letting $u(0) = v(0) = c\mathbf{e}_d$ where c is a constant and $\|\xi\|_1 = c^2 d$. We consider training algorithms GD, SGD, HB, and NAG. And the generalization performance of the solution for each training algorithm is measured by the distance $D(\theta(\infty), \theta^*) = \|\theta(\infty) - \theta^*\|_2^2$. Since \mathcal{R}^M is determined by the dataset, to control its magnitude, we build two new datasets $\mathcal{D}_\varepsilon = \{(x_{i;\varepsilon}, y_{i;\varepsilon})\}_{i=1}^d$ where $\forall i \in \{1, \dots, d\} : x_{i;\varepsilon} = \varepsilon x_i, y_{i;\varepsilon} = \varepsilon y_i$. We then train diagonal linear networks using GD and momentum-based methods HB and NAG on each dataset, respectively, and learning rate $\eta = 3 \times 10^{-2}$ and momentum factor $\mu = 0.9$. As shown in Fig. 1, as we decrease the value of ε which decreases the magnitude of \mathcal{R}^M , the generalization benefit of HB and NAG becomes more significant since their initialization mitigation effects are getting more important. Note that Fig. 1 also reveals the transition to rich regime by decreasing the initialization scales.

3.3 Comparison of HB/NAG flow and GF for biased initialization


If the initialization is biased, i.e., $u(0) \neq v(0)$, both the implicit bias of GF and that of HB and NAG flow additionally




(a) $\varepsilon = 0.6$



(a) x-axis denotes $\|\xi\|_1$



(b) $\varepsilon = 0.2$



(b) x-axis denotes φ

Figure 1: $D(\theta(\infty), \theta^*)$ for diagonal linear networks with unbiased initialization trained with different algorithms and ε (smaller ε for smaller $\nabla L(\theta(0))$). x-axis denotes $\|\xi\|_1$.

depend on $\theta(0)$ (\mathcal{R}^{GF} for GF in Proposition 2 and \mathcal{R}^{M} in Theorem 3 for HB and NAG flow) besides the hyperbolic entropy. Compared to \mathcal{R}^{GF} , \mathcal{R}^{M} also includes the explicit dependence on the initialization of gradient that is proportional to $\alpha \nabla L(\theta(0))$. Therefore, recall that α is the order of η , if $\nabla L(\theta(0)) = o(\alpha^{-1}n(1-\mu))$ and $\alpha \nabla L(\theta(0))$ is small compared to $\theta(0)$, then \mathcal{R}^{GF} is close to \mathcal{R}^{M} , leading to the fact that the difference between the implicit bias of GF and that of HB and NAG flow are mainly due to the initialization mitigation effects of HB and NAG. As a result, we can observe the generalization advantages of HB and NAG over GF (Fig. 2(a)). However, when the initialization is only slightly biased, i.e., $u(0) \neq v(0)$ and $u(0)$ is close to $v(0)$, the dependence on $\nabla L(\theta(0))$ of the solutions of HB and NAG is important and the generalization benefit of HB and NAG for sparse regression may disappear.

Numerical Experiments. We use the same dataset $\{(x_i, y_i)\}_{i=1}^d$ as in Section 3.2. We set $\eta = 10^{-1}$ and $\mu = 0.9$. To characterize the influence of the extent of the biased part of the initialization, we let $u(0) = \varphi c e_d$ and $v(0) = \varphi^{-1} c e_d$ where $\varphi \in (0, 1]$ is a constant measuring the extent of the unbiased part of the initialization. In this way, for any φ , we have $\xi_j = |u_j(0)| |v_j(0)| = c^2$. In order to verify the above theoretical claims, we conduct two sets of experiments: (i). We train diagonal linear network

Figure 2: $D(\theta(\infty), \theta^*)$ for diagonal linear networks with biased initialization trained with different algorithms and (a). different initialization scales with $\varphi = 0.03$; (b). different extents of the biased part of the initialization (smaller φ implies more biased initialization) and $\|\xi\|_1 = 0.0046$.

with different algorithms for different scales of initialization $\|\xi\|_1$ and fixed φ . As shown in Fig. 2(a), as a result of the initialization mitigation effects, HB, NAG, and SGD exhibit better generalization performance than GD for sparse regression. (ii). We fix $\|\xi\|_1$ and train diagonal linear networks with different biased initialization (different values of φ). As shown in Fig. 2(b), as we increasing φ , the initialization becomes less biased and the extra dependence on the initialization of gradient of HB and NAG outperforms their initialization mitigation effects, and, as a result, the generalization benefits of momentum-based methods disappear.

3.4 Dynamics for θ of diagonal linear networks under HB and NAG Flow

For diagonal linear networks Eq. (6), dynamics for θ under HB and NAG flow is crucial to the proof of Theorem 3, and may be of independent interest. Interestingly, different from diagonal linear networks under gradient flow where θ follows a mirror flow or stochastic gradient flow where θ follows a stochastic mirror flow with time-varying potential, due to the second-order ODE nature of HB and NAG flow as formulated in Eq. (4), θ does not directly follow a mirror flow. Instead, HB and NAG flow is special—it is $\theta + \alpha \dot{\theta}$ that

follows a mirror flow form with time-varying potential, as shown below.

Proposition 4 (Dynamics of θ for diagonal nets trained with HB and NAG flow). *For diagonal linear networks Eq. (6) trained with HB and NAG flow (Eq. (4)) and initialized as $u(0) = v(0)$ and $u(0) \odot u(0) = \xi \in \mathbb{R}^d$, let $\bar{\theta}_\alpha := \theta + \alpha \dot{\theta} \in \mathbb{R}^d$ and its j -th component be $\bar{\theta}_{\alpha;j}$, then $\bar{\theta}_\alpha$ follows a mirror flow form with time-varying potential (\mathcal{R}^M is defined in Theorem 3) in the sense that $\forall j \in \{1, \dots, d\}$:*

$$\frac{d}{dt} \nabla [Q_{\xi;j}^M(\bar{\theta}_\alpha, t) + \bar{\theta}_{\alpha;j} \mathcal{R}_j^M] = -\frac{\partial_{\theta_j} L(\theta)}{1-\mu}, \quad (11)$$

where $Q_{\xi;j}^M(\bar{\theta}_\alpha, t)$ is given by

$$Q_{\xi;j}^M(\bar{\theta}_\alpha, t) = \frac{1}{4} \left[\bar{\theta}_{\alpha;j} \operatorname{arcsinh} \left(\frac{\bar{\theta}_{\alpha;j}}{2\bar{\xi}_j(t)} \right) - \sqrt{4\bar{\xi}_j^2(t) + \bar{\theta}_{\alpha;j}^2} + 2\bar{\xi}_j(t) \right]$$

and $\bar{\xi}_j(t) = \xi_j e^{-\alpha \phi_j(t)}$ with

$$\phi_j(t) = \frac{8}{(1-\mu)^2} \int_0^t \partial_{\theta_j} L(\theta(s)) \partial_{\theta_j} L(\theta(s)) ds. \quad (12)$$

Compared to the mirror flow form of diagonal linear networks under gradient flow $d\nabla Q_\xi^{\text{GF}}(\theta)/dt = -\nabla L(\theta)$, there are three main differences in Eq. (11): (i). It is a second-order ODE since Eq. (11) can be written as $\alpha \nabla^2 Q_{\xi;j}^M(\bar{\theta}_\alpha, t) \ddot{\theta}_j + \nabla^2 Q_{\xi;j}^M(\bar{\theta}_\alpha, t) \dot{\theta}_j + \frac{\partial \nabla Q_{\xi;j}^M(\bar{\theta}_\alpha, t)}{\partial t} + \frac{\partial_{\theta_j} L(\theta)}{1-\mu} = 0$, while the dynamics of GF is a first-order ODE; (ii). It is $\bar{\theta}_\alpha$, not θ , appears in the mirror flow potential for diagonal linear networks under HB and NAG flow, and an extra term depending on the initialization of gradients is included; (iii). The hyperbolic entropy part of the mirror flow potential $Q_{\xi;j}^M(\bar{\theta}_\alpha, t)$ under HB and NAG flow is a time-varying one, and the time-varying part mainly mitigates the influence of the initialization ξ ($\bar{\xi}_j(t) \leq \xi$ for any $t \geq 0$).

3.5 Effects of Hyper-parameters for Implicit Bias

As a result of the fact that momentum-based methods (HB and NAG) add a perturbation proportional to the learning rate η to re-scaled gradient flow (Proposition 1), the difference of their implicit bias depends on η : the limit $\eta \rightarrow 0$ leads to $\bar{\xi}(\infty) \rightarrow \xi$ and, as a consequence, $Q_{\bar{\xi}(\infty)}^M \rightarrow Q_\xi^{\text{GF}}$. Therefore, for small learning rate, the implicit bias of momentum-based methods and that of GD are almost the same. This observation coincides with the experience of Rumelhart, Hinton, and Williams (1986); Kovachki and Stuart (2021) that setting momentum factor as 0 returns the same solution as reducing the learning rate when momentum factor is non-zero. The discrepancy between the implicit bias of momentum-based methods and that of GD becomes significant for moderate learning rate and momentum factor.

To verify this, we set $\|\xi\|_1 = 0.1240$, and run: (i). GD with $\eta = 10^{-2}$; (ii). HB and NAG with $\mu = 0.9$ and different η ; (iii). HB and NAG with $\eta = 10^{-2}$ and different μ . We present the generalization performance $D(\theta(t), \theta^*)$ during


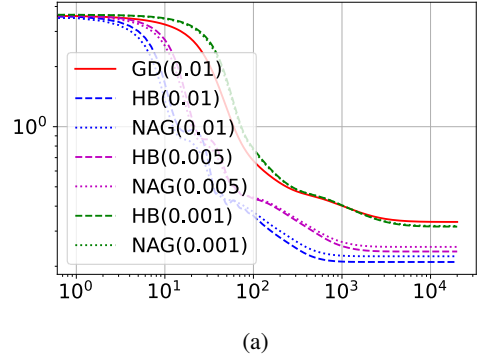


Figure 3: Diagonal nets trained with different algorithms and hyper-parameters: (a). $D(\theta(t), \theta^*)$ for different η (numbers in the brackets) and $\mu = 0.9$. (b). $D(\theta(t), \theta^*)$ for different μ (numbers in the brackets) and $\eta = 0.01$. x-axis denotes iterations.

training for each algorithm with its corresponding training parameters in Fig. 3(a) and Fig. 3(b). These results clearly reveal that both decreasing the learning rate and the momentum factor make the difference between the implicit bias of momentum-based methods and that of GD not significant. Experimental details are in Appendix.

4 Conclusion

In this paper, we have targeted on the unexplored regularization effect of momentum-based methods and we have shown that, *unlike* the single layer linear network, momentum-based methods HB and NAG flow exhibit different implicit bias compared to GD for diagonal linear networks. In particular, we reveal that HB and NAG flow induce an extra initialization mitigation effect similar to SGD that is beneficial for generalization of sparse regression and controlled by the integral of the gradients, learning rate, data matrix, and the momentum factor. In addition, the implicit bias of HB and NAG flow also depends on the initialization of both parameters and gradients explicitly, which may also hinder the generalization, while GD and SGD only depend on the initialization of parameters.

References

- Ali, A.; Dobriban, E.; and Tibshirani, R. 2020. The Implicit Regularization of Stochastic Gradient Flow for Least Squares. In *International Conference on Machine Learning*.
- Azulay, S.; Moroshko, E.; Nacson, M. S.; Woodworth, B.; Srebro, N.; Globerson, A.; and Soudry, D. 2021. On the Implicit Bias of Initialization Shape: Beyond Infinitesimal Mirror Descent. arXiv:2102.09769.
- Barrett, D.; and Dherin, B. 2021. Implicit Gradient Regularization. In *International Conference on Learning Representations*.
- Chizat, L.; and Bach, F. 2020. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*.
- Chizat, L.; Oyallon, E.; and Bach, F. 2019. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*.
- Duchi, J.; Hazan, E.; and Singer, Y. 2011. Adaptive subgradient methods for online learning and stochastic optimization. In *Journal of Machine Learning Research*.
- Even, M.; Pesme, S.; Gunasekar, S.; and Flammarion, N. 2023. (S)GD over Diagonal Linear Networks: Implicit Regularisation, Large Stepsizes and Edge of Stability. In arXiv:2302.08982.
- Ghai, U.; Hazan, E.; and Singer, Y. 2020. Exponentiated gradient meets gradient descent. In *International Conference on Algorithmic Learning Theory*.
- Ghosh, A.; Lyu, H.; Zhang, X.; and Wang, R. 2023. Implicit regularization in Heavy-ball momentum accelerated stochastic gradient descent. arXiv:2302.00849.
- Gunasekar, S.; Lee, J.; Soudry, D.; and Srebro, N. 2018. Characterizing Implicit Bias in Terms of Optimization Geometry. In *International Conference on Machine Learning*.
- Jelassi, S.; and Li, Y. 2022. Towards understanding how momentum improves generalization in deep learning. arXiv:2207.05931.
- Ji, Z.; and Telgarsky, M. 2019. Gradient Descent Aligns the Layers of Deep Linear Networks. In *International Conference on Learning Representations*.
- Kingma, D. P.; and Ba, J. 2017. Adam: A method for stochastic optimization. In arXiv:1412.6980.
- Kovachki, N. B.; and Stuart, A. M. 2021. Continuous Time Analysis of Momentum Methods. In *Journal of Machine Learning Research*.
- Li, Z.; Luo, Y.; and Lyu, K. 2021. Towards Resolving the Implicit Bias of Gradient Descent for Matrix Factorization: Greedy Low-Rank Learning. In *International Conference on Learning Representations*.
- Lyu, B.; and Zhu, Z. 2022. Implicit Bias of Adversarial Training for Deep Neural Networks. In *International Conference on Learning Representations*.
- Lyu, B.; and Zhu, Z. 2023. Implicit Bias of (Stochastic) Gradient Descent for Rank-1 Linear Neural Network. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 58166–58201. Curran Associates, Inc.
- Lyu, K.; and Li, J. 2020. Gradient Descent Maximizes the Margin of Homogeneous Neural Networks. In *International Conference on Learning Representations*.
- Nacson, M. S.; Gunasekar, S.; Lee, J.; Srebro, N.; and Soudry, D. 2019. Lexicographic and Depth-Sensitive Margins in Homogeneous and Non-Homogeneous Deep Models. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 4683–4692. PMLR.
- Nesterov, Y. 1983. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. In *Soviet Mathematics Doklady*.
- Nesterov, Y. 2014. Introductory Lectures on Convex Optimization: A Basic Course. In *Springer Publishing Company, Incorporated*.
- Papazov, H.; Pesme, S.; and Flammarion, N. 2024. Leveraging Continuous Time to Understand Momentum When Training Diagonal Linear Networks. In Dasgupta, S.; Mandt, S.; and Li, Y., eds., *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, 3556–3564. PMLR.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch. In *Advances in Neural Information Processing Systems 2017 Workshop Autodiff*.
- Pesme, S.; Pillaud-Vivien, L.; and Flammarion, N. 2021. Implicit Bias of SGD for Diagonal Linear Networks: a Provable Benefit of Stochasticity. In *Advances in Neural Information Processing Systems*.
- Pillaud-Vivien, L.; Reygner, J.; and Flammarion, N. 2020. Label Noise (stochastic) Gradient Descent Implicitly Solves the Lasso for Quadratic Parametrisation. In *Conference on Learning Theory*.
- Polyak, B. 1964. Some Methods of Speeding Up the Convergence of Iteration Methods. In *Ussr Computational Mathematics and Mathematical Physics*.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1986. Learning internal representations by error propagation. In *Parallel Distributed Processing: Explorations in Microstructures in Cognition*, volume 1: Foundations.
- Shi, B.; Du, S. S.; Jordan, M. I.; and Su, W. J. 2018. Understanding the Acceleration Phenomenon via High-Resolution Differential Equations. In arXiv: 1810.08907.
- Soudry, D.; Hoffer, E.; Nacson, M. S.; Gunasekar, S.; and Srebro, N. 2018. The Implicit Bias of Gradient Descent on Separable Data. *Journal of Machine Learning Research*, 19(70): 1–57.
- Su, W.; Boyd, S.; and Candes, E. 2014. A differential equation for modeling nesterov’s accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems*.
- Sutskever, I.; Martens, J.; Dahl, G.; and Hinton, G. 2013. On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning*.

Wang, B.; Meng, Q.; Zhang, H.; Sun, R.; Chen, W.; Ma, Z.-M.; and Liu, T.-Y. 2022. Does Momentum Change the Implicit Regularization on Separable Data? In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 26764–26776. Curran Associates, Inc.

Wibisono, A.; Roelofs, R.; Stern, M.; Srebro, N.; and Recht, B. 2017. The Marginal Value of Adaptive Gradient Methods in Machine Learning. In *Advances in Neural Information Processing Systems*.

Wibisono, A.; Wilson, A. C.; and Jordan, M. I. 2016. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, 113(47): E7351–E7358.

Wilson, A. C.; Recht, B.; and Jordan, M. I. 2016. A Lyapunov analysis of momentum methods in optimization. In *arXiv:1611.02635*.

Woodworth, B.; Gunasekar, S.; Lee, J. D.; Moroshko, E.; Savarese, P.; Golan, I.; Soudry, D.; and Srebro, N. 2020. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*.

Yun, C.; Krishnan, S.; and Mobahi, H. 2021. A Unifying View on Implicit Bias in Training Linear Neural Networks. In *International Conference on Learning Representations*.