

# Adam Reduces a Unique Form of Sharpness: Theoretical Insights Near the Minimizer Manifold

author names withheld

Under Review for the Workshop on High-dimensional Learning Dynamics, 2024

## Abstract

Despite the popularity of Adam optimizer in practice, most theoretical analyses study SGD as a proxy, and little is known about how the solutions found by Adam differ. In this paper, we show that Adam reduces a specific form of sharpness measure shaped by its adaptive updates, leading to qualitatively different solutions from SGD. When the training loss is small, Adam wanders around the manifold of minimizers and takes semi-gradients to minimize this sharpness measure in an adaptive manner, a behavior we rigorously characterize via a continuous-time approximation using stochastic differential equations. We further illustrate how this behavior differs from that of SGD in a well-studied setting: When training overparameterized models with label noise, SGD has been shown to minimize the trace of the Hessian matrix,  $\text{tr}(\mathbf{H})$ , whereas we prove that Adam minimizes  $\text{tr}(\text{Diag}(\mathbf{H})^{1/2})$  instead. In solving sparse linear regression with diagonal linear networks, Adam provably achieves better sparsity and generalization than SGD due to this difference. Finally, we note that our proof framework applies not only to Adam but also to many other adaptive gradient methods, including but not limited to RMSProp, Adam-mini, and Adalayer. This provides a unified perspective for analyzing how adaptive optimizers reduce sharpness and may offer insights for future optimizer design.

## 1. Introduction

Due to the non-convexity of the loss landscape, neural networks trained in different ways can perform very differently on the test set, even if they achieve the same training loss or accuracy [27, 36, 49, 60]. To mathematically understand the generalization of neural networks, especially for overparameterized models that admit many global minimizers, a key step is to understand the *implicit bias* of optimization methods [45, 53]. That is, beyond just minimizing the training loss, *what kinds of solutions are different optimizers implicitly biased toward?*

Many theoretical works on implicit bias focused on (full-batch) gradient descent or its continuous variant, gradient flow. This includes the works on the implicit bias towards max-margin classifiers [25, 37, 44, 53], implicit bias towards min-norm solutions [39], and equivalence to kernel methods [7, 23].

All these implicit bias characterizations hold, or can be readily extended, to the stochastic variant of gradient descent, *i.e.*, Stochastic Gradient Descent (SGD). Another line of works [3, 8, 35] demonstrated that the gradient noise in SGD induces an additional form of implicit bias that reduces the *sharpness* of the solutions, a generalization measure that has been long observed to correlate with generalization [16, 22, 26, 27]. More specifically, these works focus on the dynamics of SGD when the training loss is already small and the iterates are close to a manifold of minimizers. Li et al. [35] introduced a general framework to analyze the dynamics of SGD near the minimizer manifold, showing that SGD will not stop at arbitrary global minimizers, but drift and diffuse around

the manifold, driving the iterates towards flatter regions of the loss landscape. This behavior is mathematically characterized by a Stochastic Differential Equation (SDE), termed as *slow SDE* [19], which accurately tracks the projected dynamics of SGD near the minimizer manifold over a timescale of  $\mathcal{O}(\eta^{-2})$ . The resulting dynamics reveal that SGD behaves like a gradient method on the manifold that takes semi-gradients to minimize a specific sharpness measure determined by the Hessian and gradient noise. See [Appendix A](#) for more details.

However, SGD is rarely used directly in modern deep learning. Instead, Adaptive Gradient Methods (AGMs) have become the de facto standard for training neural networks. Among them, Adam [28] innovatively combines the moving average of the first and second moments of gradients to determine an adaptive learning rate for each parameter, and provides faster convergence and better stability than SGD across various domains [1, 10, 50, 61].

Despite the popularity of Adam, little is known about its implicit bias, especially how it is different from SGD in terms of reducing sharpness. In the literature, Ma et al. [40] made attempts to generalize the slow SDE framework from SGD to Adam, but their analysis is specific to a two-dimensional loss function and involves a quasistatic approximation that lacks full mathematical rigor. Other works, such as Gu et al. [21], Liu et al. [36], leverage insights from the slow SDE developed for SGD to interpret empirical observations with Adam, but do not provide a theoretical analysis of Adam’s own dynamics. A rigorous analysis of Adam’s implicit bias in terms of sharpness remains an open problem.

**Our Contributions.** In this paper, we show that Adam implicitly reduces a unique form of sharpness and biases the iterates towards flatter regions in a way that is different from SGD, and provide separations between SGD and Adam in concrete theoretical cases.

1. In [Section 2](#), we generalize the slow SDE for SGD to Adam. The slow SDE approximates the dynamics of Adam near the minimizer manifold, and reveals that Adam behaves like an adaptive gradient method that minimizes a unique form of sharpness by taking semi-gradients on the manifold.
2. In [Appendix D](#), we prove theoretically the generalization benefit of Adam under label noise settings. We show that under label noise setting, the implicit regularizer of Adam will reduce to  $\text{tr}(\text{Diag}(\mathbf{H})^{1/2})$  where  $\mathbf{H}$  is the Hessian matrix. Compared to the  $\text{tr}(\mathbf{H})$  of SGD, this new kind of sharpness reduction usually aligns better with sparsity regularization, thus utilizing data more efficiently when the model is required to fit a sparse ground truth. We verify this anticipation experimentally through the diagonal net setting [57]. We also demonstrated the discrepancy of the implicit biases of Adam and SGD through the matrix factorization setting in [Appendix E](#).
3. Technically, our analysis holds for a general class of adaptive gradient methods (AGMs), including Adam, RMSProp, Adam-mini, and Adalayer. We develop several new tools that can be of independent interest, including a manifold projection operator tailored for AGMs, a high-probability convergence analysis for AGMs under PL conditions that directly gives a bound on  $\mathcal{L}(\theta_k) - \mathcal{L}^*$ .

## 2. Theoretical Analysis of Adam

In this section, we generalize the slow SDE for SGD to a general class of adaptive gradient methods (AGMs), including Adam. We first present our novel slow SDE for a general class of AGMs,

including Adam, and give an intuitive explanation for our results. Then, we discuss the difficulty of directly applying the slow-SDE framework to Adam and other AGMs and how we resolve the problems.

**A General Class of Adaptive Gradient Methods.** We define a general class of AGMs as follows:

$$\begin{aligned}\mathbf{m}_{k+1} &:= \beta_1 \mathbf{m}_k + (1 - \beta_1) \nabla \ell_k(\boldsymbol{\theta}_k) \\ \mathbf{v}_{k+1} &:= \beta_2 \mathbf{v}_k + (1 - \beta_2) V \left( \nabla \ell_k(\boldsymbol{\theta}_k) \nabla \ell_k(\boldsymbol{\theta}_k)^\top \right) \\ \boldsymbol{\theta}_{k+1} &:= \boldsymbol{\theta}_k - \eta S(\mathbf{v}_{k+1}) \mathbf{m}_{k+1}.\end{aligned}$$

where  $S : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$  is  $\rho_s$ -smooth, positive definite and satisfies  $S(\mathbf{v}) \preceq \epsilon^{-1} I$  for some  $\epsilon > 0$  and any  $\mathbf{v} \in \mathbb{R}^d$ , and  $V : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$  is linear. A number of currently used optimization algorithms, such as RMSProp, Adam, Adam-mini, Adafactor<sup>1</sup>, Adalayer and AdaSGD all fit this framework. Note that we do not consider weight decays or bias corrections in these optimizers. Some examples of  $V$  and  $S$  functions are listed in Table 1, including the AdamE- $\lambda$  optimizer that will be introduced in Appendix D as a tool to tune the implicit bias of Adam.

### 2.1. Slow SDE Analysis for AGMs

Our SDE for AGMs characterizes the training dynamics near the manifold  $\Gamma$ . First we rigorously define the preconditioned projection mapping  $\Phi_S$  and the SDE projection formula as an extension to the  $\Phi$  and  $P_\zeta$  mentioned in Appendix A, after which we present the SDE for AGMs we derived.

**Definition 2.1** (Preconditioner Flow Projection). *Fix a point  $\theta_{\text{null}} \notin \Gamma$ . Given a Positive Semi-Definite matrix  $S$ . For  $x \in \mathbb{R}^d$ , consider the preconditioner flow  $\frac{dx(t)}{dt} = -S \nabla \mathcal{L}(x(t))$  with  $x(0) = x$ . We denote the preconditioner flow projection of  $x$  as  $\Phi_S(x)$ , i.e.  $\Phi_S(x) := \lim_{t \rightarrow +\infty} x(t)$  if the limit exists and belongs to  $\Gamma$ , and  $\Phi_S(x) = \theta_{\text{null}}$  otherwise.*

**Definition 2.2.** *For any  $\zeta \in \Gamma$  and any differential form  $A d\mathbf{W}_t + \mathbf{b} dt$  in Itô calculus, where  $A \in \mathbb{R}^{d \times d}$ , and  $\mathbf{b} \in \mathbb{R}^d$ . We use  $P_{\zeta, S}(A d\mathbf{W}_t + \mathbf{b} dt)$  as a shorthand for the differential form  $\partial \Phi_S(\zeta) A d\mathbf{W}_t + (\partial \Phi_S(\zeta) \mathbf{b} + \frac{1}{2} \partial^2 \Phi_S(\zeta) [A A^\top]) dt$ .*

**Definition 2.3** (Slow SDE for AGMs). *given learning rate  $\eta$ ,  $\frac{1-\beta_2}{\eta^2} = c$ ,  $\mathbf{v}_0 \in \mathbb{R}^d$ ,  $\mathbf{S}_t := S(\mathbf{v}(t))$ , and  $\zeta_0 \in \Gamma$ ,  $\mathbf{v}_0 \in \mathbb{R}^d$ , we define  $\zeta(t)$  as the solution of the following SDE with initial point  $(\zeta(0), \mathbf{v}(0)) = (\zeta_0, \mathbf{v}_0)$ :*

$$\begin{cases} d\zeta(t) = P_{\zeta, S(t)} \left( \underbrace{\Sigma_\parallel^{1/2}(\zeta(t); S(t)) d\mathbf{W}_t}_{\text{diffusion}} - \underbrace{\frac{1}{2} S(t) \nabla^3 \mathcal{L}(\zeta) [\Sigma_\diamond(\zeta(t); S(t))]}_{\text{drift}} dt \right), \\ d\mathbf{v}(t) = \underbrace{c (V(\Sigma(\zeta)) - \mathbf{v}) dt}_{\text{Preconditioner drift}}. \end{cases} \quad (1)$$

$$\Sigma_\diamond(\zeta; S) = S \Sigma(\zeta) S - \Sigma_\parallel(\zeta; S), \quad \Sigma_\parallel(\zeta; S) = \partial \Phi_S(\zeta) S \Sigma(\zeta) S \partial \Phi_S(\zeta).$$

1. We ignore update clipping, i.e. we adopt the Algorithm 2 in Shazeer and Stern [51].

Note that the drift term in  $d\zeta(t)$  can be interpreted as an *adaptive semi-gradient descent* process, in that this term drives the dynamics towards optimizing an adaptive loss function

$$\mu(\zeta, v) = \langle \nabla^2 \mathcal{L}(\zeta), \Sigma_\diamond(\zeta(t); S(t)) \rangle$$

as if  $\Sigma_\diamond(\zeta(t); S(t))$  has no dependence on  $\zeta$ ; also this gradient flow is preconditioned by a positive definite matrix  $S(t)$ . Recall that the drift term in the slow SDE for SGD can be seen as a semi-gradient descent. In the AGM framework, it takes  $\Theta(\eta^{-2})$  time for the preconditioner  $S(t)$  to make a significant (i.e.  $\Theta(1)$ ) change, which coincides with the moving speed of the slow SDE of  $\zeta$ . Therefore, compared to that of SGD, our SDE includes a new formula that tracks the motion of the preconditioner and injects adaptiveness accordingly in the semi-gradient descent process.

We could prove that  $\zeta(t)$  always stays on the manifold  $\Gamma$ . And next, we present our main theorem, and show that the above SDE in Equation (1) track the trajectory of Adam in a weak approximation sense.

**Assumption 2.1.** *The loss function  $\mathcal{L}(\cdot)$  and the matrix square root of the noise covariance  $\Sigma^{1/2}(\cdot)$  are  $C^\infty$ -smooth. Besides, we assume that  $\|\nabla \ell(\theta; \xi)\|_2$  is bounded by a constant for all  $\theta$  and  $\xi$ .*

**Assumption 2.2.**  *$\Gamma$  is a compact manifold.*

**Theorem 2.1.** *Under Assumption A.1–2.2. Let  $T > 0$  be a constant and let  $\mathbf{X}(t) = (\zeta(t), v(t))$  be the solution to Equation (1) with initial condition:*

$$\zeta(0) = \Phi(\theta_0) \in \Gamma, \quad v(0) = v_0 \in \mathbb{R}^d,$$

*and we define that the parameters of Adam as  $\bar{\mathbf{X}}_t := (\Phi_{S_t}(\theta_t), v_t)$ . For any  $C^3$ -smooth function  $g(\theta)$ ,*

$$\max_{0 \leq t \leq \frac{T}{\eta^2}} \left| \mathbb{E}[g(\bar{\mathbf{X}}_t)] - \mathbb{E}[g(\mathbf{X}(t\eta^2))] \right| = \tilde{\mathcal{O}}(\eta^{0.25}),$$

*where  $\tilde{\mathcal{O}}(\cdot)$  hides logarithmic factors and constants that are independent of  $\eta$  but may depend on  $g(\theta)$ .*

Theorem 2.1 shows that, in the small  $\eta$  regime, once Adam approaches the minimizer manifold, its long-horizon behavior within  $\tilde{\mathcal{O}}(\frac{1}{\eta^2})$  steps can be well approximated by the SDE defined in Equation (1).

## 2.2. Interpretation of The Slow SDEs for AGMs

**Adaptive Projection Operator.** Whereas Equation (2) employs a fixed projection operator  $P_\zeta$  to constrain the SDE to the manifold, the AGM slow-SDE uses an adaptive projection  $P_{\zeta, S(t)}$  that depends on the current preconditioner  $S(v(t))$ . In other words, SGD's projection is static and state-independent, but AGM's projection is state-dependent. This adaptive projection alters the way the stochastic trajectory evolves on the manifold, giving rise to a different implicit bias in AGMs versus SGD.

**Effect of the Preconditioner on the Gradient Noise Covariance.** It is well known that, near the manifold, SGD's wandering around is noise-driven. For AGMs, the situation is more subtle: First, one can show that the momentum term does not affect the implicit bias, consistent with prior theory [55]. However, the AGM trajectory is influenced by its preconditioner. Concretely, the gradient-noise covariance matrix  $\Sigma$  is filtered through the preconditioner  $S(t)$  into  $S(t)\Sigma S(t)$  and then contributes to the SDE. Over a long time horizon, this modified noise term alters the deterministic drift direction, further distinguishing AGM's dynamics from those of vanilla SGD.

### 2.3. Technical Difficulties and Proof Insights

#### 2.3.1. CONVERGENCE GUARANTEE OF AGMS

The core of our study is to consider the behavior of Adam’s implicit bias around the minimizer manifold. But before we can study this, in order to make our study meaningful, we first need to show that Adam can converge to the neighborhood of the minimizer manifold, which itself is already non-trivial. Unfortunately, Adam can not provably converge to the minimizer manifold without any constraint. In fact, the convergence issue of Adam has been debated from its birth, Reddi et al. [48] shows that Adam does not converge to the optimal solution even in some simple convex setting. Recent work [9] gives Adam’s ODE and shows that this ODE does not necessarily converge to the immobile point of the gradient flow. So We present a statement of convergence first.

**Theorem 2.2** (Convergence Bound of the AGM Framework, Stated Informally). *Let  $K = \mathcal{O}\left(\frac{1}{\eta} \log \frac{1}{\eta}\right)$ . Under mild assumptions, for any  $k < K$ ,  $\delta \in (0, 1)$ , it holds with probability at least  $1 - \delta$  that  $\mathcal{L}(\theta_k) - \mathcal{L}^* = \tilde{\mathcal{O}}(\eta)$ .*

#### 2.3.2. KEY INSIGHTS IN THE DERIVATION OF SLOW SDEs FOR AGMS

There are two obstacles preventing us from directly applying the analysis of slow SDEs from SGD to AGMs. First, the obtaining of slow SDEs requires an accurate calculation of the variation of the first-order and second-order moments of the parameters over a relatively large number of steps (a “*giant step*” in the notation of Gu et al. [19]), and in the case of SGD, due to the nature of its rotation equivariance, we can always consider its Hessian matrix as a diagonal array, as well as its corresponding minimizer manifold as a space extended by some full-space standard bases, which greatly simplifies the computation. However, it is not the case for AGMs. Due to the effect of Preconditioners  $S(\mathbf{v}_k)$ , the rotation equivariance is not satisfied here.

To resolve this, we generalize the gradient flow projection in Gu et al. [19], Li et al. [35] into a varying preconditioner flow projection. Utilizing this definition, after doing a reparameterization to the original space, we can regain the calculation simplicity in previous works [19, 35].

The second reason is that when  $\beta_2$  is too far from 1, the preconditioner moves too fast, making it very hard to characterize the change of moments. In contrast, when  $\beta_2$  is too close to 1, then the change of precondition is almost negligible, which is also not practical. To this end, we consider the case where  $1 - \beta_2 = \mathcal{O}(\eta^2)$ . And we term it “2-scheme”. The subtlety here is that this proximity does not make the change in the preconditioner negligible; rather, the change in the preconditioner affects the form of the SDE, and because the change in the preconditioner is slow enough that we can track its change.

### 3. Discussion

We show that Adam implicitly minimizes the sharpness measure  $\text{tr}(\text{Diag}(\mathbf{H})^{1/2})$ , leading to solutions and generalization behavior distinct from SGD. Our slow-SDE framework rigorously captures Adam’s adaptive semi-gradient drift near the minimizer manifold and recovers explicit separations in sparse linear regression and deep matrix factorization. Open directions include extending analysis beyond the “2-scheme” regime ( $1 - \beta_2 = \mathcal{O}(\eta^2)$ ) to intermediate regimes such as 1.5-scheme, characterizing Adam’s implicit bias once iterates exit the local manifold neighborhood, and incorporating weight-decay (e.g., AdamW) to understand its effect on the effective sharpness regularizer.

## References

- [1] Vaswani Ashish. Attention is all you need. *Advances in neural information processing systems*, 30:I, 2017.
- [2] David GT Barrett and Benoit Dherin. Implicit gradient regularization. *arXiv preprint arXiv:2009.11162*, 2020.
- [3] Guy Blanc, Neha Gupta, Gregory Valiant, and Paul Valiant. Implicit regularization for deep neural networks driven by an ornstein-uhlenbeck like process. In *Conference on learning theory*, pages 483–513. PMLR, 2020.
- [4] David Brandfonbrener and Joan Bruna. Geometric insights into the convergence of nonlinear td learning. *arXiv preprint arXiv:1905.12185*, 2019.
- [5] Matias D. Cattaneo and Boris Shigida. How memory in optimization algorithms implicitly modifies the loss, 2025. URL <https://arxiv.org/abs/2502.02132>.
- [6] Matias D. Cattaneo, Jason M. Klusowski, and Boris Shigida. On the implicit bias of adam, 2024. URL <https://arxiv.org/abs/2309.00079>.
- [7] Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in neural information processing systems*, 32, 2019.
- [8] Alex Damian, Tengyu Ma, and Jason D Lee. Label noise sgd provably prefers flat global minimizers. *Advances in Neural Information Processing Systems*, 34:27449–27461, 2021.
- [9] Steffen Dereich and Arnulf Jentzen. Convergence rates for the adam optimizer. *arXiv preprint arXiv:2407.21078*, 2024.
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [11] Aijun Du and Jinqiao Duan. Invariant manifold reduction for stochastic dynamical systems. *arXiv preprint math/0607366*, 2006.
- [12] Johannes Jisse Duistermaat and Johan AC Kolk. *Lie groups*. Springer Science & Business Media, 2012.
- [13] KJ Falconer. Differentiation of the limit mapping in a dynamical system. *Journal of the London Mathematical Society*, 2(2):356–372, 1983.
- [14] Benjamin Fehrman, Benjamin Gess, and Arnulf Jentzen. Convergence rates for the stochastic gradient descent method for non-convex objective functions. *Journal of Machine Learning Research*, 21(136):1–48, 2020.
- [15] Damir Filipović. Invariant manifolds for weak solutions to stochastic equations. *Probability theory and related fields*, 118(3):323–341, 2000.

- [16] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=6Tmlmposlrm>.
- [17] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry Vetrov, and Andrew Gordon Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns, 2018. URL <https://arxiv.org/abs/1802.10026>.
- [18] Khashayar Gatmiry, Zhiyuan Li, Tengyu Ma, Sashank Reddi, Stefanie Jegelka, and Ching-Yao Chuang. What is the inductive bias of flatness regularization? a study of deep matrix factorization models. *Advances in Neural Information Processing Systems*, 36:28040–28052, 2023.
- [19] Xinran Gu, Kaifeng Lyu, Longbo Huang, and Sanjeev Arora. Why (and when) does local sgd generalize better than sgd? *arXiv preprint arXiv:2303.01215*, 2023.
- [20] Xinran Gu, Kaifeng Lyu, Longbo Huang, and Sanjeev Arora. Why (and when) does local sgd generalize better than sgd?, 2023. URL <https://arxiv.org/abs/2303.01215>.
- [21] Xinran Gu, Kaifeng Lyu, Sanjeev Arora, Jingzhao Zhang, and Longbo Huang. A quadratic synchronization rule for distributed deep learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=yroykhkWS6>.
- [22] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 9(1):1–42, 01 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.1.1. URL <https://doi.org/10.1162/neco.1997.9.1.1>.
- [23] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [24] Stanisław Jastrzębski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in sgd. *arXiv preprint arXiv:1711.04623*, 2017.
- [25] Ziwei Ji and Matus Telgarsky. Directional convergence and alignment in deep learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 17176–17186. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/c76e4b2fa54f8506719a5c0dc14c2eb9-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/c76e4b2fa54f8506719a5c0dc14c2eb9-Paper.pdf).
- [26] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SJgIPJBFvH>.
- [27] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=HloyRlYgg>.

- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [29] Rohith Kuditipudi, Xiang Wang, Holden Lee, Yi Zhang, Zhiyuan Li, Wei Hu, Rong Ge, and Sanjeev Arora. Explaining landscape connectivity of low-cost solutions for multilayer nets. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/46a4378f835dc8040c8057beb6a2da52-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/46a4378f835dc8040c8057beb6a2da52-Paper.pdf).
- [30] Qianxiao Li, Cheng Tai, et al. Dynamics of stochastic gradient algorithms. *CoRR*, 2015.
- [31] Qianxiao Li, Cheng Tai, et al. Stochastic modified equations and adaptive stochastic gradient algorithms. In *International Conference on Machine Learning*, pages 2101–2110. PMLR, 2017.
- [32] Qianxiao Li, Cheng Tai, and Weinan E. Stochastic modified equations and dynamics of stochastic gradient algorithms i: Mathematical foundations, 2018. URL <https://arxiv.org/abs/1811.01558>.
- [33] Qianxiao Li, Cheng Tai, et al. Stochastic modified equations and dynamics of stochastic gradient algorithms i: Mathematical foundations. *Journal of Machine Learning Research*, 20(40):1–47, 2019.
- [34] Zhiyuan Li, Sathika Malladi, and Sanjeev Arora. On the validity of modeling sgd with stochastic differential equations (sdes). In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 12712–12725. Curran Associates, Inc., 2021. URL [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/69f62956429865909921fa916d61c1f8-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/69f62956429865909921fa916d61c1f8-Paper.pdf).
- [35] Zhiyuan Li, Tianhao Wang, and Sanjeev Arora. What happens after sgd reaches zero loss?—a mathematical framework. *arXiv preprint arXiv:2110.06914*, 2021.
- [36] Hong Liu, Sang Michael Xie, Zhiyuan Li, and Tengyu Ma. Same pre-training loss, better downstream: Implicit bias matters for language models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 22188–22214. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/liu23ao.html>.
- [37] Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SJeLIgBKPS>.
- [38] Kaifeng Lyu, Zhiyuan Li, and Sanjeev Arora. Understanding the generalization benefit of normalization layers: Sharpness reduction. *Advances in Neural Information Processing Systems*, 35:34689–34708, 2022.

- [39] Kaifeng Lyu, Jikai Jin, Zhiyuan Li, Simon S Du, Jason D Lee, and Wei Hu. Dichotomy of early and late phase implicit biases can provably induce grokking. In *12th International Conference on Learning Representations, ICLR 2024*, 2024.
- [40] Chao Ma, Daniel Kunin, and Lexing Ying. A quasistatic derivation of optimization algorithms’ exploration on minima manifolds, 2023. URL [https://openreview.net/forum?id=UDbNL0\\_W-3x](https://openreview.net/forum?id=UDbNL0_W-3x).
- [41] Sadhika Malladi, Kaifeng Lyu, Abhishek Panigrahi, and Sanjeev Arora. On the sdes and scaling rules for adaptive gradient algorithms. *Advances in Neural Information Processing Systems*, 35: 7697–7711, 2022.
- [42] Sadhika Malladi, Kaifeng Lyu, Abhishek Panigrahi, and Sanjeev Arora. On the sdes and scaling rules for adaptive gradient algorithms, 2024. URL <https://arxiv.org/abs/2205.10287>.
- [43] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- [44] Mor Shpigel Nacson, Suriya Gunasekar, Jason Lee, Nathan Srebro, and Daniel Soudry. Lexicographic and depth-sensitive margins in homogeneous and non-homogeneous deep models. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4683–4692. PMLR, 09–15 Jun 2019.
- [45] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.
- [46] Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- [47] Qian Qian and Xiaoyuan Qian. The implicit bias of adagrad on separable data. *Advances in Neural Information Processing Systems*, 32, 2019.
- [48] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018.
- [49] Nikunj Saunshi, Stefani Karp, Shankar Krishnan, Sobhan Miryoosefi, Sashank J. Reddi, and Sanjiv Kumar. On the inductive bias of stacking towards improving reasoning. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 71437–71464. Curran Associates, Inc., 2024. URL [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/837bc5db12f3d394d220815a7687340c-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/837bc5db12f3d394d220815a7687340c-Paper-Conference.pdf).
- [50] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [51] Noam Shazeer and Mitchell Stern. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR, 2018.

- [52] Samuel Smith, Erich Elsen, and Soham De. On the generalization benefit of noise in stochastic gradient descent. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9058–9067. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/smith20a.html>.
- [53] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19 (70):1–57, 2018.
- [54] Bohan Wang, Qi Meng, Wei Chen, and Tie-Yan Liu. The implicit bias for adaptive optimization algorithms on homogeneous neural networks. In *International Conference on Machine Learning*, pages 10849–10858. PMLR, 2021.
- [55] Runzhe Wang, Sathika Malladi, Tianhao Wang, Kaifeng Lyu, and Zhiyuan Li. The marginal value of momentum for small learning rate sgd. *arXiv preprint arXiv:2307.15196*, 2023.
- [56] Kaiyue Wen, Zhiyuan Li, Jason Wang, David Hall, Percy Liang, and Tengyu Ma. Understanding warmup-stable-decay learning rates: A river valley loss landscape perspective, 2024. URL <https://arxiv.org/abs/2410.05192>.
- [57] Blake Woodworth, Suriya Gunasekar, Jason D Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*, pages 3635–3673. PMLR, 2020.
- [58] Shuo Xie and Zhiyuan Li. Implicit bias of AdamW:  $\ell_\infty$ -norm constrained optimization. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 54488–54510. PMLR, 21–27 Jul 2024.
- [59] Chenyang Zhang, Difan Zou, and Yuan Cao. The implicit bias of adam on separable data, 2024. URL <https://arxiv.org/abs/2406.10650>.
- [60] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy8gdB9xx>.
- [61] Yushun Zhang, Congliang Chen, Tian Ding, Ziniu Li, Ruoyu Sun, and Zhi-Quan Luo. Why transformers need adam: A hessian perspective, 2024. URL <https://arxiv.org/abs/2402.16788>.

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Theoretical Analysis of Adam</b>	<b>2</b>
2.1	Slow SDE Analysis for AGMs . . . . .	3
2.2	Interpretation of The Slow SDEs for AGMs . . . . .	4
2.3	Technical Difficulties and Proof Insights . . . . .	5
2.3.1	Convergence Guarantee of AGMs . . . . .	5
2.3.2	Key Insights in the Derivation of Slow SDEs for AGMs . . . . .	5
<b>3</b>	<b>Discussion</b>	<b>5</b>
<b>A</b>	<b>Preliminaries</b>	<b>12</b>
<b>B</b>	<b>Related Work</b>	<b>14</b>
<b>C</b>	<b>Illustration of the Difference between Conventional SDE and Slow SDE</b>	<b>15</b>
<b>D</b>	<b>Adam’s Provable Generalization Benefit with Label Noise</b>	<b>15</b>
D.1	Reduction of Adam’s Implicit Regularizer with Label Noise . . . . .	15
D.2	Problem Setup: the Diagonal Net . . . . .	16
D.3	Result: Adam’s Implicit Regularizer Facilitates Sparse Ground-truth Recovery . . .	17
<b>E</b>	<b>Matrix Factorization: Adam Implicitly Regularizes Sharpness Differently</b>	<b>17</b>
E.1	Problem setup . . . . .	17
E.2	Results . . . . .	18
<b>F</b>	<b>Formal Statements of the Main Results</b>	<b>18</b>
<b>G</b>	<b>Constructing the Working Zones</b>	<b>20</b>
<b>H</b>	<b>Proof of the Convergence of AGMs</b>	<b>22</b>
<b>I</b>	<b>Proof of the SDE Approximation of AGMs</b>	<b>28</b>
I.1	Lemmas for Adaptive Manifold Projection . . . . .	29
I.2	Iteration Stays Near Manifold . . . . .	30
I.3	Moment Calculation of AGMs Near Manifold . . . . .	33
I.3.1	Moment Calculation Within a Giant Step . . . . .	36
I.4	Weak Approximation . . . . .	42
I.5	Preliminary and Additional Notations . . . . .	43
I.6	Proof of the Approximation for Slow SDE of AGMs . . . . .	44
<b>J</b>	<b>Proof of Theorems in <a href="#">Appendix D</a></b>	<b>47</b>
J.1	Proof of Adam and AdamE- $\lambda$ ’s Implicit Bias with Label Noise . . . . .	47
J.2	Proof of <a href="#">Lemma D.1</a> . . . . .	50

Table 1: Examples of  $V, S$  functions for some optimizers in the AGM Framework.

Optimizer	Function $V$	Function $S$	Remarks
Adam	$V(\mathbf{M}) = \text{diag}(\mathbf{M})$	$S(\mathbf{v}) = \text{Diag}(1/(\sqrt{\mathbf{v}} + \epsilon))$	
Adam-mini	$V(\mathbf{M})_i = \frac{1}{ B(i) } \sum_{j \in B(i)} M_{jj}$	$S(\mathbf{v}) = \text{Diag}(1/(\sqrt{\mathbf{v}} + \epsilon))$	Parameters partitioned; $i$ belongs to block $B(i)$ . $i$ belongs to layer $L(i)$ in the model.
Adalayer	$V(\mathbf{M})_i = \frac{1}{ L(i) } \sum_{j \in L(i)} M_{jj}$	$S(\mathbf{v}) = \text{Diag}(1/(\sqrt{\mathbf{v}} + \epsilon))$	
AdamE- $\lambda$	$V(\mathbf{M}) = \text{diag}(\mathbf{M})$	$S(\mathbf{v}) = \text{Diag}(1/(\mathbf{v}^{\odot \lambda} + \epsilon))$	

## Appendix A. Preliminaries

**Notations.** Unless otherwise stated, for a square matrix  $\mathbf{M}$ ,  $\text{diag}(\mathbf{M})$  denotes the vector consisting of its diagonal entries. The notation  $\text{Diag}$  has two usages: For a vector  $\mathbf{v}$ ,  $\text{Diag}(\mathbf{v})$  denotes the diagonal matrix with  $\mathbf{v}$  on its diagonal; and for a square matrix  $\mathbf{M}$ ,  $\text{Diag}(\mathbf{M})$  denotes the diagonal matrix that only keeps  $\mathbf{M}$ 's diagonal entries and equals 0 elsewhere, i.e.  $\text{Diag}(\mathbf{M}) \stackrel{\text{def}}{=} \text{Diag}(\text{diag}(\mathbf{M}))$ . For two vectors  $\mathbf{u}, \mathbf{v}$  with the same dimension  $d$ ,  $\mathbf{u} \odot \mathbf{v}$  denotes element-wise multiplication  $(u_1 v_1, \dots, u_d v_d)$ . For any exponent  $p$ ,  $\mathbf{v}^{\odot p}$  denotes element-wise exponentiation, i.e.  $\mathbf{v}^{\odot p} = (v_1^p, \dots, v_d^p)$ , and  $\sqrt{\mathbf{v}}$  means  $\mathbf{v}^{\odot 1/2}$ .

For a mapping  $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ , we denote the Jacobian with respect to  $\boldsymbol{\theta} \in \mathbb{R}^d$  as  $\partial F(\boldsymbol{\theta}) \in \mathbb{R}^{d \times d}$ , and  $\partial^2 F(\boldsymbol{\theta})$  the second-order derivative at  $\boldsymbol{\theta}$ , which is a third-order tensor. Given a matrix  $\mathbf{M} \in \mathbb{R}^{d \times d}$ , we use the notation  $\partial^2 F(\boldsymbol{\theta})[\mathbf{M}]$  to denote the second-order directional derivative of  $F$  at  $\boldsymbol{\theta}$  in the direction  $\mathbf{M}$ , defined as  $\partial^2 F(\boldsymbol{\theta})[\mathbf{M}] := \sum_{i \in [d]} \left\langle \frac{\partial^2 F_i}{\partial \boldsymbol{\theta}^2}, \mathbf{M} \right\rangle \mathbf{e}_i$ , where  $F_i$  represents the  $i$ -th element in  $F$ , and  $\mathbf{e}_i$  is the  $i$ -th vector of the standard basis. When the context is clear, we write  $\partial^2(\nabla \mathcal{L})(\boldsymbol{\theta})[\mathbf{M}]$  as  $\nabla^3 \mathcal{L}(\boldsymbol{\theta})[\mathbf{M}]$  for brevity.

**Loss Functions.** Define  $\ell(\boldsymbol{\theta}; \xi)$  as the loss function for a data sample  $\xi$  for a model with parameters  $\boldsymbol{\theta}$ . Define  $\mathcal{L}(\boldsymbol{\theta}) := \mathbb{E}_{\xi \sim \mathcal{S}}[\ell(\boldsymbol{\theta}; \xi)]$  as the training loss function, where  $\mathcal{S}$  is the training dataset and  $\xi \sim \mathcal{S}$  means the data sample  $\xi$  is drawn from  $\mathcal{S}$  uniformly at random. Let  $\mathcal{L}^* := \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathcal{L}(\boldsymbol{\theta})$  be the minimum of training loss. Let  $\mathcal{Z}(\boldsymbol{\theta})$  be the distribution of gradient noise  $\nabla \ell(\boldsymbol{\theta}; \xi) - \nabla \mathcal{L}(\boldsymbol{\theta})$ , which is a random variable that depends on  $\boldsymbol{\theta}$ . We define  $\Sigma(\boldsymbol{\theta}) := \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}(\boldsymbol{\theta})}[\mathbf{z} \mathbf{z}^\top]$  as the noise covariance matrix of gradients at  $\boldsymbol{\theta}$ .

**SGD and Adam.** SGD is an iterative method that starts from an initial point  $\boldsymbol{\theta}_0$  and updates the parameters as  $\boldsymbol{\theta}_{k+1} := \boldsymbol{\theta}_k - \eta \nabla \ell_k(\boldsymbol{\theta}_k)$  for all  $k \geq 0$ , where  $\eta$  is the learning rate,  $\ell_k(\boldsymbol{\theta})$  is the loss function for the data sample  $\xi_k$  sampled at step  $k$ . Adam [28] is a popular optimizer that updates the parameters as:

$$\begin{aligned}
 \mathbf{m}_{k+1} &:= \beta_1 \mathbf{m}_k + (1 - \beta_1) \nabla \ell_k(\boldsymbol{\theta}_k) \\
 \mathbf{v}_{k+1} &:= \beta_2 \mathbf{v}_k + (1 - \beta_2) \nabla \ell_k(\boldsymbol{\theta}_k)^{\odot 2} \\
 \theta_{k+1,i} &:= \theta_{k,i} - \eta \frac{m_{k+1,i}}{\sqrt{v_{k+1,i}} + \epsilon} \quad \text{for all } i \in [d].
 \end{aligned}$$

Note that in practice, it is common to normalize  $\mathbf{m}_{k+1}$  and  $\mathbf{v}_{k+1}$  by  $1 - \beta_1^{k+1}$  and  $1 - \beta_2^{k+1}$  respectively before the division. However, this normalization quickly becomes neglectable when  $k$  is large, so we ignore it for simplicity.

**SDE First-Order Approximation For SGD.** A *stochastic differential equation* (SDE) is an extension of an ordinary differential equation that incorporates random perturbations, and is widely used to model systems under the influence of noise. An SDE on  $\mathbb{R}^d$  takes the form  $d\theta_t = b(\theta_t)dt + \sigma(\theta_t)dW_t$  where  $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is the drift vector field,  $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times m}$  is the diffusion matrix, and  $\{W_t\}_{t \geq 0}$  is an  $m$ -dimensional Wiener process. A line of works [24, 30, 31, 33, 34, 52] used the following SDE to serve as a first-order approximation of SGD, which we refer to as the *conventional SDE*:

$$d\theta_t = -\nabla \mathcal{L}(\theta_t)dt + \sqrt{\eta} \Sigma^{1/2}(\theta_t)dW_t,$$

where the stochastic integral is taken in the Itô sense. For an introduction to Itô calculus, see Oksendal [46]. Later, Malladi et al. [42] extended this type of SDE to Adam. Besides these conventional SDEs, below we introduce another type of SDE, slow SDE, that can more explicitly capture the implicit bias of SGD near a manifold of minimizers.

**Manifold Assumption.** Before going into the slow SDE, we introduce the *manifold assumption*. Previous studies [17, 29] have found that low-loss solutions are in fact connected to each other, a phenomenon known as mode connectivity. Wen et al. [56] provided empirical evidence that the training dynamics of language model training usually happen in a structure similar to a river valley, where many low-loss solutions lie in the bottom of the valley. Motivated by these observations, many previous works [14, 19, 35, 37] assumed that the minimizers of the training loss function are not isolated points but connected and form a manifold  $\Gamma$ :

**Assumption A.1.**  $\Gamma$  is  $\mathcal{C}^\infty$ -smooth,  $(d - m)$ -dimensional submanifold of  $\mathbb{R}^d$ , where any  $\zeta \in \Gamma$  is a local minimizer of  $\mathcal{L}$ . For all  $\zeta \in \Gamma$ ,  $\text{rank}(\nabla^2 \mathcal{L}(\zeta)) = m$ . Additionally, there exists an open neighborhood of  $\Gamma$ , denoted as  $U$ , such that  $\Gamma = \arg \min_{\theta \in U} \mathcal{L}(\theta)$ .

With this assumption, if an optimization process converges and the learning rate  $\eta$  is sufficiently small, then the process will be trapped near some minimizer manifold which we denote by  $\Gamma$ .

**Slow SDE.** A line of works [3, 8, 35] studied the dynamics of SGD near the manifold  $\Gamma$  and showed that SGD has an implicit bias towards flatter minimizers on  $\Gamma$ . This effect cannot be directly seen from conventional SDEs, so Li et al. [35] derived a new type of SDE approximation, called slow SDE, that can explicitly capture this effect. See Appendix C for an illustration of the difference between conventional SDEs and slow SDEs. Here we introduce the slow SDE for SGD following the formulation in Gu et al. [21]. For ease of presentation, we define the following projection operators  $\Phi, P_\zeta$  for points and differential forms respectively. Consider the gradient flow  $\frac{dx(t)}{dt} = -\nabla \mathcal{L}(x(t))$  with  $x(0) = x$ , and fix some point  $\theta_{\text{null}} \notin \Gamma$ , we define the gradient flow projection of any  $x$ ,  $\Phi(x)$ , as  $\lim_{t \rightarrow +\infty} x(t)$  if the limit exists and belongs to  $\Gamma$ , and  $\theta_{\text{null}}$  otherwise. It can be shown by simple calculus [35] that  $\partial\Phi(\zeta)$  equals the projection matrix onto the tangent space of  $\Gamma$  at  $\zeta$ . We decompose the noise covariance  $\Sigma(\zeta)$  for  $\zeta \in \Gamma$  into two parts: the noise in the tangent space  $\Sigma_{\parallel}(\zeta) := \partial\Phi(\zeta)\Sigma(\zeta)\partial\Phi(\zeta)$  and the noise in the rest  $\Sigma_{\diamond}(\zeta) := \Sigma(\zeta) - \Sigma_{\parallel}(\zeta)$ .

For any  $\zeta \in \Gamma$ , matrix  $A$  and vector  $b$ , we use  $P_\zeta(A dW_t + b dt)$  to denote  $\Phi(\zeta + A dW_t + b dt) - \Phi(\zeta)$ , which equals  $\partial\Phi(\zeta)A dW_t + (\partial\Phi(\zeta)b + \frac{1}{2}\partial^2\Phi(\zeta)[AA^\top]) dt$  by Itô calculus.  $P_\zeta$  can be interpreted as projecting an infinitesimal step from  $\zeta$ , so that  $\zeta$  after taking the projected step does not leave the manifold  $\Gamma$ . Now we are ready to state the SDE for Local SGD.

**Definition A.1** (Slow SDE for SGD). Given  $\eta > 0$  and  $\zeta_0 \in \Gamma$ , define  $\zeta(t)$  as the solution of the following SDE with initial condition  $\zeta(0) = \zeta_0$ :

$$d\zeta(t) = \underbrace{P_\zeta\left(\Sigma_\parallel^{1/2}(\zeta)d\mathbf{W}_t\right)}_{(a) \text{ diffusion}} - \underbrace{\frac{1}{2}\nabla^3\mathcal{L}(\zeta)[\widehat{\Sigma}_\diamond(\zeta)]dt}_{(b) \text{ drift}}. \quad (2)$$

Here  $\widehat{\Sigma}_\diamond(\zeta)$  is defined as  $\sum_{i,j: \lambda_i \neq 0 \vee \lambda_j \neq 0} \frac{1}{\lambda_i + \lambda_j} \langle \Sigma_\diamond(\zeta), \mathbf{v}_i \mathbf{v}_j^\top \rangle \mathbf{v}_i \mathbf{v}_j^\top$ , where  $\{\mathbf{v}_i\}_{i=1}^d$  is an orthonormal eigenbasis of  $\nabla^2\mathcal{L}(\zeta)$  with corresponding eigenvalues  $\lambda_1, \dots, \lambda_d$ .

**Interpretation of the Slow SDE for SGD: Semi-gradient Descent** This SDE on the minimizer manifold  $\Gamma$  splits naturally into a *diffusion* term  $P_\zeta(\Sigma_\parallel^{1/2}(\zeta)d\mathbf{W}_t)$  injecting noise in the tangent space, and a *drift* term  $-\frac{1}{2}P_\zeta(\nabla^3\mathcal{L}(\zeta)[\widehat{\Sigma}_\diamond(\zeta)]dt)$  that can be seen as the negative *semi-gradient* of the following sharpness measure:

$$\mu(\zeta) := \left\langle \nabla^2\mathcal{L}(\zeta), \widehat{\Sigma}_\diamond(\zeta) \right\rangle.$$

Here we use the word “semi-gradient” [4, 43] because it is not exactly the gradient of  $\mu(\zeta)$  but only the gradient with respect to the first argument of the inner product. More specifically, define  $\mu(\zeta_1, \zeta_2) := \left\langle \nabla^2\mathcal{L}(\zeta_1), \widehat{\Sigma}_\diamond(\zeta_2) \right\rangle$ , then the drift term is essentially  $-\frac{1}{2}\nabla_{\zeta_1}\mu(\zeta_1, \zeta_2)|_{\zeta_1=\zeta, \zeta_2=\zeta}$  after projecting onto the tangent space of  $\Gamma$  at  $\zeta$ . In other words, SGD near manifold takes semi-gradients to minimize the implicit regularizer  $\langle \nabla^2\mathcal{L}(\zeta), \widehat{\Sigma}_\diamond(\zeta) \rangle$  but pretend  $\widehat{\Sigma}_\diamond(\zeta)$  to be fixed, i.e. ignore the dependency of  $\widehat{\Sigma}_\diamond(\zeta)$  on  $\zeta$ .

**Example: Noisy Ellipse.** We provide a toy example to illustrate the phenomenon described by the slow SDE for SGD: there are two parameters  $x, y$  and an elliptical loss with label noise  $\mathcal{L}(x, y) = \frac{1}{2} \left( \frac{(x+y)^2}{2a^2} + \frac{(y-x)^2}{2b^2} - 1 - \delta \right)^2$ . The label noise  $\delta$  is sampled uniformly from  $\{-0.5, 0.5\}$  at every step. As depicted in Fig. 4, SGD moves towards flatter minimizers after reaching the manifold. The same phenomenon can be observed for Adam, but Adam converges to a different minimizer that is closer to the axis (or, “sparser” in the parameter space). Understanding the difference between SGD and Adam is the main focus of this paper.

## Appendix B. Related Work

**Implicit Bias of SGD.** Parallel work on *implicit gradient regularization* (IGR) derives higher-order terms for full-batch GD [2] and extends to Adam [5, 6]. While Cattaneo et al. [6] argued that Adam anti-regularizes sharpness when  $\beta_1 < \beta_2$ , our  $\mathcal{O}(\eta^{-2})$ -time SDE analysis shows Adam still regularizes sharpness under these settings, overturning their conclusion.

**Implicit Bias of Adam.** Despite Adam’s widespread use, its implicit bias remains underexplored. Qian and Qian [47] and Xie and Li [58] analyzed AdaGrad and AdamW, but these techniques do not apply directly to Adam. Wang et al. [54] showed Adam’s regularizer matches SGD’s under restrictive gradient-magnitude assumptions, and Zhang et al. [59] treated only linearly separable data, limiting practical relevance.

**Slow SDE Approximation.** To capture long-term behavior, we adopt the *slow SDE* technique of Li et al. [35] and Gu et al. [20]. Standard SDE approximations (6, 32, 34, 42) focus on the  $\tilde{\mathcal{O}}(\eta^{-1})$  convergence phase and fail on the manifold. In contrast, slow SDEs peel off convergence to track the  $\mathcal{O}(\eta^{-2})$  manifold dynamics accurately.

### Appendix C. Illustration of the Difference between Conventional SDE and Slow SDE

In this section, we illustrate the difference between conventional SDE and slow SDE. In Fig. 8, let  $\Gamma$  denotes a 1D manifold, then the discrete iteration of the optimization process can be seen as successive steps (orange, Fig. 5) that starts from  $A$ , first converge to some point  $B$  in  $\Gamma$  and then move along  $\Gamma$  to  $C$ .

The main intuition behind slow SDE is that the whole process  $A \rightarrow B \rightarrow C$  can actually be decomposed into two motions: a convergence motion  $A \rightarrow H$  (dashed, Fig. 7) and an implicit regularization motion  $H \rightarrow B \rightarrow C$ . The convergence motion is fast and dominates the dynamics during the convergence phase, but it fades out as soon as convergence phase ends; meanwhile the slow, implicit regularization motion starts to dominate.

The conventional SDE approximates the convergence phase only, whose one unit time corresponds to  $\tilde{O}(\eta^{-1})$  steps (Fig. 6). In contrast, slow SDE manages to separate the slow implicit regularization motion from the fast convergence, and approximate the implicit regularization only (Fig. 7).

**Remark.** The projection methods (which project  $A \rightarrow B \rightarrow C$  to  $H \rightarrow B \rightarrow C$ ) vary in the analysis of different optimizers. Intuitively, the projection should reflect the converging direction driven by a clean (without noise) and continuous version of the optimizer. In SGD the projection is gradient flow; but in Adam we need to consider the pre-conditioning effect caused by  $1/\sqrt{v} + \epsilon$ , so we add an SDE to track the pre-conditioner, and define a pre-conditioned gradient flow for projection.

### Appendix D. Adam’s Provable Generalization Benefit with Label Noise

In this section, we will prove that under label noise setting, the implicit regularizer of Adam reduces to a simpler form that aligns better with sparsity regularizations, and then verify experimentally.

#### D.1. Reduction of Adam’s Implicit Regularizer with Label Noise

On an  $\ell_2$ -regression task on dataset  $\{z_i, y_i\}_{i=1}^n$ , adding *label noise* means adding a noise sampled i.i.d. from  $\{\pm\delta\}$  to any true label  $y$  before feeding forward to the network. A crucial property of the label noise setting is that when  $\theta \in \Gamma$ ,  $\Sigma \equiv \alpha \nabla^2 \mathcal{L}$  for some constant  $\alpha$  [3], which simplifies the setting and has been largely used [3, 8, 19, 35] to analyze the implicit bias of SGD and other optimizers.

**Theorem D.1** (Adam’s Implicit Bias with Label Noise, Stated Informally). *Adam’s SDE becomes an ODE under the label noise setting, and when  $\epsilon$  is small, the fixed point of this ODE must satisfy  $\nabla \text{tr}(\text{Diag}(\mathbf{H})^{1/2}) = 0$ .*

**Proof Sketch.** Under label noise setting, SDE Eq. (1) will be greatly simplified. In fact, the diffusion term of the slow SDE would equal zero, and the drift term could be simplified to

$$\begin{cases} dv(t) = c(V(\Sigma(\zeta)) - v) dt, \\ d\zeta(t) = -\frac{\alpha}{2} S(v) \nabla^3 \mathcal{L}(\zeta) [S(v)] dt, \end{cases} \quad (3)$$

the proof of which is deferred to Appendix. With our SDE becoming an ODE, we consider the fixed point of this ODE, which should satisfy  $v = V(\Sigma(\zeta))$  and  $\nabla^3 \mathcal{L}(\zeta) [S(v)] = 0$  since  $S(v)$  is invertible. Denote  $\mathbf{H} = \nabla^2 \mathcal{L}(\zeta) = \Sigma(\zeta)/\alpha$ . In the case of Adam,  $v = \text{diag}(\Sigma) = \alpha \cdot \text{diag}(\mathbf{H})$ , and  $S(v) = \text{Diag}(1/(\sqrt{v} + \epsilon))$ . Then we integrate by parts and obtain  $\nabla^3 \mathcal{L}(\zeta) [S(v)] = \nabla [\langle \mathbf{H}, S(v) \rangle] - \nabla (S(v)) [\mathbf{H}]$ . A straightforward simplification gives the result.  $\square$

The proof of this theorem also inspires us of a simple way to directly adjust the implicit bias of Adam. Specifically, for any  $\lambda \in [0, 1]$ , we define *AdamE- $\lambda$*  as an optimizer identical with Adam, except that  $S(\mathbf{v}) = \text{Diag}(1/(\mathbf{v}^{\odot \lambda} + \epsilon))$ . Obviously AdamE- $\frac{1}{2}$  reduces to Adam and all AdamE- $\lambda$ 's belong to the AGM framework. To compute the implicit bias of AdamE- $\lambda$  with label noise, we can apply the same method as in [Theorem D.1](#), and the result is stated below.

**Theorem D.2** (AdamE- $\lambda$ 's Implicit Bias with Label Noise, Stated Informally). *For  $\lambda \in [0, 1]$ , AdamE- $\lambda$ 's SDE becomes an ODE under the label noise setting, and when  $\epsilon$  is small, the fixed point of this ODE must satisfy  $\nabla \text{tr}(\text{Diag}(\mathbf{H})^{1-\lambda}) = 0$ .*

[Theorem D.2](#) indicates that tuning the exponent of the second-order moment in adam will exactly result in tuning the exponent of  $\text{diag}(\nabla^2 \mathcal{L}(\zeta))$  in the implicit bias. When  $\lambda = 0$ , the implicit bias reduces to that of SGD, and AdamE also gets rid of the effect of second-order moments and reduces to SGD with momentum, which coincides perfectly. Next, we will relate the implicit bias with sparsity and compare the performance of Adam, AdamE and SGD in a simple experimental setup.

## D.2. Problem Setup: the Diagonal Net

In this section, we adopt the *diagonal linear network* (diagonal net) setting proposed by Woodworth et al. [\[57\]](#) as an experimental setting, which is also used by Li et al. [\[35\]](#) to study the implicit bias of SGD.

**Setting (Diagonal Net with Label Noise):** Let  $\mathbf{w}^* \in \mathbb{R}^d$  be an unknown  $\kappa$ -sparse ground truth vector. Let  $\{(z_i, y_i)\}_{i \in [n]}$  be the training dataset where each  $z_i \stackrel{\text{i.i.d.}}{\sim} \text{Unif}\{\pm 1\}^d$ , and each  $y_i$  is generated by  $\langle z_i, \mathbf{w}^* \rangle$ . Our parameter is defined as  $\boldsymbol{\theta} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} \in \mathbb{R}^{2d}$ . For any function  $g$  defined on  $\mathbb{R}^{2d}$ , we write  $g(\boldsymbol{\theta})$  and  $g(\mathbf{u}, \mathbf{v})$  exchangeably. The loss function is defined as:

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\boldsymbol{\theta}), \quad \text{where } \mathcal{L}_i(\boldsymbol{\theta}) = \frac{1}{2} (\langle z_i, \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2} \rangle - y_i)^2$$

where a label noise is added to the true label  $y$  during training. This setting can be viewed as using estimation  $\hat{\mathbf{w}} = \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}$  to approximate the ground truth vector  $\mathbf{w}^*$  of a linear regression task. Note that  $d \gg n$  here so the model is highly overparameterized: Theoretically, Li et al. [\[35\]](#) proved that  $n = \mathcal{O}(\kappa \ln d)$  is enough for SGD to recover ground truth, and we will later show experimentally that less than 1000 training pairs is required for both Adam and SGD to achieve a low test loss when  $d = 10000$ . The manifold is defined as wherever zero train loss is achieved, i.e.  $\Gamma = \{\boldsymbol{\theta} | \langle z_i, \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2} \rangle = y_i, \forall i \in [n]\}$ .

This setting allows us to relate the implicit bias directly to the sparsity of the output. It's straightforward to verify that  $\nabla^2 \mathcal{L}(\boldsymbol{\theta}) = \frac{4}{n} \sum_{i=1}^n \begin{pmatrix} z_i \odot \mathbf{u} \\ -z_i \odot \mathbf{v} \end{pmatrix} \begin{pmatrix} z_i \odot \mathbf{u} \\ -z_i \odot \mathbf{v} \end{pmatrix}^\top$  when  $\boldsymbol{\theta} \in \Gamma$ , so  $\text{diag}(\nabla^2 \mathcal{L}(\boldsymbol{\theta})) = 4\boldsymbol{\theta}^{\odot 2}$ . Then note the following property:

**Lemma D.1.** *Let some optimum  $\boldsymbol{\theta}$  satisfy that  $\boldsymbol{\theta} \in \arg \min_{\boldsymbol{\theta}' \in \Gamma} \text{tr}(\text{Diag}(\mathbf{H})^{1/2})$ , then we also have  $\boldsymbol{\theta} \in \arg \min_{\boldsymbol{\theta}' \in \Gamma} \|\mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}\|_{1/2}$ . Similarly, for any  $e_0 \in (0, 1]$ , if  $\boldsymbol{\theta} \in \arg \min_{\boldsymbol{\theta}' \in \Gamma} \text{tr}(\text{Diag}(\mathbf{H})^{e_0})$ , then  $\boldsymbol{\theta} \in \arg \min_{\boldsymbol{\theta}' \in \Gamma} \|\mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}\|_{e_0}$ .*

This is because only  $\hat{\mathbf{w}} = \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}$  matters in the evaluation of train loss, so if  $u_i \neq 0$  and  $v_i \neq 0$  for some  $i$ , then we can decrease the absolute value of both  $u_i$  and  $v_i$  while keeping  $u_i^2 - v_i^2$  unchanged, and we will get another optimum with smaller  $\text{tr}(\text{Diag}(\mathbf{H})^{e_0})$ . Thus  $u_i = 0$  or  $v_i = 0$  for any  $i$ . When this holds, we have  $\text{tr}(\text{Diag}(\mathbf{H})^{e_0}) = 4^{e_0} \|\boldsymbol{\theta}^{\odot 2e_0}\|_1 = (4 \|\hat{\mathbf{w}}\|_{e_0})^{e_0}$ . Therefore,

implicitly regularizing  $\text{tr}(\text{Diag}(\mathbf{H})^{e_0})$  can be viewed as regularizing the  $\ell_{e_0}$  norm of the output: Adam regularizes  $\ell_{0.5}$ , SGD regularizes  $\ell_1$ , and AdamE- $\lambda$  regularizes  $\ell_{1-\lambda}$  norm of the output. Just as lasso ( $\ell_1$ ) regression’s advantage over ridge ( $\ell_2$ ) regression in sparse ground truth recovery, we argue that Adam and AdamE with large  $\lambda$ ’s will recover ground truth more efficiently than SGD, and AdamE with small  $\lambda$ ’s on this task.

### D.3. Result: Adam’s Implicit Regularizer Facilitates Sparse Ground-truth Recovery

We plot the results of the experiment in Fig. 11. We gradually increase the number of training points, and train Adam, SGD and AdamE with different configurations until convergence. We identify a training configuration as ‘recovered the groundtruth’ when the test loss ends up below 1. As depicted in Fig. 9, Adam’s test loss plunges towards zero around  $n_{\text{train}} = 420$ , while SGD’s test loss decreases gradually with the increase of training data. As an attempt to interpolate between different implicit biases, we also train AdamE with different  $\lambda$ ’s. Fig. 10 shows that AdamE-0.001’s performance is similar to that of SGD, and all AdamE with larger  $\lambda$ ’s exhibit the same sudden recovery behavior as Adam.

**Takeaway.** In the diagonal net setting, Adam’s unique implicit bias aligns better with the fundamental target of reducing the sparsity of the model’s output, which facilitates the recovery of the sparse ground truth compared to SGD, and this improvement mainly arises from the fact that Adam takes the second order moment into consideration. Starting from SGD, even if we introduce the second-order moment in the preconditioner for a little bit, it could result in significant assistance in sparse ground truth recovery.

## Appendix E. Matrix Factorization: Adam Implicitly Regularizes Sharpness Differently

The diagonal-net experiments in Appendix D showed that Adam’s implicit bias towards *sparsity* improves generalization relative to SGD. We now turn to supply the potentially negative impact of Adam’s implicit bias in another controlled setting: **deep matrix factorization with label noise**, where the relevant implicit regularizers are analytically tractable. In this task, Adam is expected to minimize  $\text{tr}(\text{Diag}(\mathbf{H})^{1/2})$  rather than  $\text{tr}(\mathbf{H})$ . Leveraging existing theory, we therefore predict that (i) Adam will converge to a solution with  $\text{tr}(\mathbf{H})$  larger—but  $\text{tr}(\text{Diag}(\mathbf{H})^{1/2})$  smaller—than SGD’s solution, and (ii) once training reaches the interpolation regime, Adam will *generalize worse* than vanilla SGD in the presence of label noise. Our experiments confirm both predictions (Figure 12).

### E.1. Problem setup

Consider an  $L$ -layer linear network with parameters  $\mathbf{W} = (\mathbf{W}_1, \dots, \mathbf{W}_L)$ , where  $\mathbf{W}_i \in \mathbb{R}^{d_i \times d_{i-1}}$  and  $d_i \geq \min\{d_0, d_L\}$  for all  $i$ . Let  $\mathbf{M}^* \in \mathbb{R}^{d_L \times d_0}$  be a rank- $r$  ground-truth matrix, and observe  $n$  i.i.d. linear measurements  $\{(\mathbf{A}_i, b_i)\}_{i=1}^n$  generated by  $b_i = \langle \mathbf{A}_i, \mathbf{M}^* \rangle$ . With label noise and mini-batch size  $B$  the empirical loss at step  $t$  is

$$\mathcal{L}_t(\mathbf{W}) = \frac{1}{B} \sum_{i \in \mathcal{B}_t} (\langle \mathbf{A}_i, \mathbf{W}_L \cdots \mathbf{W}_1 \rangle - b_i + \xi_{t,i})^2,$$

where  $\mathcal{B}_t$  is a fresh batch of size  $B$ , and  $\xi_{t,i} \sim \mathcal{N}(0, \sigma^2)$  are independent across  $(t, i)$ .

**Implicit regularization.** With small learning rates and additive label noise, SGD asymptotically minimizes  $\text{tr}(\mathbf{H})$  once it reaches the zero-loss manifold. In matrix factorization, minimizing  $\text{tr}(\mathbf{H})$  is nearly equivalent to minimizing the nuclear norm of the recovered matrix [18], which promotes low rank and hence better generalization when  $\mathbf{M}^*$  is low rank. Adam, however, implicitly minimizes  $\text{tr}(\text{Diag}(\mathbf{H})^{1/2})$ ; it therefore converges to a different point, typically with larger  $\text{tr}(\mathbf{H})$  and reduced generalization.

## E.2. Results

Our SGD setup follows Section 7 of Gatmiry et al. [18]. For Adam, we use the standard hyperparameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and learning rate  $10^{-3}$ ; all other settings are identical to SGD.

Figure 12 (top row) shows the evolution of curvature metrics. Adam drives  $\text{tr}(\text{Diag}(\mathbf{H})^{1/2})$  sharply downward while  $\text{tr}(\mathbf{H})$  remains high and even non-monotone, confirming that Adam does *not* target overall Hessian trace. Correspondingly, the bottom row shows that Adam attains a higher test MSE despite identical training error—evidence that its implicit bias is detrimental in this setting.

**Takeaway.** In deep matrix factorization with label noise, Adam’s preference for minimizing the diagonal curvature leads it to sharper—and less generalizable—solutions than SGD, reinforcing that Adam’s implicit regularization differs qualitatively from SGD’s and can hurt performance when overall curvature matters.

## Appendix F. Formal Statements of the Main Results

In this section, we give the formal versions of the main results stated in Section 2, where we presented the two principal theorems:

1. The AGM iterates converge to a neighborhood of the manifold (Theorem 2.2);
2. Moreover, once the iterates enter this neighborhood, their dynamics over  $\mathcal{O}(\eta^{-2})$  discrete steps can be accurately tracked by a slow SDE (Theorem 2.1).

Recall that in the AGM framework, the transition from  $\boldsymbol{\theta}_k$  to  $\boldsymbol{\theta}_{k+1}$  is defined as:

$$\begin{aligned}\mathbf{m}_{k+1} &:= \beta_1 \mathbf{m}_k + (1 - \beta_1) \nabla \ell_k(\boldsymbol{\theta}_k) \\ \mathbf{v}_{k+1} &:= \beta_2 \mathbf{v}_k + (1 - \beta_2) V \left( \nabla \ell_k(\boldsymbol{\theta}_k) \nabla \ell_k(\boldsymbol{\theta}_k)^\top \right) \\ \boldsymbol{\theta}_{k+1} &:= \boldsymbol{\theta}_k - \eta S(\mathbf{v}_{k+1}) \mathbf{m}_{k+1}.\end{aligned}$$

Before formalizing these two results, we introduce some technical assumptions first.

**Assumption F.1.**  $\mathcal{L}$  is  $\mathcal{C}^5$ -smooth on  $\mathbb{R}^d$  and is bounded from below, i.e.  $\mathcal{L}^* = \inf_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}) > -\infty$ .

**Assumption F.2.**  $\mathcal{L}$  is  $\rho$ -smooth, i.e.  $\forall \boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^d$ :

$$\|\nabla \mathcal{L}(\boldsymbol{\theta}_1) - \nabla \mathcal{L}(\boldsymbol{\theta}_2)\|_2 \leq \rho \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2$$

**Assumption F.3.** We assume that all  $\nabla \mathcal{L}(\boldsymbol{\theta})$  and  $\|z_k\|_2$  are L2-bounded; Specifically, there exists some constant  $R$  s.t.  $\forall \boldsymbol{\theta}$  and  $\forall k \in \mathbb{N}$ ,  $\|\nabla \mathcal{L}(\boldsymbol{\theta})\|_2 \leq R$  and  $\|\nabla \ell_k(\boldsymbol{\theta})\|_2 \leq R$ .

Since the gradients and gradient noises are assumed to be bounded and  $S$  and  $V$  both satisfy Lipschitz-ness, we reasonably assume that all outputs from  $S$ ,  $S^{-1}$  and  $V$  are bounded. Specifically, denote  $R_0, R_1$  as constants satisfying  $S^{-1}(\mathbf{v}) \preceq R_0$  and  $\mathbf{v} \preceq R_1$  (i.e. each entry of  $\mathbf{v}$  is smaller than  $R_1$ ) for all  $\mathbf{v}$  in the iteration. For all algorithms listed in [Table 1](#), setting  $R_1 = R^2$  is enough. Next, we add some additional assumptions on functions  $S$  and  $V$  that are also met by all optimizers in the AGM framework:

**Assumption F.4.** For any  $\mathbf{g} \in \mathbb{R}^d$ , each entry of  $V(\mathbf{g}\mathbf{g}^\top)$  is non-negative.

**Assumption F.5.** The function  $S$  is  $\mathcal{C}^4$ -smooth on  $\{\mathbf{v} : 0 \preceq \mathbf{v} \preceq R_1\}$ .

Note that  $S$  is also Lipschitz from this assumption, since  $\{\mathbf{v} : 0 \preceq \mathbf{v} \preceq R_1\}$  is a compact set, and the derivative of  $S$  is bounded. Lastly, since  $\eta = o(1)$  but  $1 - \beta_1 = \mathcal{O}(1)$ , we add a reasonable assumption as follows, which helps bound the convergence rate:

**Assumption F.6.** Define  $\gamma := 1 - \frac{2\eta\mu(1-\beta_1)}{R_0}$ . We assume that  $\gamma \geq \frac{\beta_1}{0.95}$ .

Note that in real-world applications,  $\beta_1$  is usually set to 0.9 in the area of NLP or 0.95 in CV, while  $\eta$  is usually smaller than  $1e-2$ , so our assumption aligns with common practice.

**Theorem F.1** (Convergence to a Near-Manifold Neighborhood). *Under [Assumption F.2](#), [Assumption F.3](#), and [Assumption F.6](#), let  $\Gamma$  be a local minimizer manifold. There exists a constant  $\epsilon > 0$ , such that with learning rate  $\eta$ , for any starting point of AGM  $\boldsymbol{\theta}_0 \in \Gamma^\epsilon$  and any  $\delta \in (0, 1)$ , it holds for some  $K = \mathcal{O}(\frac{1}{\eta} \log \frac{1}{\eta})$  with probability at least  $1 - \delta$  that*

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}_K) - \mathcal{L}^* &= \mathcal{O}(\eta \log \frac{1}{\eta\delta}), \\ \|\boldsymbol{\theta}_K - \Phi_{S_K}(\boldsymbol{\theta}_K)\|_2 &= \mathcal{O}(\sqrt{\eta \log \frac{1}{\eta\delta}}). \end{aligned}$$

The above theorem guarantees that the AGMs iterate converge to a neighborhood of the manifold. Building on this result, we now proceed to formally state our main SDE approximation theorem.

**Theorem F.2** (Slow SDE tracks AGMs' trajectory in a weak approximation sense). *Under conditions in [Theorem F.1](#), [Assumption A.1–2.2](#), [Assumption F.1](#), and [Assumption F.5](#), we denote by  $(\boldsymbol{\theta}_0, \mathbf{v}_0)$  the state reached by AGMs by [Theorem F.1](#) such that*

$$\|\boldsymbol{\theta}_K - \Phi_{S_K}(\boldsymbol{\theta}_K)\|_2 = \mathcal{O}(\sqrt{\eta \log \frac{1}{\eta\delta}}).$$

Then, let  $T > 0$  be a constant and let  $\mathbf{X}(t) = (\boldsymbol{\zeta}(t), \mathbf{v}(t))$  be the solution to [Equation \(1\)](#) with initial condition:

$$\boldsymbol{\zeta}(0) = \Phi(\boldsymbol{\theta}_0) \in \Gamma, \quad \mathbf{v}(0) = \mathbf{v}_0 \in \mathbb{R}^d,$$

and we define that the parameters of Adam as  $\bar{\mathbf{X}}_t := (\Phi_{S_t}(\boldsymbol{\theta}_t), \mathbf{v}_t)$ . For any  $C^3$ -smooth function  $g(\boldsymbol{\theta})$ ,

$$\max_{0 \leq t \leq \frac{T}{\eta^2}} \left| \mathbb{E}[g(\bar{\mathbf{X}}_t)] - \mathbb{E}[g(\mathbf{X}(t\eta^2))] \right| = \tilde{\mathcal{O}}(\eta^{0.25}),$$

where  $\tilde{\mathcal{O}}(\cdot)$  hides logarithmic factors and constants that are independent of  $\eta$  but may depend on  $g(\boldsymbol{\theta})$ .

## Appendix G. Constructing the Working Zones

Note that it is generally hard to ensure that a pre-conditioned gradient flow defined by  $\frac{d\boldsymbol{\theta}(t)}{dt} = -\mathbf{S}\nabla\mathcal{L}(\boldsymbol{\theta}(t))$  will converge to some particular minimizer manifold, but it becomes possible when we constrain the discussion inside some small neighborhood of a manifold. So in this subsection, we assume  $\Gamma$  to be some local minimizer manifold, and construct “working zones” that are close enough to  $\Gamma$  such that iterations inside the working zones will be captured by the manifold and obtain certain properties that support the analysis of slow SDE.

**Definition G.1** (Polyak-Łojasiewicz Condition). *For some  $\mu > 0$ , we say some function  $\mathcal{L} : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $\mu$ -Polyak-Łojasiewicz condition (abbreviated as  $\mu$ -PL), if and only if  $\forall \boldsymbol{\theta} \in \mathbb{R}^d$ :*

$$2\mu(\mathcal{L}(\boldsymbol{\theta}) - \mathcal{L}^*) \leq \|\nabla\mathcal{L}(\boldsymbol{\theta})\|_2^2.$$

**Definition G.2** (Neighborhood of a Manifold). *For some manifold  $\Gamma$ , the neighborhood of  $\Gamma$  with radius  $\epsilon$ , denoted by  $\Gamma^\epsilon$  is the set of points  $\boldsymbol{\theta}$  such that  $\exists \boldsymbol{\zeta} \in \Gamma$*

$$\|\boldsymbol{\theta} - \boldsymbol{\zeta}\|_2 \leq \epsilon.$$

**Lemma G.1.** *Assume that  $C_1 < C_2$  are two positive constants, and  $\mathcal{L}$  be a function that satisfies both  $\rho$ -smoothness and  $\mu$ -PL. For any matrix  $\mathbf{S}$  satisfying  $C_1\mathbf{I} \preceq \mathbf{S} \preceq C_2\mathbf{I}$ , consider the preconditioned gradient flow  $\frac{d\boldsymbol{\theta}(t)}{dt} = -\mathbf{S}\nabla\mathcal{L}(\boldsymbol{\theta}(t))$  starting at  $\boldsymbol{\theta}(0) = \boldsymbol{\theta}_0$ . For any  $T > 0$ , we have  $\|\boldsymbol{\theta}_0 - \boldsymbol{\theta}(T)\|_2 \leq \frac{2C_2}{\sqrt{2\mu}C_1} \sqrt{\mathcal{L}(\boldsymbol{\theta}_0) - \mathcal{L}^*}$ .*

**Proof** Since  $C_1\mathbf{I} \preceq \mathbf{S} \preceq C_2\mathbf{I}$ , we have  $\|\mathbf{S}\nabla\mathcal{L}(\boldsymbol{\theta})\|_2 \leq C_2\|\nabla\mathcal{L}(\boldsymbol{\theta})\|_2$  and  $\langle \nabla\mathcal{L}(\boldsymbol{\theta}), \mathbf{S}\nabla\mathcal{L}(\boldsymbol{\theta}) \rangle \geq C_1\|\nabla\mathcal{L}(\boldsymbol{\theta})\|_2^2$  for any  $\boldsymbol{\theta}$ , which implies

$$\langle \nabla\mathcal{L}(\boldsymbol{\theta}), \mathbf{S}\nabla\mathcal{L}(\boldsymbol{\theta}) \rangle \geq \frac{C_1}{C_2} \|\nabla\mathcal{L}(\boldsymbol{\theta})\|_2 \|\mathbf{S}\nabla\mathcal{L}(\boldsymbol{\theta})\|_2.$$

Then for any  $t < T$  we have

$$\begin{aligned} \frac{d}{dt} \sqrt{\mathcal{L}(\boldsymbol{\theta}(t)) - \mathcal{L}^*} &= \frac{1}{2} (\mathcal{L}(\boldsymbol{\theta}(t)) - \mathcal{L}^*)^{-\frac{1}{2}} \cdot \left\langle \nabla\mathcal{L}(\boldsymbol{\theta}(t)), \frac{d\boldsymbol{\theta}(t)}{dt} \right\rangle \\ &\leq -\frac{C_1}{2C_2} (\mathcal{L}(\boldsymbol{\theta}(t)) - \mathcal{L}^*)^{-\frac{1}{2}} \cdot \|\nabla\mathcal{L}(\boldsymbol{\theta}(t))\|_2 \left\| \frac{d\boldsymbol{\theta}(t)}{dt} \right\|_2 \\ &\leq -\frac{C_1}{2C_2} (\mathcal{L}(\boldsymbol{\theta}(t)) - \mathcal{L}^*)^{-\frac{1}{2}} \cdot \sqrt{2\mu(\mathcal{L}(\boldsymbol{\theta}(t)) - \mathcal{L}^*)} \left\| \frac{d\boldsymbol{\theta}(t)}{dt} \right\|_2 \\ &= -\frac{\sqrt{2\mu}C_1}{2C_2} \left\| \frac{d\boldsymbol{\theta}(t)}{dt} \right\|_2. \end{aligned}$$

Integrating both sides gives us

$$\begin{aligned} \sqrt{\mathcal{L}(\boldsymbol{\theta}_0) - \mathcal{L}^*} &\geq \frac{\sqrt{2\mu}C_1}{2C_2} \int_0^T \left\| \frac{d\boldsymbol{\theta}(t)}{dt} \right\|_2 dt \\ &\geq \frac{\sqrt{2\mu}C_1}{2C_2} \|\boldsymbol{\theta}_0 - \boldsymbol{\theta}(T)\|_2. \end{aligned}$$

The above equations complete the proof.

We construct nested working zones  $(\Gamma^{\epsilon_1}, \Gamma^{\epsilon_2}, \Gamma^{\epsilon_3}, \Gamma^{\epsilon_4})$ , where we denote the minimal distance of the minimizer manifold  $\Gamma$  and any other local minimizer manifold as  $\epsilon_4$ , i.e. no other minimizer manifolds exist in  $\Gamma^{\epsilon_4}$ , the  $\epsilon_4$ -neighborhood of  $\Gamma$ . Then we give the following lemma with respect to  $\Gamma^{\epsilon_1}, \Gamma^{\epsilon_2}$  and  $\Gamma^{\epsilon_3}$ .

**Lemma G.2** (Working Zone Lemma). *There exist positive constants  $\epsilon_1, \epsilon_2, \epsilon_3$  such that  $\epsilon_1 < \epsilon_2 < \epsilon_3 < \epsilon_4$  and  $\Gamma^{\epsilon_1}, \Gamma^{\epsilon_2}, \Gamma^{\epsilon_3}$  satisfy the following properties:*

1.  $\mathcal{L}$  is  $\mu$ -PL in  $\Gamma^{\epsilon_3}$  for some  $\mu > 0$ .
2. For any matrix  $\mathbf{S} \in \mathbb{R}^{d \times d}$  such that  $\frac{1}{R_0} \mathbf{I} \preceq \mathbf{S} \preceq \frac{1}{\epsilon} \mathbf{I}$ , any gradient flow preconditioned by  $\mathbf{S}$  and starting from  $\Gamma^{\epsilon_2}$  will converge to some point in  $\Gamma$ .
3. If we view  $\Phi_{S(\mathbf{v})}(\boldsymbol{\theta})$  as a function defined on the support

$$\mathcal{X}^{\epsilon_2} := \{(\mathbf{v}, \boldsymbol{\theta}) : 0 \preceq \mathbf{v} \preceq R_1, \boldsymbol{\theta} \in \Gamma^{\epsilon_2}\},$$

then  $\Phi_{S(\mathbf{v})}(\boldsymbol{\theta})$  is  $\mathcal{C}^4$  on  $\mathcal{X}^{\epsilon_2}$ .

4. For any  $\boldsymbol{\theta} \in \Gamma^{\epsilon_1}$ , the  $\epsilon_1$ -neighborhood of  $\boldsymbol{\theta}$  belongs to  $\Gamma^{\epsilon_2}$ .

**Proof** By Lemma H.3 in Lyu et al. [38], there exists an  $\epsilon_3$ -neighborhood of  $\Gamma$  where  $\mathcal{L}$  is  $\mu$ -PL for some  $\mu > 0$ . WLOG we can let  $\epsilon_3 < \epsilon_4$ . The existence of  $\epsilon_1$  is trivial once  $\epsilon_2$  is determined: Simply set  $\epsilon_1 = \epsilon_2/2$ . So we focus on the construction of  $\Gamma^{\epsilon_2}$  below.

Let  $C_1 = 1/R_0$  and  $C_2 = 1/\epsilon$ . Let  $\tilde{\epsilon}_2$  be some constant such that  $\tilde{\epsilon}_2 + \sqrt{\frac{\rho}{\mu}} \cdot \frac{C_2}{C_1} \tilde{\epsilon}_2 < \epsilon_3$ . For any starting point  $\boldsymbol{\theta}_0 \in \Gamma^{\tilde{\epsilon}_2}$ , and any preconditioning matrix  $\mathbf{S}$  satisfying  $C_1 \mathbf{I} \preceq \mathbf{S} \preceq C_2 \mathbf{I}$ , assume on the contrary that the preconditioned gradient flow starting from  $\boldsymbol{\theta}(0) = \boldsymbol{\theta}_0$  will leave  $\Gamma^{\epsilon_3}$  at some finite time. Then let  $T = \inf \{t : \boldsymbol{\theta}(t) \notin \Gamma^{\epsilon_3}\} < \infty$ . Using Lemma G.1 and combining the  $\mu$ -PL condition, we conclude that  $\|\boldsymbol{\theta}_0 - \boldsymbol{\theta}(T)\|_2 \leq \frac{2C_2}{\sqrt{2\mu}C_1} \sqrt{\mathcal{L}(\boldsymbol{\theta}_0) - \mathcal{L}^*} \leq \frac{2C_2}{\sqrt{2\mu}C_1} \cdot \sqrt{\frac{\mu}{2}} \|\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*\|_2 = \sqrt{\frac{\rho}{\mu}} \cdot \frac{C_2}{C_1} \|\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*\|_2$  for any  $\boldsymbol{\theta}^* \in \Gamma$ . Hence  $\boldsymbol{\theta}(T) \in \Gamma^{\epsilon_3}$ , a contradiction.

Define a function  $f(\mathbf{v}, \boldsymbol{\theta}) : \mathbb{R}^{2d} \rightarrow \mathbb{R}^{2d}$  as

$$f(\mathbf{v}, \boldsymbol{\theta}) = (\mathbf{v}, -S(\mathbf{v})\nabla\mathcal{L}(\boldsymbol{\theta})),$$

then  $f$  is  $\mathcal{C}^4$  on  $\{(\mathbf{v}, \boldsymbol{\theta}) : 0 \preceq \mathbf{v} \preceq R_1, \boldsymbol{\theta} \in \Gamma^{\tilde{r}}\}$ . Let  $\tilde{r}$  be a constant such that  $\tilde{r} > \tilde{\epsilon}_2$ . Substituting  $f_0 = f$ ,  $r = \sqrt{\tilde{r}^2 + d \cdot R_1^2}$ ,  $x_0 = (\mathbf{v}_0, \boldsymbol{\theta}_0)$  such that each entry of  $\mathbf{v}_0$  is  $R_1/2$  and  $\boldsymbol{\theta}_0$  be arbitrary point in  $\Gamma$ , and  $B = \{(\mathbf{v}, \boldsymbol{\theta}) : 0 \preceq \mathbf{v} \preceq R_1, \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq \tilde{r}\}$  into Lemma B.4 in Duistermaat and Kolk [12], we conclude that there exists some constant  $\delta$  such that the mapping  $\gamma_\delta(\mathbf{v}, \boldsymbol{\theta})$  defined by:

$$\boldsymbol{\theta}(0) = \boldsymbol{\theta}, \quad \frac{d\boldsymbol{\theta}(t)}{dt} = -S(\mathbf{v})\nabla\mathcal{L}(\boldsymbol{\theta}(t)), \quad \gamma_\delta(\mathbf{v}, \boldsymbol{\theta}) = \boldsymbol{\theta}(\delta)$$

is well-defined and  $\mathcal{C}^4$  on  $\{(\mathbf{v}, \boldsymbol{\theta}) : 0 \preceq \mathbf{v} \preceq R_1, \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq \tilde{r}\}$ . Note that we require a slight modification of the original proof since  $B$  is now a factorization of a hypercube and a ball instead of a ball, but the convexity of  $B$  is preserved, hence the modification is trivial.

Note that the constant  $\delta$  can be independent with  $\boldsymbol{\theta}_0$  since  $\|\nabla\mathcal{L}\|_2$  and  $\|\nabla^2\mathcal{L}\|_2$  can be uniformly bounded. Take  $\epsilon_2 = 0.9\tilde{\epsilon}_2$ , then for any  $\boldsymbol{\theta} \in \Gamma^{\epsilon_2}$ , a small open neighborhood of  $\boldsymbol{\theta}$  stays in the  $\tilde{\epsilon}_2$ -neighborhoods of two different points on  $\Gamma$ . Taking union of all  $\boldsymbol{\theta}_0 \in \Gamma$ , we conclude that  $\gamma_\delta$  is  $\mathcal{C}^4$  on  $\mathcal{X}^{\epsilon_2}$ . Finally, we use Theorem 6.4 in Falconer [13] to conclude that  $\Phi_{S(\mathbf{v})}(\boldsymbol{\theta})$  is  $\mathcal{C}^4$  on  $\mathcal{X}^{\epsilon_2}$ .

## Appendix H. Proof of the Convergence of AGMs

In this section, we aim to prove [Theorem 2.2](#). First, we introduce some necessary notations.

Recall that  $\gamma = 1 - \frac{2\eta\mu(1-\beta_1)}{R_0}$ . In this part, we will prove that the loss value of Adam converges to  $\tilde{\mathcal{O}}(\gamma^K + \eta)$  within  $K$  steps with high probability. If we substitute  $K = \Theta\left(\frac{1}{\eta} \log \frac{1}{\eta}\right)$ , this will recover [Theorem 2.2](#); However, this convergence analysis works for any  $K = \mathcal{O}(\text{poly}(1/\eta))$ , and substituting  $K = \Theta(\eta^{-2})$  will give us a high probability guarantee that the iteration stays near manifold in the whole scope of our analysis.

Also, we introduce some necessary assumptions used in our proof. To start with, we prove that the descent direction of each step does not veer off the direction of a preconditioned gradient descent, and the mismatch term can also be constrained by a list of martingales. After that, we can ensure a decay in the loss function every step, with some small perturbations that can be dealt with using Azuma-Hoeffding's inequality.

**Lemma H.1.** *Define  $\tilde{\mathbf{v}}_k := \beta_2 \mathbf{v}_{k-1} + (1 - \beta_2) \mathbb{E}_{k-1}[V(\mathbf{g}_{k-1} \mathbf{g}_{k-1}^\top)]$ . There exist constants  $C_{1a}$ ,  $C_{1b}$  such that for any  $k \geq 1$ ,*

$$\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{U}_k \rangle = \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})^\top S(\tilde{\mathbf{v}}_k) \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}) - Y_k - X_k,$$

where  $Y_k$  and  $X_k$  are two  $\mathcal{F}_k$ -measurable random variables such that:

1.  $|Y_k| \leq C_{1a} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})\|_2 \cdot \eta^2$  a.s.
2.  $|X_k| \leq C_{1b} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})\|_2$  a.s., and  $\mathbb{E}_{k-1}[X_k] = 0$ .

**Proof** We first peel the  $S(\tilde{\mathbf{v}}_k)$  part off the  $S(\mathbf{v}_k)$  term:

$$\begin{aligned} \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{U}_k \rangle &= \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\mathbf{v}_k) \mathbf{g}_{k-1} \rangle \\ &= \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\tilde{\mathbf{v}}_k) \mathbf{g}_{k-1} \rangle + \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), (S(\mathbf{v}_k) - S(\tilde{\mathbf{v}}_k)) \mathbf{g}_{k-1} \rangle. \end{aligned}$$

Define  $Y_k$  as  $Y_k = -\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), (S(\mathbf{v}_k) - S(\tilde{\mathbf{v}}_k)) \mathbf{g}_{k-1} \rangle$ , then it holds almost surely that  $|Y_k| \leq \|\nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})\|_2 \| (S(\mathbf{v}_k) - S(\tilde{\mathbf{v}}_k)) \mathbf{g}_{k-1} \|_2$ . Since  $S$  is Lipschitz,  $V$  is linear and

$$\|\tilde{\mathbf{v}}_k - \mathbf{v}_k\|_2 = (1 - \beta_2) \|\mathbb{E}_{k-1}[V(\mathbf{g}_{k-1} \mathbf{g}_{k-1}^\top)] - V(\mathbf{g}_{k-1} \mathbf{g}_{k-1}^\top)\|_2,$$

we conclude that  $|Y_k| \leq C_{1a} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})\|_2 \cdot \eta^2$  a.s. for some constant  $C_{1a}$ . The rest term  $\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\tilde{\mathbf{v}}_k) \mathbf{g}_{k-1} \rangle$  can also be decomposed into a deterministic part and a random part as:

$$\begin{aligned} \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\tilde{\mathbf{v}}_k) \mathbf{g}_{k-1} \rangle &= \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\tilde{\mathbf{v}}_k) (\nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}) + \mathbf{z}_{k-1}) \rangle \\ &= \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})^\top S(\tilde{\mathbf{v}}_k) \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}) + \left\langle \mathbf{z}_{k-1}, S(\tilde{\mathbf{v}}_k)^\top \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}) \right\rangle. \end{aligned}$$

Now we only need to let  $X_k = \langle \mathbf{z}_{k-1}, S(\tilde{\mathbf{v}}_k)^\top \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}) \rangle$ . It's easy to see that  $\mathbb{E}_{k-1}[X_k] = 0$  and  $|X_k| \leq C_{1b} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})\|_2$  a.s. for some constant  $C_{1b}$ , which completes the proof.

**Lemma H.2** (Descent Lemma of the AGM Framework). *For any  $k \geq 1$  it holds that*

$$\mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}(\boldsymbol{\theta}_{k-1}) \leq C_2 \eta^2 - \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1}), \mathbf{U}_i \rangle$$

for some constant  $C_2$ .

**Proof** From the smoothness of  $\mathcal{L}$  we have

$$\mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}(\boldsymbol{\theta}_{k-1}) \leq -\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \eta \mathbf{u}_k \rangle + \frac{\rho \eta^2}{2} \|\mathbf{u}_k\|_2^2.$$

If  $k = 1$ , then  $\mathbf{m}_k = (1 - \beta_1) \mathbf{g}_{k-1}$ , so  $\mathbf{u}_k = (1 - \beta_1) \mathbf{U}_k$ , and the statement trivially holds as long as  $C_2 \geq \frac{\rho}{2} \|\mathbf{u}_k\|_2^2$ . If  $k > 1$ , then the  $-\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{u}_k \rangle$  term can be expanded as

$$\begin{aligned} -\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{u}_k \rangle &= -\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\mathbf{v}_k) \mathbf{m}_k \rangle \\ &= -\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\mathbf{v}_k) (\beta_1 \mathbf{m}_{k-1} + (1 - \beta_1) \mathbf{g}_{k-1}) \rangle \\ &= -\beta_1 \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\mathbf{v}_k) \mathbf{m}_{k-1} \rangle - (1 - \beta_1) \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), S(\mathbf{v}_k) \mathbf{g}_{k-1} \rangle \\ &= -\beta_1 \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-2}), S(\mathbf{v}_{k-1}) \mathbf{m}_{k-1} \rangle - (1 - \beta_1) \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{U}_k \rangle \\ &\quad - \beta_1 \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}) - \nabla \mathcal{L}(\boldsymbol{\theta}_{k-2}), S(\mathbf{v}_{k-1}) \mathbf{m}_{k-1} \rangle \\ &\quad - \beta_1 \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), (S(\mathbf{v}_k) - S(\mathbf{v}_{k-1})) \mathbf{m}_{k-1} \rangle \\ &\leq -\beta_1 \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-2}), S(\mathbf{v}_{k-1}) \mathbf{m}_{k-1} \rangle - (1 - \beta_1) \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{U}_k \rangle \\ &\quad + \beta_1 \|\nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}) - \nabla \mathcal{L}(\boldsymbol{\theta}_{k-2})\|_2 \|S(\mathbf{v}_{k-1}) \mathbf{m}_{k-1}\|_2 \\ &\quad + \beta_1 \|\nabla \mathcal{L}(\boldsymbol{\theta}_{k-1})\|_2 \|(S(\mathbf{v}_k) - S(\mathbf{v}_{k-1})) \mathbf{m}_{k-1}\|_2. \end{aligned}$$

Note that a single step of update on  $\boldsymbol{\theta}$  and  $\mathbf{v}$  is small since

$$\begin{aligned} \boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1} &= \eta \mathbf{u}_k, \\ \mathbf{v}_k - \mathbf{v}_{k-1} &= \beta_2 \mathbf{v}_{k-1} + (1 - \beta_2) V \left( \mathbf{g}_{k-1} \mathbf{g}_{k-1}^\top \right) - \mathbf{v}_{k-1} \\ &= (1 - \beta_2) \left( V \left( \mathbf{g}_{k-1} \mathbf{g}_{k-1}^\top \right) - \mathbf{v}_{k-1} \right) \end{aligned}$$

which implies that  $\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\|_2 = \mathcal{O}(\eta)$  and  $\|\mathbf{v}_k - \mathbf{v}_{k-1}\|_2 = \mathcal{O}(\eta^2)$ . We then leverage the smoothness of  $\nabla \mathcal{L}$  and  $S$  to conclude that there exists some constant  $\tilde{C}_2$  such that

$$-\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{u}_k \rangle \leq -\beta_1 \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-2}), \mathbf{u}_{k-1} \rangle - (1 - \beta_1) \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{U}_k \rangle + \beta_1 \tilde{C}_2 \eta.$$

Giving that  $\mathbf{u}_0 = \mathbf{0}$ , we can expand this formula iteratively as

$$\begin{aligned} -\langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{u}_k \rangle &\leq -\beta_1 \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-2}), \mathbf{u}_{k-1} \rangle - (1 - \beta_1) \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{U}_k \rangle + \beta_1 \tilde{C}_2 \eta \\ &\leq -\beta_1^2 \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-3}), \mathbf{u}_{k-2} \rangle - \beta_1 (1 - \beta_1) \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-2}), \mathbf{U}_{k-1} \rangle \\ &\quad - (1 - \beta_1) \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{k-1}), \mathbf{U}_k \rangle + \beta_1 \tilde{C}_2 \eta + \beta_1^2 \tilde{C}_2 \eta \\ &\leq \dots \\ &\leq -(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1}), \mathbf{U}_i \rangle + \beta_1^{k-i+1} \tilde{C}_2 \eta \\ &\leq \frac{\beta_1}{1 - \beta_1} \tilde{C}_2 \eta - (1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1}), \mathbf{U}_i \rangle. \end{aligned}$$

Plugging in, we get

$$\begin{aligned}\mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}(\boldsymbol{\theta}_{k-1}) &\leq \frac{\beta_1}{1 - \beta_1} \tilde{C}_2 \eta^2 + \frac{\rho \eta^2}{2} \|\mathbf{u}_k\|_2^2 - \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1}), \mathbf{U}_i \rangle \\ &\leq C_2 \eta^2 - \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1}), \mathbf{U}_i \rangle,\end{aligned}$$

for some constant  $C_2$ .

**Lemma H.3.** Recall that  $\gamma = 1 - \frac{2\eta\mu(1-\beta_1)}{R_0}$ . For any  $k \geq 0$ , we have:

$$\mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}^* \leq \gamma^k (\mathcal{L}(\boldsymbol{\theta}_0) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{i=1}^k X_i \sum_{j=i}^k \gamma^{k-j} \beta_1^{j-i} + C_3 \eta$$

for some constant  $C_3$ .

**Proof** We start from [Lemma H.2](#) and plug in [Lemma H.1](#):

$$\begin{aligned}\mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}(\boldsymbol{\theta}_{k-1}) &\leq C_2 \eta^2 - \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} \langle \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1}), \mathbf{U}_i \rangle \\ &= C_2 \eta^2 - \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} \left( \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})^\top S(\tilde{\mathbf{v}}_i) \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1}) - Y_i - X_i \right).\end{aligned}$$

Note that  $S(\tilde{\mathbf{v}}_i) \succeq \frac{1}{R_0}$ , so  $\nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})^\top S(\tilde{\mathbf{v}}_i) \nabla \mathcal{L}(\boldsymbol{\theta}_{i-1}) \geq \frac{1}{R_0} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})\|_2^2$  for any  $i$ . Combining with the  $\mu$ -PL property  $\|\nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})\|_2^2 \geq 2\mu (\mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^*)$ , we have

$$\mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}(\boldsymbol{\theta}_{k-1}) \leq C_2 \eta^2 - \frac{2\eta\mu(1 - \beta_1)}{R_0} \sum_{i=1}^k \beta_1^{k-i} (\mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} (Y_i + X_i).$$

Since  $|Y_i| \leq C_{1a} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})\|_2 \cdot \eta^2$  for every  $i$ , the effect of  $Y$  is negligible:

$$\left| \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} Y_i \right| \leq C_{1a} \eta^3 \cdot \max_{i=0}^k \{\|\nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})\|_2\} = o(\eta^2),$$

and we can absorb it into the  $C_2 \eta^2$  term to write out that

$$\begin{aligned}\mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}^* &\leq \tilde{C}_3 \eta^2 + \mathcal{L}(\boldsymbol{\theta}_{k-1}) - \mathcal{L}^* - \frac{2\eta\mu(1 - \beta_1)}{R_0} \sum_{i=1}^k \beta_1^{k-i} (\mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} X_i \\ &\leq \tilde{C}_3 \eta^2 + \left( 1 - \frac{2\eta\mu(1 - \beta_1)}{R_0} \right) (\mathcal{L}(\boldsymbol{\theta}_{k-1}) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} X_i \\ &= \tilde{C}_3 \eta^2 + \gamma (\mathcal{L}(\boldsymbol{\theta}_{k-1}) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} X_i\end{aligned}$$

for some constant  $\tilde{C}_3$ . Note that we can expand the  $\mathcal{L}(\boldsymbol{\theta}_{k-1}) - \mathcal{L}^*$  term iteratively to obtain a generic formula for  $\mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}^*$ :

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}_k) - \mathcal{L}^* &\leq \gamma (\mathcal{L}(\boldsymbol{\theta}_{k-1}) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} X_i + \tilde{C}_3 \eta^2 \\ &\leq \gamma^k (\mathcal{L}(\boldsymbol{\theta}_0) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{j=1}^k \gamma^{k-j} \sum_{i=1}^j \beta_1^{j-i} X_i + \sum_{j=1}^k \gamma^{k-j} \tilde{C}_3 \eta^2 \\ &\leq \gamma^k (\mathcal{L}(\boldsymbol{\theta}_0) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{i=1}^k X_i \sum_{j=i}^k \gamma^{k-j} \beta_1^{j-i} + C_3 \eta, \end{aligned}$$

where  $C_3 = \tilde{C}_3 \cdot \frac{R_0}{2\mu(1-\beta_1)}$ .

**Lemma H.4.** *Let  $k \leq K = \mathcal{O}(\text{poly}(1/\eta))$  and  $f : (\{0, 1, \dots, k-1\} \times (0, 1)) \rightarrow \mathbb{R}^+$  be a function. Let  $\{X_i\}_{i=1}^k$  be any martingale difference sequence such that:*

1.  $X_i$  is  $\mathcal{F}_i$ -measurable and  $\mathbb{E}_{i-1}[X_i] = 0$ ;
2.  $|X_i| \leq C_{1b} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})\|_2$  a.s.

for any  $i \in [k]$ . If for any  $i \in [k]$  and  $\delta \in (0, 1)$ , it holds with probability  $1 - \delta$  that

$$\mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^* \leq f(i, \delta),$$

then  $\forall \delta \in (0, 1)$ , with probability  $1 - \delta$ , we have  $\mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^* \leq f(i, \frac{\delta}{2k})$  for all  $i \in [k]$ , and that

$$\left| \sum_{i=1}^k \gamma^{k-i} X_i \right| \leq C_4 \sqrt{\sum_{i=1}^k \gamma^{2k-2i} f\left(i, \frac{\delta}{2k}\right) \log \frac{4}{\delta}}$$

for some constant  $C_4$ .

**Remark.** The  $\{X_i\}$  here may not necessarily equal the  $\{X_i\}$  defined in [Lemma H.1](#); we just make it general to benefit future steps. In fact, when we leverage this lemma later, we will multiply that of [Lemma H.1](#) by some scalar  $\in (0, 1)$ .

**Proof** Note that  $\sum_{i=1}^k \gamma^{k-i} X_i$  is a sum of martingale differences. Moreover, since  $\mathcal{L}$  is  $\rho$ -smooth and  $\exists C_{1b}$  s.t. every  $|X_i|$  is bounded by  $C_{1b} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})\|_2$  ([Lemma H.1](#)), we have

$$\begin{aligned} |X_i| &\leq C_{1b} \|\nabla \mathcal{L}(\boldsymbol{\theta}_{i-1})\|_2 \\ &\leq C_{1b} \sqrt{2\rho (\mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^*)} \\ &\leq C_{1b} \sqrt{2\rho f(i, \delta')} \quad \text{if } \mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^* \leq f(i, \delta'). \end{aligned}$$

Since  $\mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^* \leq f(i, \delta')$  holds with probability  $1 - \delta'$  instead of probability 1, we create a new martingale difference sequence that masks out all the positions that exceed the bound. Specifically, we define  $X'_{i,\delta'}$  as:

$$X'_{i,\delta'} = \begin{cases} X_i & \text{if } \mathcal{L}(\boldsymbol{\theta}_{i-1}) - \mathcal{L}^* \leq f(i, \delta'), \\ 0 & \text{else.} \end{cases}$$

This ensures that  $\left|X'_{i,\delta'}\right| \leq C_{1b}\sqrt{2\rho f(i,\delta')}$  a.s. Then Azuma-Hoeffding's inequality gives us that for any  $\epsilon'$ ,

$$\mathbb{P}\left[\left|\sum_{i=1}^k \gamma^{k-i} X'_{i,\delta'}\right| \geq \epsilon'\right] \leq 2 \exp\left(\frac{-\epsilon'^2}{4 \sum_{i=1}^k C_{1b}^2 \gamma^{2k-2i} \rho f(i,\delta')}\right),$$

denoting the right hand side as  $\frac{\delta}{2}$  gives that for any  $\delta$ , with probability  $1 - \frac{\delta}{2}$ ,

$$\left|\sum_{i=1}^k \gamma^{k-i} X'_{i,\delta'}\right| \leq \sqrt{4 \sum_{i=1}^k C_{1b}^2 \gamma^{2k-2i} \rho f(i,\delta') \log \frac{4}{\delta}}.$$

Let  $\delta' = \frac{\delta}{2k}$ , by union bound,  $\mathcal{L}(\theta_{i-1}) - \mathcal{L}^* \leq f\left(i, \frac{\delta}{2k}\right)$  for all  $i \in [k]$  with probability  $1 - \frac{\delta}{2}$ , which also implies  $X'_{i,\delta'} = X_i$  for all  $i \in [k]$ . So with probability  $1 - \delta$ , the following two statements hold simultaneously for all  $i \in [k]$ :

$$\mathcal{L}(\theta_{i-1}) - \mathcal{L}^* \leq f\left(i, \frac{\delta}{2k}\right)$$

and

$$\begin{aligned} \left|\sum_{i=1}^k \gamma^{k-i} X_i\right| &\leq \sqrt{4 \sum_{i=1}^k C_{1b}^2 \gamma^{2k-2i} \rho f(i,\delta') \log \frac{4}{\delta}} \\ &= C_4 \sqrt{\sum_{i=1}^k \gamma^{2k-2i} f\left(i, \frac{\delta}{2k}\right) \log \frac{4}{\delta}}, \end{aligned}$$

where  $C_4 = 2C_{1b}\sqrt{\rho}$ .

**Lemma H.5** (Convergence Bound of the AGM Framework). *Recall that  $\frac{\beta_1}{\gamma} \leq 0.95$ . Let  $K = \mathcal{O}(\text{poly}(1/\eta))$ . Under mild restrictions on  $K$ , for any  $k \leq K$ ,  $\delta \in (0, 1)$ , it holds with probability at least  $1 - \delta$  that*

$$\mathcal{L}(\theta_k) - \mathcal{L}^* \leq \left(C_{5a}\gamma^k + C_{5b}\eta\right) \log \frac{K}{\delta}$$

for some constants  $C_{5a}$  and  $C_{5b}$ .

**Proof** We denote the bound with  $1 - \delta$  probability as  $f(k, \delta) := (C_{5a}\gamma^k + C_{5b}\eta) \log \frac{K}{\delta}$ , where the constants  $C_{5a}, C_{5b}$  will be specified by us later. We prove by induction. When  $k = 0$ , we need

$$(C_{5a} + C_{5b}\eta) \left(\log K + \log \frac{1}{\delta}\right) \geq \mathcal{L}(\theta_0) - \mathcal{L}^*,$$

where setting  $C_{5a} \geq \frac{\mathcal{L}(\theta_0) - \mathcal{L}^*}{\log K}$  suffices. Now assume that the statement holds for  $0, 1, \dots, k-1$ . From Lemma H.3, we have

$$\mathcal{L}(\theta_k) - \mathcal{L}^* \leq \gamma^k (\mathcal{L}(\theta_0) - \mathcal{L}^*) + \eta(1 - \beta_1) \sum_{i=1}^k X_i \sum_{j=i}^k \gamma^{k-j} \beta_1^{j-i} + C_3\eta.$$

We can bound the coefficients by

$$\begin{aligned} \sum_{j=i}^k \gamma^{k-j} \beta_1^{j-i} &= \gamma^{k-i} \sum_{j=0}^{k-i} \left( \frac{\beta_1}{\gamma} \right)^j \\ &\leq \gamma^{k-i} \cdot \frac{1}{1 - \frac{\beta_1}{\gamma}} \\ &\leq 20\gamma^{k-i}, \end{aligned}$$

where the last inequality is due to our assumption  $\frac{\beta_1}{\gamma} \leq 0.95$ . Let  $\tilde{X}_i := \frac{\sum_{j=i}^k \gamma^{k-j} \beta_1^{j-i}}{20\gamma^{k-i}} X_i$ , then  $\{\tilde{X}_i\}_{i=1}^k$  is also a martingale difference sequence and  $|\tilde{X}_i| \leq |X_i| \leq C_{1b} \|\nabla \mathcal{L}(\theta_i)\|_2$  a.s. From [Lemma H.4](#), with probability  $1 - \delta$ ,  $\mathcal{L}(\theta_{i-1}) - \mathcal{L}^* \leq f(i, \frac{\delta}{2k})$  holds for all  $i \in [k]$  and  $|\sum_{i=1}^k \gamma^{k-i} \tilde{X}_i| \leq C_4 \sqrt{\sum_{i=1}^k \gamma^{2k-2i} f(i, \frac{\delta}{2k}) \log \frac{4}{\delta}}$  holds. If this happens, we have

$$\begin{aligned} &\eta(1 - \beta_1) \sum_{i=1}^k X_i \sum_{j=i}^k \gamma^{k-j} \beta_1^{j-i} \\ &\leq 20C_4\eta(1 - \beta_1) \sqrt{\sum_{i=1}^k \gamma^{2k-2i} f\left(i, \frac{\delta}{2k}\right) \log \frac{4}{\delta}} \\ &\leq 20C_4\eta(1 - \beta_1) \sqrt{\sum_{i=1}^k \gamma^{2k-2i} (C_{5a}\gamma^i + C_{5b}\eta) \log \frac{2kK}{\delta} \log \frac{4}{\delta}} \\ &\leq 20C_4\eta(1 - \beta_1) \sqrt{\sum_{i=1}^k \gamma^{2k-2i} C_{5a}\gamma^i + \sum_{i=1}^k \gamma^{2k-2i} C_{5b}\eta} \cdot \sqrt{\log \frac{2K^2}{\delta} \log \frac{4}{\delta}} \\ &\leq 20C_4\eta(1 - \beta_1) \sqrt{\frac{C_{5a}\gamma^k}{1 - \gamma} + \frac{C_{5b}\eta}{1 - \gamma^2}} \cdot \sqrt{\log \frac{2K^2}{\delta} \log \frac{4}{\delta}} \\ &\leq 20C_4\eta(1 - \beta_1) \left( \sqrt{\frac{C_{5a}\gamma^k}{1 - \gamma}} + \sqrt{\frac{C_{5b}\eta}{1 - \gamma}} \right) \cdot \sqrt{\log \frac{2K^2}{\delta} \log \frac{4}{\delta}}. \end{aligned}$$

As long as  $K \geq \max\{2\delta^2, 4\}$  (which is a mild restriction on  $K$ ), we have  $\sqrt{\log \frac{2K^2}{\delta} \log \frac{4}{\delta}} \leq \sqrt{3 \log^2 \frac{K}{\delta}}$ . Plugging in  $\frac{1}{1-\gamma} = \frac{R_0}{2\mu(1-\beta_1)} \cdot \frac{1}{\eta}$ , we have

$$\begin{aligned} & \eta(1-\beta_1) \sum_{i=1}^k X_i \sum_{j=i}^k \gamma^{k-j} \beta_1^{j-i} \\ & \leq 20C_4(1-\beta_1) \left( \sqrt{\frac{C_{5a}R_0}{2\mu(1-\beta_1)}} \cdot \sqrt{\eta\gamma^k} + \sqrt{\frac{C_{5b}R_0}{2\mu(1-\beta_1)}} \cdot \eta \right) \cdot \sqrt{3} \log \frac{K}{\delta} \\ & \leq 10C_4 \sqrt{\frac{6C_{5a}R_0(1-\beta_1)}{\mu}} \cdot \sqrt{\eta\gamma^k} \log \frac{K}{\delta} + 10C_4 \sqrt{\frac{6C_{5b}R_0(1-\beta_1)}{\mu}} \cdot \eta \log \frac{K}{\delta} \\ & \leq (C_{5c}\gamma^k + C_{5d}\eta) \log \frac{K}{\delta}, \end{aligned}$$

where  $C_{5c} = 5C_4 \sqrt{6C_{5a}R_0(1-\beta_1)/\mu}$  and  $C_{5d} = 5C_4 \sqrt{6C_{5a}R_0(1-\beta_1)/\mu} + 10C_4 \sqrt{6C_{5b}R_0(1-\beta_1)/\mu}$ . Now as long as  $K \geq e\delta$  (so that  $\log \frac{K}{\delta} \geq 1$ ), we have

$$\begin{aligned} \mathcal{L}(\theta_k) - \mathcal{L}^* & \leq \gamma^k (\mathcal{L}(\theta_0) - \mathcal{L}^*) + \eta(1-\beta_1) \sum_{i=1}^k X_i \sum_{j=i}^k \gamma^{k-j} \beta_1^{j-i} + C_3\eta \\ & \leq \gamma^k (\mathcal{L}(\theta_0) - \mathcal{L}^*) + (C_{5c}\gamma^k + C_{5d}\eta) \log \frac{K}{\delta} + C_3\eta \\ & \leq (C_{5c} + \mathcal{L}(\theta_0) - \mathcal{L}^*) \gamma^k \log \frac{K}{\delta} + (C_{5d} + C_3) \eta \log \frac{K}{\delta}. \end{aligned}$$

To complete the induction, we need  $C_{5a}, C_{5b}$  satisfy

$$\begin{cases} C_{5a} & \geq C_{5c} + \mathcal{L}(\theta_0) - \mathcal{L}^* & = 5C_4 \sqrt{\frac{6C_{5a}R_0(1-\beta_1)}{\mu}} + \mathcal{L}(\theta_0) - \mathcal{L}^* \\ C_{5b} & \geq C_{5d} + C_3 & = 5C_4 \sqrt{\frac{6C_{5a}R_0(1-\beta_1)}{\mu}} + 10C_4 \sqrt{\frac{6C_{5b}R_0(1-\beta_1)}{\mu}} + C_3. \end{cases}$$

Notice that the right-hand side grows at the rate of the square root of  $C_{5a}$  and  $C_{5b}$ , so there must exist some feasible constants  $C_{5a}$  and  $C_{5b}$ . Summarizing, under mild restrictions  $K \geq \max\{2\delta^2, e\delta, 4\}$ , the statement  $\mathcal{L}(\theta_k) - \mathcal{L}^* \leq (C_{5a}\gamma^k + C_{5b}\eta) \log \frac{K}{\delta}$  holds with probability  $1 - \delta$ , completing the induction.

**Proof** [Proof of [Theorem F.1](#)] By [Lemma G.2](#), there exists some constant  $\epsilon_3$  such that  $\mathcal{L}$  is  $\mu$ -PL in  $\Gamma^{\epsilon_3}$ . Letting  $\gamma^K = \mathcal{O}(\eta)$  gives  $K = \mathcal{O}(\frac{1}{\eta} \log \frac{1}{\eta})$ , completing the proof.

## Appendix I. Proof of the SDE Approximation of AGMs

In this section, we present a detailed derivation of our slow SDE approximation of the AGM framework as shown in [Theorem 2.1](#). First, we formalize our settings and introduce some additional notations that will be used in our proof. In the AGM framework, the transition from  $\theta_k$  to  $\theta_{k+1}$  is defined as:

$$\begin{aligned} \mathbf{m}_{k+1} &:= \beta_1 \mathbf{m}_k + (1-\beta_1) \nabla \ell_k(\theta_k) \\ \mathbf{v}_{k+1} &:= \beta_2 \mathbf{v}_k + (1-\beta_2) V \left( \nabla \ell_k(\theta_k) \nabla \ell_k(\theta_k)^\top \right) \\ \theta_{k+1} &:= \theta_k - \eta S(\mathbf{v}_{k+1}) \mathbf{m}_{k+1}, \end{aligned}$$

where  $S : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$  is  $\rho_s$ -smooth, positive definite (p.d) and satisfies  $S(v) \preceq \epsilon^{-1}I$  for some  $\epsilon > 0$  and any  $v \in \mathbb{R}^d$ , and  $V : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$  is linear. The algorithm starts from the initial state  $\theta_0$ , and we set  $m_0 = v_0 = \mathbf{0}$ . For every  $k \geq 0$ , we use **step**  $k + 1$  to refer to the process of obtaining the noisy gradient  $\nabla \ell_k(\theta_k)$  and performing the above three updates. For any  $k \geq 0$ , to simplify the notation, we denote that

$$\begin{aligned} g_k &:= \nabla \ell_k(\theta_k), \quad z_k := \ell_k(\theta_k) - \mathcal{L}(\theta_k) \sim \mathcal{Z}(\theta_k), \quad S_k := S(v_k), \\ U_{k+1} &:= S(v_k)g_k, \quad u_{k+1} := S(v_k)m_{k+1}, \quad \phi_k := \Phi_{S_k}(\theta_k) \end{aligned}$$

**Time**  $k$  refers to the time right before step  $k + 1$  happens, i.e. the time right after we get  $\theta_k$ . We also define  $\{\mathcal{F}_k\}$  as the natural filtration generated by the history of optimization, where each  $\mathcal{F}_k = \sigma(\theta_0, z_0, \dots, z_{k-1})$  can be interpreted as “all the information available up to time  $k$ ”. We use the notation  $\mathbb{E}_k$  to denote the expectation conditioned on  $\mathcal{F}_k$ .

**Remark I.1.** Without causing confusion, we reword the definition of  $\theta_0$  and  $v_0$  here. In the following calculation in [Appendix I](#),  $\theta_0$  and  $v_0$  do not represent the parameters that are initialized at the real beginning of training. According to [Theorem F.1](#), we define  $\theta_0$  as the parameter near the minimizer manifold such that  $\|\theta_0 - \phi_0\|_2 = \mathcal{O}(\sqrt{\eta \log \frac{1}{\eta}})$ , and  $v_0$  is the velocity vector as the corresponding time step as  $\theta_0$ . Our SDE approximation then describes AGMs’ dynamics after reaching such a state  $(\theta_0, v_0)$ .

### I.1. Lemmas for Adaptive Manifold Projection

Before we characterize the projections, we introduce some properties of the preconditioned projection function in this part.

**Lemma I.1** (Adaption of Lemma C.2 in Li et al. [35]). *For any  $x \in \mathbb{R}^d$ , and any p.d matrix  $S \in \mathbb{R}^{d \times d}$ , it holds that  $\partial \Phi_S(x) S \nabla \mathcal{L}(x) = 0$ , and*

$$\partial^2 \Phi_S(x) [S \nabla \mathcal{L}(x), S \nabla \mathcal{L}(x)] = -\partial \Phi_S(x) S \nabla^2 \mathcal{L}(x) S \nabla \mathcal{L}(x).$$

**Proof** We consider a trajectory starting from  $x(0) = x$ , with an ODE  $\frac{dx(t)}{dt} = -S \nabla \mathcal{L}(x(t))$ , thus by the definition of  $\Phi_S$ , we have  $\Phi_S(x) = \Phi_S(x(t))$ , then we have

$$\frac{d\Phi_S(x(t))}{dt} = -\partial \Phi_S(x) S \nabla \mathcal{L}(x) = 0.$$

Further, we take the second derivative of  $\Phi_S(x(t))$  with respect to  $t$

$$\frac{d^2 \Phi_S(x(t))}{dt^2} = \partial^2 \Phi_S(x) [S \nabla \mathcal{L}(x), S \nabla \mathcal{L}(x)] + \partial \Phi_S(x) S \nabla^2 \mathcal{L}(x) S \nabla \mathcal{L}(x) = 0.$$

Taking  $t = 0$  completes the proof.

**Lemma I.2.** *For any  $x \in \Gamma$ , and a p.d matrix  $S$ , it holds that  $\partial \Phi_S(x) \nabla^2 \mathcal{L}(x) = \mathbf{0}$ .*

**Proof** From Lemma C.1 in Li et al. [35], we have for  $u \in T_x(\Gamma)$ ,  $\nabla^2 \mathcal{L}(x)u = 0$ , and for  $u \in T_x^\perp(\Gamma)$ , it is direct corollary of Lemma 4.3 in Gu et al. [20] that

$$\partial \Phi_S(x) S u = 0.$$

The above identity completes the proof.

**Lemma I.3.** *For any  $x \in \Gamma$ ,  $u, v \in \mathbb{R}^d$ ,  $p.d$  matrix  $S$ , and  $v \in T_x(\Gamma)$ , it holds that*

$$\partial^2 \Phi_S(x)[uv^T] = -\partial \Phi_S(x) S \partial^2(\nabla \mathcal{L})(x)[\nabla^2 \mathcal{L}(x)^\dagger S^{-1} uv^T] - S^{-1} \nabla^2 \mathcal{L}(x)^\dagger \partial^2(\nabla \mathcal{L})(x)[S \partial \Phi(x) uv^T].$$

**Proof** We define  $P := S^{1/2}$ . And we do a reparameterization as  $x' := P^{-1}x$ ,  $\mathcal{L}'(x) := \mathcal{L}(Px)$ , then we have

$$\begin{aligned} \partial \Phi'(x') &= P \partial \Phi_S(Px) P \\ \nabla^2 L'(x') &= P \nabla^2 L(Px) P \\ \partial^2(\nabla L')(x')[M] &= P \partial^2(\nabla L)(Px)[PMP] \\ \partial^2 \Phi'(x')[M] &= P \partial^2 \Phi(x)[PMP]. \end{aligned}$$

Notice that in the space of  $x'$ , the adaptive projection mapping  $\Phi_S$  turns into a fixed gradient flow projection. And this allows us to directly apply Lemma C.4 in Li et al. [35], which gives

$$\partial^2 \Phi'(x')[v, u] = -\partial \Phi'(x') \partial^2(\nabla \mathcal{L}')(x')[v, \nabla^2 \mathcal{L}'(x')^\dagger u] - \nabla^2 \mathcal{L}'(x')^\dagger \partial^2(\nabla \mathcal{L}')(x')[v, \partial \Phi'(x')u].$$

A slight modification using the above transformations gives

$$\begin{aligned} \partial^2 \Phi_S(x)[Pv, Pu] &= -\partial \Phi_S(x) S \partial^2(\nabla \mathcal{L})(x)[Pv, \nabla^2 \mathcal{L}(x)^\dagger S^{-1} Pu] \\ &\quad - S^{-1} \nabla^2 \mathcal{L}(x)^\dagger \partial^2(\nabla \mathcal{L})(x)[Pv, S \partial \Phi(x) Pu]. \end{aligned}$$

We now redefine  $u = Pu$ ,  $v = Pv$ , and we organize the above equation

$$\partial^2 \Phi_S(x)[uv^T] = -\partial \Phi_S(x) S \partial^2(\nabla \mathcal{L})(x)[\nabla^2 \mathcal{L}(x)^\dagger S^{-1} uv^T] - S^{-1} \nabla^2 \mathcal{L}(x)^\dagger \partial^2(\nabla \mathcal{L})(x)[S \partial \Phi(x) uv^T].$$

We completes the proof.

## I.2. Iteration Stays Near Manifold

Now we begin the final preparations before deriving the slow SDE near the manifold. Denote  $K := \lfloor T\eta^{-2} \rfloor$  be the total number of steps in our analysis. Let  $\beta$  be some constant in  $(0, 0.5)$ , whose exact value will be specified later. First, we bound the movement of projected steps by showing that  $\phi$  shifts no more than  $\tilde{\mathcal{O}}(\eta^{0.5-0.5\beta})$  within  $\Delta K := \eta^{-1-\beta}$  steps, demonstrating the “slowness” of the dynamics of AGMs after the projection.

**Lemma I.4.** *If  $\phi_k$  stays inside  $\Gamma^{\epsilon_2}$  for any  $k \in [0, K]$ , then for any  $\delta = \mathcal{O}(\text{poly}(\eta))$ , with probability  $1 - \delta$ , for any  $k \in [0, K - \Delta K]$ ,  $\Delta k \in [\Delta K]$ ,*

$$\|\phi_{k+\Delta k} - \phi_k\|_2 \leq C_6 \eta^{0.5-0.5\beta} \sqrt{\log \frac{1}{\eta\delta}}$$

for some constant  $C_6$ .

**Proof** Recall that  $\Phi_{S(v)}(\theta)$  is  $\mathcal{C}^4$  on  $\mathcal{X}^{\epsilon_2} := \{(v, \theta) : 0 \preceq v \preceq R_1, \theta \in \Gamma^{\epsilon_2}\}$ , since  $\mathcal{X}^{\epsilon_2}$  is compact,  $\Phi_{S(v)}(\theta)$  is then bounded and Lipschitz on  $\mathcal{X}^{\epsilon_2}$ . Similarly,  $\partial \Phi_{S(v)}(\theta)$  is bounded and Lipschitz on

$\mathcal{X}^{\epsilon_2}$ . For any  $k \in [0, K)$ , let  $\bar{k} = k - 2 \log_{\beta_1} \eta$ , we have:

$$\begin{aligned}
 \phi_{k+1} - \phi_k &= \Phi_{S(v_{k+1})}(\theta_{k+1}) - \Phi_{S(v_k)}(\theta_k) \\
 &= \Phi_{S(v_{\bar{k}})}(\theta_{k+1}) - \Phi_{S(v_{\bar{k}})}(\theta_k) + \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right) \\
 &= \partial \Phi_{S(v_{\bar{k}})}(\theta_k)(\theta_{k+1} - \theta_k) + \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right) \\
 &= \partial \Phi_{S(v_{\bar{k}})}(\theta_k)(\eta S(v_k) \mathbf{m}_k) + \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right) \\
 &= \partial \Phi_{S(v_{\bar{k}})}(\theta_{\bar{k}})(\eta S(v_{\bar{k}}) \mathbf{m}_k) + \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right),
 \end{aligned}$$

where the second equality comes from the fact that one step of update on  $\mathbf{v}$  is of  $\mathcal{O}(\eta^2)$  and the Lipschitzness of  $S$  and  $\Phi$ , the third equality comes from  $\|\theta_{k+1} - \theta_k\|_2 = \mathcal{O}(\eta)$ , and the last equality follows from the boundedness and Lipschitzness of  $\partial \Phi$ . We can decompose  $\mathbf{m}_k$  as:

$$\begin{aligned}
 \mathbf{m}_k &= (1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} (\nabla \mathcal{L}(\theta_i) + \mathbf{z}_i) + \mathcal{O}(\eta^2) \\
 &= (1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} \left( \nabla \mathcal{L}(\theta_{\bar{k}}) + \mathcal{O}\left(\eta \log \frac{1}{\eta}\right) \right) + (1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} \mathbf{z}_i + \mathcal{O}(\eta^2).
 \end{aligned}$$

A key observation is that  $\partial \Phi_{S(v_{\bar{k}})}(\theta_{\bar{k}}) S(v_{\bar{k}}) \nabla \mathcal{L}(\theta_{\bar{k}}) = 0$  from [Lemma I.2](#), which allows us to view  $\phi_{k+1} - \phi_k$  as  $\sum_{i=\bar{k}}^k \tilde{\mathbf{z}}_{k,i} + \mathcal{O}(\eta^2 \log \frac{1}{\eta})$  where  $\tilde{\mathbf{z}}_{k,i} = \partial \Phi_{S(v_{\bar{k}})}(\theta_{\bar{k}})(\eta(1 - \beta_1) \beta_1^{k-i} S(v_{\bar{k}}) \mathbf{z}_i)$ . Note that  $\tilde{\mathbf{z}}_{k,i}$  is  $\mathcal{F}_{i+1}$ -measurable and its mean is  $\mathbf{0}$ , since  $\tilde{\mathbf{z}}_{k,i}$  just applies a linear tensor transformation to  $\mathbf{z}_i$ . If we define a constant  $C_{6a} := \sup \{ \|\partial \Phi_{S(v)}(\theta)\|_2 \mid (\mathbf{v}, \theta) \in \mathcal{X}^{\epsilon_2} \} \cdot (1 - \beta_1) \cdot \epsilon^{-1}$  that is independent of  $k$  and  $i$ , then  $\|\tilde{\mathbf{z}}_{k,i}\|_2$  is almost surely bounded by  $\eta \beta_1^{k-i} C_{6a} \|\mathbf{z}_i\|_2$ .

For any  $k \in [0, K - \Delta K]$  and  $\Delta k \in [\Delta K]$ , we have

$$\begin{aligned}
 \phi_{k+\Delta k} - \phi_k &= \sum_{j=k}^{k+\Delta k-1} (\phi_{j+1} - \phi_j) \\
 &= \sum_{j=k}^{k+\Delta k-1} \left( \sum_{i=j-2 \log_{\beta_1} \eta}^j \tilde{\mathbf{z}}_{j,i} + \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right) \right) \\
 &= \sum_{i=k-2 \log_{\beta_1} \eta}^{k+\Delta k-1} \sum_{j=i}^{\min\{k+\Delta k-1, j+2 \log_{\beta_1} \eta\}} \tilde{\mathbf{z}}_{j,i} + \tilde{\mathcal{O}}(\eta^{1-\beta})
 \end{aligned}$$

Denote  $\mathbf{Z}_i := \sum_{j=i}^{\min\{k+\Delta k-1, j+2 \log_{\beta_1} \eta\}} \tilde{\mathbf{z}}_{j,i}$ , then each  $\mathbf{Z}_i$  is a linear transformation of  $\mathbf{z}_i$  so it is with zero mean, and also  $\|\mathbf{Z}_i\|_2 \leq \eta \cdot \frac{C_{6a}}{1-\beta_1} \|\mathbf{z}_i\|_2 \leq \eta \cdot \frac{C_{6a} R}{1-\beta_1}$  a.s. Azuma-Hoeffding's inequality then

gives that for any  $\delta = \mathcal{O}(\text{poly}(\eta))$ , with probability  $1 - \delta$ ,

$$\begin{aligned}\phi_{k+\Delta k} - \phi_k &\leq \sqrt{2\eta^2 \left( \frac{C_{6a}R}{1 - \beta_1} \right)^2 \cdot (R_{\text{grp}}H + 2\log_{\beta_1} \eta) \cdot \log \frac{2}{\delta}} \\ &\leq C_{6b} \sqrt{\eta^{1-\beta} \log \frac{2}{\delta}}\end{aligned}$$

for some constant  $C_{6b}$ . Finally, plugging in  $\delta' = \frac{\delta}{K \cdot \Delta K}$  and taking union bound over all  $k \in [0, K - \Delta K]$  and  $\Delta k \in [\Delta K]$  gives the theorem.

With the concentration bounds so far, we can show that the dynamics behaves ‘well’ during the whole iteration, and we formalize this idea below.

**Definition I.1** ( $\delta$ -good). *For any  $\delta = \mathcal{O}(\text{poly}(\eta))$  and any step  $\hat{K} \in [K]$ , we define step  $\hat{K}$  to be  $\delta$ -good if and only if the simultaneous establishment of the following statements:*

1. *For any  $k \in [0, \hat{K}]$ ,  $\phi_k \in \Gamma$  and  $\|\theta_k - \phi_k\|_2 \leq C_{8a} \sqrt{\eta \log \frac{1}{\eta\delta}}$ .*
2. *For any  $k \in [0, \hat{K} - \Delta K]$ ,  $\Delta k \in [\Delta K]$ ,  $\|\phi_{k+\Delta k} - \phi_k\|_2 \leq C_{8b} \eta^{0.5-0.5\beta} \sqrt{\log \frac{1}{\eta\delta}}$ .*

Here  $C_{8a} = \frac{4C_2}{\sqrt{2\mu}C_1} \cdot \sqrt{(C_{5a} + C_{5b})}$  and  $C_{8b} = C_6\sqrt{2}$  are two constants.

**Lemma I.5.** *When  $\eta$  is sufficiently small, with probability  $1 - \eta^{2025}$ , the event  $\eta^{2025}$ -good holds for any step  $\hat{K}$  in  $[K]$ .*

**Proof** Denote  $\delta := \eta^{2025}$ . From [Lemma H.5](#), with probability  $1 - \delta/2$ , all  $k \in [0, K]$  satisfy  $\mathcal{L}(\theta_k) - \mathcal{L}^* \leq (C_{5a} + C_{5b})\eta \log \frac{2K}{\delta}$ . Combining [Lemma G.1](#), this implies  $\|\theta_k - \phi_k\|_2 \leq \frac{2C_2}{\sqrt{2\mu}C_1} \cdot \sqrt{(C_{5a} + C_{5b})\eta \log \frac{2K}{\delta}}$  for any  $k \in [0, K]$ . When  $\eta$  is small enough such that  $\frac{2C_2}{\sqrt{2\mu}C_1} \cdot \sqrt{(C_{5a} + C_{5b})\eta \log \frac{2K}{\delta}} + \eta R/\epsilon < \epsilon_2$ , any  $\phi_k \in \Gamma$  with  $k \geq 0$  will imply  $\phi_{k+1} \in \Gamma$ , since  $\theta_{k+1}$  cannot escape  $\Gamma^{\epsilon_2}$ . Giving  $\phi_0 \in \Gamma$  and using induction, we conclude that all  $\phi_k \in \Gamma$  for  $k \geq 0$ .

When the above holds, the requirement of [Lemma I.4](#) is met. Then with probability  $1 - \delta/2$ , for any  $k \in [0, K - \Delta K]$ ,  $\Delta k \in [\Delta K]$ , we have  $\|\phi_{k+\Delta k} - \phi_k\|_2 \leq C_6 \eta^{0.5-0.5\beta} \sqrt{\log \frac{2}{\eta\delta}}$ .

Finally, we just take the union of [Lemma H.5](#) and [Lemma I.4](#). With  $\log \frac{2K}{\delta} \leq 4 \log \frac{1}{\eta\delta}$  and  $\log \frac{2}{\eta\delta} \leq 2 \log \frac{1}{\eta\delta}$  (which are mild restrictions since  $\eta$  is small), we have the theorem.

We have proved that our iteration will behave well with high probability, but chances still exist that the iteration is driven out of working zones and becomes intractable. We define a well-behaved sequence that manually redirects the iteration when extreme cases happen.

**Definition I.2** (Well-behaved Sequence). *Denote the event of step  $k$  being  $\eta^{2025}$ -good as  $\mathcal{E}_k$ . Let  $\phi_{\text{null}}$  be a fixed point on  $\Gamma$ . Starting from  $\hat{\theta}_0 = \theta_0$  and  $\hat{v}_0 = v_0$ , we define a sequence of  $(\hat{\theta}_k, \hat{v}_k, \hat{m}_k)$  as follows:*

$$\begin{aligned}\hat{m}_{k+1} &:= \beta_1 \hat{m}_k + (1 - \beta_1)(\nabla \mathcal{L}(\hat{\theta}_k) + z_k) \\ \hat{v}_{k+1} &:= \beta_2 \hat{v}_k + (1 - \beta_2)V((\nabla \mathcal{L}(\hat{\theta}_k) + z_k)(\nabla \mathcal{L}(\hat{\theta}_k) + z_k)^\top) \\ \hat{\theta}_{k+1} &:= \mathbf{1}_{\mathcal{E}_k} \theta_{k+1} + \mathbf{1}_{\bar{\mathcal{E}}_k} \phi_{\text{null}},\end{aligned}$$

where  $\mathbf{1}$  is the indicator function and  $\mathbf{1}_{\mathcal{E}} = 1$  if event  $\mathcal{E}$  happens or  $\mathbf{1}_{\mathcal{E}} = 0$  otherwise.

Note that the update of  $\hat{\theta}_k$  can be written as

$$\begin{aligned} \hat{\theta}_{k+1} &:= \hat{\theta}_k - \eta S(\hat{v}_{k+1}) \hat{m}_{k+1} \\ &\quad - \underbrace{\mathbf{1}_{\bar{\mathcal{E}}_k} (\hat{\theta}_k - \eta S(\hat{v}_{k+1}) \hat{m}_{k+1}) + \mathbf{1}_{\bar{\mathcal{E}}_k} \phi_{\text{null}}}_{:= e_k}. \end{aligned}$$

### I.3. Moment Calculation of AGMs Near Manifold

**Additional Notations.** To utilize the analysis framework in Gu et al. [20], we first introduce some notations needed. Consistent with Gu et al. [20], we pretend that AGMs proceed with  $H = \frac{1}{\eta}$  local steps, as a single worker (without multiple workers). We denote every  $H$  steps as one round. Next, we define a “giant step”, which encompasses  $R_{\text{grp}} = \frac{1}{\eta^\beta}$  rounds, corresponding to  $R_{\text{grp}} \cdot H$  steps. We consider a total timescope of  $\frac{T}{\eta^2}$  steps, which corresponds to  $\frac{T}{\eta^{1-\beta}}$  giant steps.

For any  $0 \leq s < R_{\text{grp}}$  and  $0 \leq t \leq H$ , we use  $\hat{\theta}_t^{(s)}$  and  $\hat{\theta}_k$  (where  $k = sH + t$ ) exchangeably to denote the parameter we get on the  $t$ -th local step of round  $s$ , which is also the  $k$ -th global step. Also note that for any  $0 \leq s < R_{\text{grp}}$ ,  $\hat{\theta}_H^{(s)}$  and  $\hat{\theta}_0^{(s+1)}$  refer to the same thing. We define the notation  $\hat{v}_t^{(s)}$ ,  $\hat{m}_t^{(s)}$  and  $\mathcal{E}_t^{(s)}$  in the same way as we did for  $\theta$ .

We further introduce a list of notations:

$$\begin{aligned} \hat{g}_t^{(s)} &:= \nabla \ell_t^{(s)}(\hat{\theta}_t^{(s)}), \hat{S}_k = S(\hat{v}_k), \hat{S}_t^{(s)} := S(\hat{v}_t^{(s)}), \hat{S}^{(s)} := \hat{S}_0^{(s)}, \hat{\phi}^{(s)} := \Phi_{\hat{S}^{(s)}}(\hat{\theta}_0^{(s)}), \\ \hat{x}_t^{(s)} &:= \hat{\theta}_t^{(s)} - \hat{\phi}^{(s)}, \Delta \hat{\phi}^{(s)} := \hat{\phi}^{(s)} - \hat{\phi}^{(0)}, \Sigma_0 := \Sigma(\hat{\phi}^{(0)}), P_{\parallel} := \partial \Phi_{\hat{S}^{(0)}}(\hat{\phi}^{(0)}), P_{\perp} := I - P_{\parallel}, \\ \hat{q}_t^{(s)} &:= \mathbb{E}[\hat{x}_t^{(s)}], \hat{A}_t^{(s)} := \mathbb{E}[\hat{x}_t^{(s)} \hat{x}_t^{(s)\top}], \hat{B}_t^{(s)} := \mathbb{E}[\hat{x}_t^{(s)} \Delta \hat{\phi}^{(s)\top}]. \end{aligned}$$

**Corollary I.1.** *There exist constants  $C_{9a}, C_{9b}, C_{9c}$  such that for all  $0 \leq s < R_{\text{grp}}, 0 \leq t \leq H$ ,*

$$\begin{aligned} \|\hat{x}_t^{(s)}\|_2 &\leq C_{9a} \sqrt{\eta \log \frac{1}{\eta}}, \\ \|\hat{\theta}_t^{(s)} - \hat{\theta}_0^{(s)}\|_2 &\leq C_{9b} \sqrt{\eta \log \frac{1}{\eta}}, \\ \|\hat{\phi}^{(s)} - \hat{\phi}^{(0)}\|_2 &\leq C_{9c} \eta^{0.5-0.5\beta} \sqrt{\log \frac{1}{\eta}}. \end{aligned}$$

**Proof** Substituting  $\delta = \eta^{2025}$ . When  $\mathcal{E}$  holds, this follows directly from the definition of  $\delta$ -goodness; Otherwise, all  $\hat{\theta}$  and  $\hat{\phi}$  are equal, and these quantities are equal to  $\mathbf{0}$ .

**Impact of Momentum.** Our conclusion regarding to the impact of Momentum on the implicit bias is similar to the conclusion in Wang et al. [55]: It does not impact the implicit bias. Further, our analysis is based on moment methods and can give exact error bounds. First, we state some technical lemmas in order to show that introducing momentum will not cause the gradient to deviate too much from itself, i.e.  $\mathbb{E}[\hat{m}_t]$  is close to  $\mathbb{E}[\hat{g}_t]$ . Once this guarantee is established, we can replace  $\hat{m}_t$  with  $\hat{g}_t$  in the moment calculation to simplify it. The general idea of the proof is to show that if  $i$  is close to  $t$ , then  $\mathbb{E}[\nabla \mathcal{L}(\hat{\theta}_{i-1})]$  will become close to  $\mathbb{E}[\nabla \mathcal{L}(\hat{\theta}_{t-1})]$ , and if  $i$  is far from  $t$ , then the contribution of  $\mathbb{E}[\nabla \mathcal{L}(\hat{\theta}_{i-1})]$  would be negligible in  $\mathbb{E}[\hat{m}_t]$ .

**Lemma I.6.** *For any  $k \geq 0$ , we have*

$$\|\mathbb{E}[\nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_{k+1}) - \nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_k)]\|_2 \leq C_{10}\eta^{1.5}$$

for some constant  $C_{10}$ .

**Proof** We have

$$\begin{aligned} \nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_{k+1}) - \nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_k) &= \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_k)(\hat{\boldsymbol{\theta}}_{k+1} - \hat{\boldsymbol{\theta}}_k) + \mathcal{O}\left(\|\hat{\boldsymbol{\theta}}_{k+1} - \hat{\boldsymbol{\theta}}_k\|_2^2\right) \\ &= \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_k)(\hat{\boldsymbol{\theta}}_{k+1} - \hat{\boldsymbol{\theta}}_k) + \mathcal{O}(\eta^2) + \mathcal{O}(\|\mathbf{e}_k\|_2), \end{aligned}$$

since  $\|\hat{\boldsymbol{\theta}}_{k+1} - \hat{\boldsymbol{\theta}}_k\|_2 = \|\eta S(\hat{\mathbf{v}}_k) \hat{\mathbf{m}}_k - \mathbf{e}_{3k}\|_2 = \mathcal{O}(\eta) + \mathcal{O}(\|\mathbf{e}_{3k}\|_2)$ . Let  $\bar{k} = k - \log_{\beta_1}(\eta)$  be a threshold that is logarithmically close to  $k$ , then we have

$$\begin{aligned} \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_k)(\hat{\boldsymbol{\theta}}_{k+1} - \hat{\boldsymbol{\theta}}_k) &= \left( \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_{\bar{k}}) + \mathcal{O}\left(\|\hat{\boldsymbol{\theta}}_k - \hat{\boldsymbol{\theta}}_{\bar{k}}\|_2\right) \right) (\hat{\boldsymbol{\theta}}_{k+1} - \hat{\boldsymbol{\theta}}_k) \\ &= \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_{\bar{k}})(\hat{\boldsymbol{\theta}}_{k+1} - \hat{\boldsymbol{\theta}}_k) + \mathcal{O}(\eta \cdot \log_{\beta_1}(\eta) \cdot \eta) + \mathcal{O}(\|\mathbf{e}_k\|_2) \\ &= \eta \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_{\bar{k}}) S(\hat{\mathbf{v}}_{k+1}) \hat{\mathbf{m}}_{k+1} + \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right) + \mathcal{O}(\|\mathbf{e}_k\|_2). \end{aligned}$$

Recentering the Hessian term to  $\hat{\boldsymbol{\theta}}_{\bar{k}}$  allows us to take  $\mathbb{E}_{\bar{k}}$  on  $S(\hat{\mathbf{v}}_{k+1}) \hat{\mathbf{m}}_{k+1}$ :

$$\mathbb{E} \left[ \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_{\bar{k}}) S(\hat{\mathbf{v}}_{k+1}) \hat{\mathbf{m}}_{k+1} \right] = \mathbb{E} \left[ \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_{\bar{k}}) \mathbb{E}_{\bar{k}} [S(\hat{\mathbf{v}}_{k+1}) \hat{\mathbf{m}}_{k+1}] \right].$$

After that, notice that

$$\begin{aligned} \|\mathbb{E}_{\bar{k}} [S(\hat{\mathbf{v}}_{k+1}) \hat{\mathbf{m}}_{k+1}]\|_2 &= \|\mathbb{E}_{\bar{k}} [S(\mathbb{E}_{\bar{k}}[\hat{\mathbf{v}}_{k+1}]) \hat{\mathbf{m}}_{k+1}]\|_2 + \mathcal{O}(\|\hat{\mathbf{v}}_{k+1} - \mathbb{E}_{\bar{k}}[\hat{\mathbf{v}}_{k+1}]\|_2) \\ &= \|S(\mathbb{E}_{\bar{k}}[\hat{\mathbf{v}}_{k+1}]) \mathbb{E}_{\bar{k}}[\hat{\mathbf{m}}_{k+1}]\|_2 + \mathcal{O}(\|\hat{\mathbf{v}}_{k+1} - \mathbb{E}_{\bar{k}}[\hat{\mathbf{v}}_{k+1}]\|_2) \\ &= \underbrace{\mathcal{O}(\|\mathbb{E}_{\bar{k}}[\hat{\mathbf{m}}_{k+1}]\|_2)}_{D_1} + \underbrace{\mathcal{O}(\|\hat{\mathbf{v}}_{k+1} - \mathbb{E}_{\bar{k}}[\hat{\mathbf{v}}_{k+1}]\|_2)}_{D_2} \end{aligned}$$

since  $S$  and  $\hat{\mathbf{m}}$  are both bounded by constant scale. We figure out the orders of these two terms respectively:

$$\begin{aligned} D_1 &= \|\mathbb{E}_{\bar{k}} \left[ \beta_1^{k-\bar{k}+1} \hat{\mathbf{m}}_{\bar{k}} + (1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} \hat{\mathbf{g}}_i \right]\|_2 \\ &= \mathcal{O}\left(\beta_1^{\log_{\beta_1}(\eta)}\right) + \|\mathbb{E}_{\bar{k}} \left[ (1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} \hat{\mathbf{g}}_i \right]\|_2 \\ &= \mathcal{O}(\eta) + \|\mathbb{E}_{\bar{k}} \left[ (1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} \nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_i) \right]\|_2 \\ &= \mathcal{O}(\eta) + \mathcal{O}(\eta^{0.5}) = \mathcal{O}(\eta^{0.5}) \end{aligned}$$

since  $\nabla \mathcal{L}$  is uniformly bounded by  $\mathcal{O}(\eta^{0.5})$  after convergence (see [Lemma H.5](#)); And

$$\begin{aligned} D_2 &= (1 - \beta_2) \sum_{i=\bar{k}}^k \beta_2^{k-i} \left( V(\hat{\mathbf{g}}_i \hat{\mathbf{g}}_i^\top) - \mathbb{E}_{\bar{k}} \left[ V(\hat{\mathbf{g}}_i \hat{\mathbf{g}}_i^\top) \right] \right) \\ &= \mathcal{O}(b_2 \cdot (k - \bar{k})) \\ &= \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right), \end{aligned}$$

since  $V$  is bounded by a constant scale. Now combining the above together, we have

$$\begin{aligned} \|\mathbb{E}[\nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_{k+1}) - \nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_k)]\|_2 &= \eta \mathbb{E} \left[ \nabla^2 \mathcal{L}(\hat{\boldsymbol{\theta}}_{\bar{k}}) \mathbb{E}_{\bar{k}} [S(\hat{\mathbf{v}}_{k+1}) \hat{\mathbf{m}}_{k+1}] \right] + \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right) \\ &\quad + \mathcal{O}(\mathbb{E}[\|e_k\|_2]) \\ &= \eta \cdot \mathcal{O}(D_1 + D_2) + \mathcal{O}\left(\eta^2 \log \frac{1}{\eta}\right) + \mathcal{O}(\eta^{2025}) \\ &= \mathcal{O}(\eta^{1.5}), \end{aligned}$$

which concludes the proof.

With [Lemma I.6](#), we are ready to deduce the closeness between  $\mathbb{E}[\hat{\mathbf{m}}_k]$  and  $\mathbb{E}[\hat{\mathbf{g}}_k]$ .

**Lemma I.7.** *For any  $k \geq 2 \log_{\beta_1}(\eta)$ , let  $\bar{k} = k - 2 \log_{\beta_1}(\eta)$ , we have*

$$\|\mathbb{E}_{\bar{k}}[\hat{\mathbf{m}}_{k+1} - \hat{\mathbf{g}}_{k+1}]\|_2 \leq C_{11} \eta^{1.5} \log \frac{1}{\eta}, \quad a.s.$$

Note that this also implies that  $\|\mathbb{E}[\hat{\mathbf{m}}_{k+1} - \hat{\mathbf{g}}_{k+1}]\|_2 \leq C_{11} \eta^{1.5} \log \frac{1}{\eta}$ .

**Proof** Expanding  $\mathbb{E}_{\bar{k}}[\hat{\mathbf{m}}_{k+1}]$ , we have

$$\begin{aligned} \mathbb{E}_{\bar{k}}[\hat{\mathbf{m}}_{k+1}] &= \mathbb{E}_{\bar{k}} \left[ (1 - \beta_1) \sum_{i=1}^k \beta_1^{k-i} \hat{\mathbf{g}}_i \right] \\ &= (1 - \beta_1) \sum_{i=1}^{\bar{k}-1} \beta_1^{k-i} \hat{\mathbf{g}}_i + (1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} \mathbb{E}_{\bar{k}}[\hat{\mathbf{g}}_i] \\ &= (1 - \beta_1) \underbrace{\sum_{i=1}^{\bar{k}-1} \beta_1^{k-i} \hat{\mathbf{g}}_i}_{:=E_1} + (1 - \beta_1) \underbrace{\sum_{i=\bar{k}}^k \beta_1^{k-i} \nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_i)}_{:=E_2} \end{aligned}$$

Note that  $E_1$  is negligible:

$$\begin{aligned}
 \|E_1\|_2 &= \|(1 - \beta_1) \sum_{i=1}^{\bar{k}-1} \beta_1^{k-i} \hat{\mathbf{g}}_i\|_2 \\
 &= (1 - \beta_1) \sum_{i=1}^{\bar{k}-1} \beta_1^{k-i} \cdot \mathcal{O}(1) \\
 &\leq (1 - \beta_1) \sum_{i=2 \log_{\beta_1}(\eta)}^{\infty} \beta_1^i \cdot \mathcal{O}(1) \\
 &= \mathcal{O}\left(\beta_1^{2 \log_{\beta_1}(\eta)}\right) = \mathcal{O}(\eta^2),
 \end{aligned}$$

and that  $E_2$  is close to  $\nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_k)$ :

$$\begin{aligned}
 \|E_2 - \mathbb{E}[\nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_k)]\|_2 &= \|(1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} \mathbb{E}[\nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_i)] - \mathbb{E}[\nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_k)]\|_2 \\
 &= \|(1 - \beta_1) \sum_{i=\bar{k}}^k \beta_1^{k-i} \mathbb{E}[\nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_i) - \nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_k)]\|_2 + \mathcal{O}(\eta^2) \\
 &\leq (1 - \beta_1) \cdot (k - \bar{k}) \cdot C_{10} \eta^{1.5} + \mathcal{O}(\eta^2). \quad (\text{by Lemma I.6})
 \end{aligned}$$

Combining the results of  $E_1$  and  $E_2$  gives

$$\begin{aligned}
 \|\mathbb{E}_{\bar{k}}[\hat{\mathbf{m}}_k - \hat{\mathbf{g}}_k]\|_2 &\leq \|E_1\|_2 + \|E_2 - \mathbb{E}[\nabla \mathcal{L}(\hat{\boldsymbol{\theta}}_k)]\|_2 \\
 &\leq (1 - \beta_1) \cdot 2 \log_{\beta_1}(\eta) \cdot C_{10} \eta^{1.5} + \mathcal{O}(\eta^2) \\
 &\leq C_{11} \eta^{1.5} \log \frac{1}{\eta}
 \end{aligned}$$

for some constant  $C_{11}$ , which completes the proof.

### I.3.1. MOMENT CALCULATION WITHIN A GIANT STEP

In this part, we aim to give the change of first and second moments of  $\phi$  and  $\hat{\mathbf{v}}$ , which is the basis of deriving the SDE for AGMs.

Now there are only a few preparations left before we get into the direct part of the moment calculation. For all  $0 \leq s < R_{\text{grp}}, 0 \leq t \leq H$ . Note that  $\|\hat{\mathbf{v}}_{k+1} - \hat{\mathbf{v}}_k\|_2 = (1 - \beta_2) \|V(\hat{\mathbf{g}}_k \hat{\mathbf{g}}_k^\top) - \hat{\mathbf{v}}_k\|_2 = \mathcal{O}(b_2) = \mathcal{O}(\eta^2)$ , so combining with the Lipschitzness of  $S$  gives  $\|\hat{\mathbf{S}}_{k_2} - \hat{\mathbf{S}}_{k_1}\|_2 = \mathcal{O}((k_2 - k_1)\eta^2)$  for any  $k_2 > k_1$  and  $k_2 - k_1 = o(\eta^{-2})$ . Next, we begin our moment calculation analysis starting from the update in one step.

**Lemma I.8.** *For all  $2 \log_{\beta}(\eta) \leq k \leq R_{\text{grp}}H$ , we have*

$$\mathbb{E}[\hat{\boldsymbol{\theta}}_{k+1}] = \mathbb{E}[\hat{\boldsymbol{\theta}}_k - \eta \hat{\mathbf{S}}_0 \hat{\mathbf{g}}_k] + \mathcal{O}(\eta^{2.5-\beta}).$$

**Proof** We write the update rule:

$$\begin{aligned}
 \hat{\theta}_{k+1} &= \hat{\theta}_k - \eta \hat{S}_k \hat{m}_{k+1} - e_k \\
 &= \hat{\theta}_k - \eta \left[ \hat{S}_k \hat{g}_k + \hat{S}_k (\hat{m}_{k+1} - \hat{g}_k) \right] - e_k \\
 &= \hat{\theta}_k - \eta \left[ \hat{S}_0 \hat{g}_k + \underbrace{(\hat{S}_k - \hat{S}_0) \hat{g}_k}_{\Delta \hat{\theta}_1} + \underbrace{\hat{S}_k (\hat{m}_{k+1} - \hat{g}_k)}_{\Delta \hat{\theta}_2} \right] - e_k.
 \end{aligned}$$

We can prove that  $\Delta \hat{\theta}_1$  and  $\Delta \hat{\theta}_2$  are small in expectation. If  $k = 0$  then  $\Delta \hat{\theta}_1 = \mathbf{0}$ ; and if  $k > 0$ , we can decompose  $\mathbb{E} [\Delta \hat{\theta}_1]$  as:

$$\begin{aligned}
 \mathbb{E} [\Delta \hat{\theta}_1] &= \mathbb{E} \left[ (\hat{S}_{k-1} - \hat{S}_0) \hat{g}_k + (\hat{S}_k - \hat{S}_{k-1}) \hat{g}_k \right] \\
 &= \mathbb{E} \left[ (\hat{S}_{k-1} - \hat{S}_0) \nabla \mathcal{L}(\hat{\theta}_k) \right] + \mathbb{E} \left[ (\hat{S}_k - \hat{S}_{k-1}) \hat{g}_k \right] \\
 &= \mathcal{O}((k-1)\eta^2 \cdot \eta^{0.5}) + \mathcal{O}(\eta^2) \\
 &= \mathcal{O}(H \cdot R_{\text{grp}} \cdot \eta^{2.5} + \eta^2) \\
 &= \mathcal{O}(\eta^{1.5-\beta}).
 \end{aligned}$$

Here, the second equality holds because  $z_k$  is conditioned on time  $k$ , when  $\hat{S}_{k-1}$  has already been determined. For  $\Delta \hat{\theta}_2$ , let  $\bar{k} = k - 2 \log_{\beta_1}(\eta)$ , we have

$$\begin{aligned}
 \mathbb{E} [\Delta \hat{\theta}_2] &= \mathbb{E} \left[ \hat{S}_{\bar{k}-1} (\hat{m}_{k+1} - \hat{g}_k) + \mathcal{O} \left( \eta^2 \log \frac{1}{\eta} \right) \right] \\
 &= \mathbb{E} \left[ \hat{S}_{\bar{k}-1} \mathbb{E}_{\bar{k}} [(\hat{m}_{k+1} - \hat{g}_k)] \right] + \mathcal{O} \left( \eta^2 \log \frac{1}{\eta} \right) \\
 &= \mathcal{O} \left( \eta^{1.5} \log \frac{1}{\eta} \right) + \mathcal{O} \left( \eta^2 \log \frac{1}{\eta} \right) \\
 &= \mathcal{O} \left( \eta^{1.5} \log \frac{1}{\eta} \right),
 \end{aligned}$$

where the second-to-last equality follows from Lemma I.7. Finally, we have

$$\begin{aligned}
 \mathbb{E} [\hat{\theta}_{k+1}] &= \mathbb{E} [\hat{\theta}_k - \eta \hat{S}_0 \hat{g}_k] + \mathcal{O}(\eta^{2.5-\beta}) + \mathcal{O} \left( \eta^{2.5} \log \frac{1}{\eta} \right) + \mathcal{O}(\eta^{2025}) \\
 &= \mathbb{E} [\hat{\theta}_k - \eta \hat{S}_0 \hat{g}_k] + \mathcal{O}(\eta^{2.5-\beta}),
 \end{aligned}$$

which concludes the proof.

After getting the update rule of  $\hat{\theta}_k$ , we then derive the moment during the single round with  $H$  steps. To this end, we recap our modification of manifold projection from a ‘‘Gradient Flow’’ manner to ‘‘Preconditioner Flow’’ manner in Definition 2.1.

**Definition I.3** (Pre-conditioner Flow Projection). *Fix a point  $\theta_{null} \notin \Gamma$ . Given a Positive Semi-Definite matrix  $\hat{S}$ . For  $x \in \mathbb{R}^d$ , consider the pre-conditioner flow  $\frac{dx(t)}{dt} = -\hat{S}\nabla\mathcal{L}(x(t))$  with  $x(0) = x$ . We denote the pre-conditioner flow projection of  $x$  as  $\Phi_{\hat{S}}(x)$ , i.e.  $\Phi_{\hat{S}}(x) := \lim_{t \rightarrow +\infty} x(t)$  if the limit exists and belongs to  $\Gamma$ , and  $\Phi_{\hat{S}}(x) = \theta_{null}$  otherwise.*

We decompose the preconditioner matrix in the very begining of the giant step as  $\hat{S}_0 = \hat{S}(\hat{v}_0) = PP$ , where  $P = \hat{S}^{1/2}$ . Then we give the first moment calculation of  $\hat{\phi}$  in the following lemma.

**Lemma I.9.** *The expectation of the change of the manifold projection every round is*

$$\mathbb{E} \left[ \hat{\phi}^{(s+1)} - \hat{\phi}^{(s)} \right] = \begin{cases} -\frac{H\eta^2}{2} \hat{S}_0 \partial^2 \nabla \mathcal{L}(\phi_{(0)}) [P \mathcal{V}_{\nabla^2 \mathcal{L}'(\phi'_{(0)})} (P \Sigma_0 P) P] + \tilde{\mathcal{O}}(\eta^{1.5-\beta}), & R_0 < s < R_{\text{grp}} \\ \tilde{\mathcal{O}}(\eta), & s \leq R_0 \end{cases}$$

where  $R_0 := \max \left\{ \left\lceil \frac{10}{\lambda_{\max} \alpha} \log \frac{1}{\eta} \right\rceil, \left\lceil 2 \log_{1/\beta} \frac{1}{\eta} \right\rceil \right\}$ .

**Proof** First we consider the scenario when  $R_0 < s < R_{\text{grp}}$ . By Lemma I.8, the update rule holds. And we consider an auxiliary process  $\{\hat{\theta}'_t\}$ . Let  $L'(\mathbf{x}) := L(P\mathbf{x})$ , then

$$\begin{aligned} \nabla L'(\mathbf{x}) &= P \nabla L(P\mathbf{x}) \\ \nabla^2 L'(\mathbf{x}) &= P \nabla^2 L(P\mathbf{x}) P \\ \Sigma'(\mathbf{x}) &= P \Sigma(P\mathbf{x}) P \\ \partial^2(\nabla L')(\mathbf{x})[M] &= P \partial^2(\nabla L)(P\mathbf{x})[PMP]. \end{aligned}$$

For an one-step GD update, we have that

$$\begin{aligned} \hat{\theta}'_{t+1} &= \hat{\theta}'_t - \eta \nabla \mathcal{L}'(\hat{\theta}'_t) \\ &= \hat{\theta}'_t - \eta P \nabla \mathcal{L}(P\hat{\theta}'_t). \end{aligned}$$

Similarly, we define  $\mathbf{A}'^{(s)}_t := \mathbb{E}[\mathbf{x}'^{(s)}_t \mathbf{x}'^{(s)\top}_t]$ ,  $\mathbf{q}'^{(s)}_t := \mathbb{E}[\mathbf{x}'^{(s)}_t]$ , and  $\mathbf{B}'^{(s)}_t := \mathbb{E}[\mathbf{x}'^{(s)}_t \Delta \phi'^{(s)\top}_t]$ , where  $\phi(\mathbf{x})$  is the gradient flow projection of point  $\mathbf{x}$ .

Now we are interested in the update of  $P\hat{\theta}'$ , which is

$$P\hat{\theta}'_{t+1} = P\hat{\theta}'_t - \eta \hat{S}_0 \nabla \mathcal{L}(P\hat{\theta}'_t). \quad (4)$$

We now define  $\hat{\theta}'_t := P^{-1} \hat{\theta}_t$ , then combining Equation (4) and Lemma I.8 gives

$$\mathbf{q}'^{(s)}_{t+1} = \mathbf{q}'^{(s)}_{t+1} - \eta \nabla \mathcal{L}'(\hat{\theta}'^{(s)}_t) + \mathcal{O}(\eta^{2.5-\beta}).$$

Now we can apply the results in Lemma I.36 from Gu et al. [20] for the update of  $\hat{\theta}'$ , with loss function  $\mathcal{L}'(\hat{\theta})$ , number of workers  $k = 1$  and manifold projection  $\Phi'(\hat{\theta})$ , which gives

$$\begin{aligned} P \mathbb{E} \left[ \phi'^{(s+1)} - \phi'^{(s)} \right] &= \mathbb{E} \left[ \phi^{(s+1)} - \phi^{(s)} \right] \\ &= -\frac{H\eta^2}{2} PP \partial^2 \nabla \mathcal{L}(\phi_{(0)}) [P \mathcal{V}_{\nabla^2 \mathcal{L}'(\phi'_{(0)})} (P \Sigma_0 P) P] + \tilde{\mathcal{O}}(\eta^{1.5-\beta}), \end{aligned}$$

where the first equation uses the fact that  $P\phi'(\hat{\theta}') = \phi(\hat{\theta})$ , and it can be verified with the definitions of  $\phi$ ,  $\phi'$ , and  $\hat{\theta}'$ .

The proof when  $s \leq R_0$  is a direct conclusion of Lemma I.36 in Gu et al. [20] since the  $R_0 \propto \log \frac{1}{\eta}$  in our case.

**Corollary I.2.** *The expectation of the change of manifold projection every round is:*

$$\mathbb{E} [\phi^{(s+1)} - \phi^{(s)}] = \begin{cases} \frac{H\eta^2}{2} \hat{\mathbf{S}}_0 \partial^2 \Phi_{\hat{\mathbf{S}}_0}(\phi^{(0)}) [\hat{\mathbf{S}}_0 \Sigma_0 \hat{\mathbf{S}}_0] + \tilde{\mathcal{O}}(\eta^{1.5-\beta}), & R_0 < s < R_{\text{grp}} \\ \tilde{\mathcal{O}}(\eta), & s \leq R_0 \end{cases}$$

**Proof** Notice that for the preconditioner projection, we also have the corresponding transformation

$$\begin{aligned} \partial \Phi'(\mathbf{x}') &= \mathbf{P} \partial \Phi_{\hat{\mathbf{S}}}(\mathbf{P}\mathbf{x}) \mathbf{P} \\ \partial^2 \Phi'(\mathbf{x}') [\mathbf{M}] &= \mathbf{P} \partial^2 \Phi(\mathbf{x}) [\mathbf{P}\mathbf{M}\mathbf{P}]. \end{aligned}$$

The above two equations and Lemma I.36 in Gu et al. [20] complete the proof.

**Lemma I.10.** *The second moment of the change of manifold projection every round is*

$$\mathbb{E} [(\phi^{(s+1)} - \phi^{(s)})(\phi^{(s+1)} - \phi^{(s)})^\top] = \begin{cases} H\eta^2 \hat{\mathbf{S}}_0 \mathbf{P}_{\parallel} \hat{\mathbf{S}}_0 \Sigma_0 \hat{\mathbf{S}}_0 \mathbf{P}_{\parallel} \hat{\mathbf{S}}_0 + \tilde{\mathcal{O}}(\eta^{1.5-\beta}), & R_0 < s < R_{\text{grp}} \\ \tilde{\mathcal{O}}(\eta), & s \leq R_0 \end{cases}$$

where  $R_0 := \max \left\{ \left\lceil \frac{10}{\lambda_{\max} \alpha} \log \frac{1}{\eta} \right\rceil, \left\lceil 2 \log_{1/\beta} \frac{1}{\eta} \right\rceil \right\}$ .

**Proof** According to Lemma I.37 in Gu et al. [20], we could write the second moment for  $\hat{\boldsymbol{\theta}}'$  as

$$\mathbb{E} [(\phi'^{(s+1)} - \phi'^{(s)})(\phi'^{(s+1)} - \phi'^{(s)})^\top] = \begin{cases} H\eta^2 \Sigma'_{0,\parallel} + \tilde{\mathcal{O}}(\eta^{1.5-\beta}), & R_0 < s < R_{\text{grp}} \\ \tilde{\mathcal{O}}(\eta), & s \leq R_0. \end{cases}$$

Notice that

$$\begin{aligned} \Sigma'_{0,\parallel} &:= \partial \Phi'(\phi'^{(0)}) \Sigma'_0 \partial \Phi'(\phi'^{(0)}) \\ &= \mathbf{P} \partial \Phi(\phi^{(0)}) \mathbf{P} \mathbf{P} \Sigma_0 \mathbf{P} \mathbf{P} \partial \Phi(\phi^{(0)}) \mathbf{P}. \end{aligned}$$

When  $R_0 \leq s < R_{\text{grp}}$ ,

$$\begin{aligned} \mathbb{E} [(\phi^{(s+1)} - \phi^{(s)})(\phi^{(s+1)} - \phi^{(s)})^\top] &= \mathbb{E} [\mathbf{P}(\phi'^{(s+1)} - \phi'^{(s)})(\phi'^{(s+1)} - \phi'^{(s)})^\top \mathbf{P}] \\ &= \hat{\mathbf{S}}_0 \mathbf{P}_{\parallel} \hat{\mathbf{S}}_0 \Sigma_0 \hat{\mathbf{S}}_0 \mathbf{P}_{\parallel} \hat{\mathbf{S}}_0. \end{aligned}$$

The proof when  $s \leq R_0$  is a direct conclusion of Lemma I.37 in Gu et al. [20] since the  $R_0 \propto \log \frac{1}{\eta}$  in our case.

Then we give the moment change of  $\phi$  within a single giant step.

**Theorem I.1.** *Given  $\|\hat{\boldsymbol{\theta}}^{(0)} - \phi^{(0)}\|_2 = \mathcal{O}(\sqrt{\eta \log \frac{1}{\eta}})$ , for  $0 < \beta < 0.5$ , the first and second moments of  $\Delta \phi^{(R_{\text{grp}})} := \phi^{(R_{\text{grp}})} - \phi^{(0)}$  are as follows:*

$$\begin{aligned} \mathbb{E}[\Delta \phi^{(R_{\text{grp}})}] &= \frac{\eta^{1-\beta}}{2} \hat{\mathbf{S}}_0 \partial^2 \Phi_{\hat{\mathbf{S}}_0}(\phi^{(0)}) [\hat{\mathbf{S}}_0 \Sigma_0 \hat{\mathbf{S}}_0] + \tilde{\mathcal{O}}(\eta^{1.5-2\beta}) + \tilde{\mathcal{O}}(\eta), \\ \mathbb{E}[\Delta \phi^{(R_{\text{grp}})}]^\top &= \eta^{1-\beta} \hat{\mathbf{S}}_0 \Sigma_{\parallel}(\phi^{(0)}, \hat{\mathbf{S}}^{(0)}) \hat{\mathbf{S}}_0 + \tilde{\mathcal{O}}(\eta^{1.5-1.5\beta}) + \tilde{\mathcal{O}}(\eta), \end{aligned}$$

where  $\Sigma_{\parallel}(\phi^{(0)}, \hat{\mathbf{S}}^{(0)}) := \mathbf{P}_{\parallel} \hat{\mathbf{S}}_0 \Sigma_0 \hat{\mathbf{S}}_0 \mathbf{P}_{\parallel}$ .

**Proof** First we prove the first moment change as

$$\begin{aligned}
 \mathbb{E}[\Delta\phi^{(R_{\text{grp}})}] &= \mathbb{E}\left[\sum_{s=0}^{R_{\text{grp}}-1} \phi^{(s+1)} - \phi^{(s)}\right] \\
 &= \sum_{s=0}^{R_0} \mathbb{E}[\phi^{(s+1)} - \phi^{(s)}] + \sum_{s=R_0+1}^{R_{\text{grp}}-1} \mathbb{E}[\phi^{(s+1)} - \phi^{(s)}] \\
 &= \frac{\eta^{1-\beta}}{2} \hat{\mathbf{S}}_0 \partial^2 \Phi_{\hat{\mathbf{S}}_0}(\phi^{(0)})[\hat{\mathbf{S}}_0 \mathbf{\Sigma}_0 \hat{\mathbf{S}}_0] + \tilde{\mathcal{O}}(\eta^{1.5-2\beta}) + \tilde{\mathcal{O}}(\eta).
 \end{aligned}$$

The last equation is a direct conclusion of [Corollary I.2](#).

And for the second moment, we have

$$\begin{aligned}
 \mathbb{E}\left[\left(\sum_{s=0}^{R_{\text{grp}}-1} \phi^{(s+1)} - \phi^{(s)}\right) \left(\sum_{s=0}^{R_{\text{grp}}-1} \phi^{(s+1)} - \phi^{(s)}\right)^\top\right] &= \sum_{s=0}^{R_{\text{grp}}-1} \mathbb{E}[(\phi^{(s+1)} - \phi^{(s)})(\phi^{(s+1)} - \phi^{(s)})^\top] \\
 &\quad + \sum_{s \neq s'} \mathbb{E}[(\phi^{(s+1)} - \phi^{(s)})] \mathbb{E}[(\phi^{(s'+1)} - \phi^{(s')})^\top] \\
 &= \eta^{1-\beta} \hat{\mathbf{S}}_0 \mathbf{\Sigma}_\parallel(\phi^{(0)}, \hat{\mathbf{S}}^{(0)}) \hat{\mathbf{S}}_0 + \tilde{\mathcal{O}}(\eta^{1.5-1.5\beta}) + \tilde{\mathcal{O}}(\eta),
 \end{aligned}$$

where the last equation uses  $\mathbb{E}[(\phi^{(s+1)} - \phi^{(s)})] \mathbb{E}[(\phi^{(s'+1)} - \phi^{(s')})^\top] = \tilde{\mathcal{O}}(\eta^2)$ .

Next, we proceed with the updates of  $v$ .

**Lemma I.11.** *Given  $c := \frac{1-\beta_2}{\eta^2}$ , and we have*

$$\mathbb{E}\left[\hat{\mathbf{v}}_0^{(R_{\text{grp}})} - \hat{\mathbf{v}}_0^{(0)}\right] = c\eta^{1-\beta} \left(V\left(\mathbf{\Sigma}_0^{(0)}\right) - \hat{\mathbf{v}}_0^{(0)}\right) + \mathcal{O}\left(\eta^{1.5-1.5\beta}\right).$$

**Proof** We have

$$\begin{aligned}
 \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} &= \hat{\mathbf{v}}_H^{(s)} - \hat{\mathbf{v}}_0^{(s)} \\
 &= \beta_2^H \hat{\mathbf{v}}_0^{(s)} + (1 - \beta_2) \sum_{i=1}^H \beta_2^{H-i} \mathbf{V}\left(\hat{\mathbf{g}}_i^{(s)} \hat{\mathbf{g}}_i^{(s)\top}\right) - \hat{\mathbf{v}}_0^{(s)} \\
 &= (\beta_2^H - 1) \hat{\mathbf{v}}_0^{(0)} + (1 - \beta_2) \sum_{i=1}^H \beta_2^{H-i} \mathbf{V}\left(\hat{\mathbf{g}}_i^{(s)} \hat{\mathbf{g}}_i^{(s)\top}\right).
 \end{aligned}$$

Note that

$$\begin{aligned}
 \mathbb{E}\left[\hat{\mathbf{g}}_i^{(s)} \hat{\mathbf{g}}_i^{(s)\top}\right] &= \mathbb{E}\left[\mathbf{\Sigma}(\hat{\boldsymbol{\theta}}_i^{(s)})\right] \\
 &= \mathbb{E}\left[\mathbf{\Sigma}(\phi_0^{(0)} + \mathbf{x}_i^{(s)})\right] \\
 &= \mathbb{E}\left[\mathbf{\Sigma}(\phi_0^{(0)}) + \mathcal{O}\left(\eta^{0.5-0.5\beta}\right)\right] \\
 &= \mathbf{\Sigma}_0^{(0)} + \mathcal{O}\left(\eta^{0.5-0.5\beta}\right).
 \end{aligned}$$

Combining with the linearity of  $\mathbf{V}$ , we conclude that

$$\begin{aligned}\mathbb{E} \left[ \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right] &= (\beta_2^H - 1) \hat{\mathbf{v}}_0^{(0)} + (1 - \beta_2^H) \mathbf{V} \left( \boldsymbol{\Sigma}_0^{(0)} \right) + \mathcal{O} \left( \eta^{1.5-0.5\beta} \right) \\ \mathbb{E} \left[ \hat{\mathbf{v}}_0^{(s+1)} \right] &= \beta_2^H \hat{\mathbf{v}}_0^{(s)} + (1 - \beta_2^H) \mathbf{V} \left( \boldsymbol{\Sigma}_0^{(0)} \right) + \mathcal{O} \left( \eta^{1.5-0.5\beta} \right).\end{aligned}$$

To transfer from  $\hat{\mathbf{v}}_0^{(0)}$  to arbitrary  $\hat{\mathbf{v}}_0^{(s)}$ , we simply expand to get the result:

$$\begin{aligned}\mathbb{E} \left[ \hat{\mathbf{v}}_0^{(s)} \right] &= \beta_2^{sH} \hat{\mathbf{v}}_0^{(0)} + \left[ (1 - \beta_2^H) \mathbf{V} \left( \boldsymbol{\Sigma}_0^{(0)} \right) + \mathcal{O} \left( \eta^{1.5-0.5\beta} \right) \right] \left( 1 + \beta_2^H + \beta_2^{2H} + \dots + \beta_2^{(s-1)H} \right) \\ &= \beta_2^{sH} \hat{\mathbf{v}}_0^{(0)} + \left[ (1 - \beta_2^H) \mathbf{V} \left( \boldsymbol{\Sigma}_0^{(0)} \right) \right] \left( \frac{1 - \beta_2^{sH}}{1 - \beta_2^H} \right) + \mathcal{O} \left( \eta^{1.5-0.5\beta} \right) \cdot \mathcal{O} \left( \eta^{-\beta} \right) \\ &= \beta_2^{sH} \hat{\mathbf{v}}_0^{(0)} + (1 - \beta_2^{sH}) \mathbf{V} \left( \boldsymbol{\Sigma}_0^{(0)} \right) + \mathcal{O} \left( \eta^{1.5-1.5\beta} \right).\end{aligned}$$

Thus we have

$$\mathbb{E} \left[ \hat{\mathbf{v}}_0^{(R_{\text{grp}})} - \hat{\mathbf{v}}_0^{(0)} \right] = c\eta^{1-\beta} \left( \mathbf{V} \left( \boldsymbol{\Sigma}_0^{(0)} \right) - \hat{\mathbf{v}}_0^{(0)} \right) + \mathcal{O} \left( \eta^{1.5-1.5\beta} \right).$$

where the last equation uses the fact that  $1 - \beta_2^{R_{\text{grp}}H} = 1 - (1 - c\eta^{1-\beta}) + O(\eta^{2-2\beta}) = c\eta + O(\eta^2)$ .

Also, for the second moment change of  $\hat{\mathbf{v}}$ , we get the following lemma

**Lemma I.12.** *The second moment change of  $\hat{\mathbf{v}}$  over a giant step is*

$$\mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(R_{\text{grp}})} - \hat{\mathbf{v}}_0^{(0)} \right) \left( \hat{\mathbf{v}}_0^{(R_{\text{grp}})} - \hat{\mathbf{v}}_0^{(0)} \right)^\top \right] = \mathcal{O}(\eta^{2-\beta}).$$

**Proof**

$$\begin{aligned}\mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right) \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right)^\top \right] &= \mathbb{E} \left[ \left( (\beta_2^H - 1) + (1 - \beta_2) \sum_{i=1}^H \beta_2^{H-i} \mathbf{V} \left( \hat{\mathbf{g}}_i^{(s)} \hat{\mathbf{g}}_i^{(s)\top} \right) \right) \right. \\ &\quad \left. \left( (\beta_2^H - 1) + (1 - \beta_2) \sum_{i=1}^H \beta_2^{H-i} \mathbf{V} \left( \hat{\mathbf{g}}_i^{(s)} \hat{\mathbf{g}}_i^{(s)\top} \right) \right)^\top \right] \\ &= \mathcal{O} \left( (1 - \beta_2^H)^2 \right) = \mathcal{O} \left( \eta^2 \right).\end{aligned}$$

$$\begin{aligned}\mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(R_{\text{grp}})} - \hat{\mathbf{v}}_0^{(0)} \right) \left( \hat{\mathbf{v}}_0^{(R_{\text{grp}})} - \hat{\mathbf{v}}_0^{(0)} \right)^\top \right] &= \mathbb{E} \left[ \left( \sum_{s=0}^{R_{\text{grp}}-1} \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right) \right) \left( \sum_{s=0}^{R_{\text{grp}}-1} \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right)^\top \right) \right] \\ &= \sum_{s=0}^{R_{\text{grp}}-1} \mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right) \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right)^\top \right] \\ &\quad + \sum_{s \neq s'} \mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right) \right] \mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(s'+1)} - \hat{\mathbf{v}}_0^{(s')} \right)^\top \right] \\ &= \mathcal{O}(\eta^{2-\beta}).\end{aligned}$$

The last equation uses

$$\mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right) \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right)^\top \right] = \mathcal{O}(\eta^2),$$

and

$$\mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(s+1)} - \hat{\mathbf{v}}_0^{(s)} \right) \right] \mathbb{E} \left[ \left( \hat{\mathbf{v}}_0^{(s'+1)} - \hat{\mathbf{v}}_0^{(s')} \right)^\top \right] = \mathcal{O}(3 - 3\beta).$$

The above equation completes the proof.

#### I.4. Weak Approximation

After we get the first and second moment changes within a giant step, we now utilize the moment calculation to prove [Theorem 2.1](#). First, we recall our slow SDE for AGMs

$$\begin{cases} d\zeta(t) = P_{\zeta, \mathbf{S}(t)} \left( \Sigma_{\parallel}^{1/2}(\zeta(t); \mathbf{S}(t)) d\mathbf{W}_t - \frac{1}{2} \mathbf{S}(t) \nabla^3 \mathcal{L}(\zeta) [\Sigma_{\diamond}(\zeta(t); \mathbf{S}(t))] dt \right), \\ d\mathbf{v}(t) = c(V(\Sigma(\zeta)) - \mathbf{v}) dt. \end{cases}$$

We then open the projection mapping  $P_{\zeta, \mathbf{S}(t)}$  as

$$\begin{cases} d\zeta = \partial \Phi_{\mathbf{S}(\mathbf{v})}(\zeta) \mathbf{S}(\mathbf{v}) \Sigma^{1/2}(\zeta) d\mathbf{W}_t + \frac{1}{2} \partial^2 \Phi_{\mathbf{S}(\mathbf{v})}(\zeta) [\mathbf{S}(\mathbf{v}) \Sigma(\zeta) \mathbf{S}(\mathbf{v})] dt, \\ d\mathbf{v}(t) = c(V(\Sigma(\zeta)) - \mathbf{v}) dt. \end{cases} \quad (5)$$

Now it suffices to prove the SDE in [Equation \(5\)](#) tracks the trajectory in AGMs within  $\mathcal{O}(\frac{1}{\eta^2})$  steps in a weak approximation sense.

First, we have to show that the solution of [Equation \(5\)](#) is close in the minimizer manifold

**Lemma I.13.** *Let  $\mathbf{X}(t) := (\zeta(t)^\top, \mathbf{v}(t)^\top)^\top$  be the solution of [Equation \(5\)](#) with  $\zeta(0) \in \Gamma$ , and  $\mathbf{v}(0) \in \mathbb{R}^d$ , then we have that  $\zeta(t) \in \Gamma$  for all  $t \geq 0$ .*

**Proof** According to Du and Duan [11], Filipović [15], for a closed manifold  $\mathcal{M}$  to be viable for the SDE  $d\mathbf{X}(t) = \mathbf{A}(\mathbf{X}(t))d\mathbf{W}_t + \mathbf{b}(\mathbf{X}(t))dt$ , where  $\mathbf{A}(\cdot) : \mathbb{R}^{2d} \rightarrow \mathbb{R}^{2d \times 2d}$  and  $\mathbf{b}(\cdot) : \mathbb{R}^{2d} \rightarrow \mathbb{R}^{2d}$  are locally Lipschitz, it suffices to show that the following Nagumo type consistency condition holds:

$$\mu(\mathbf{x}) := \mathbf{b}(\mathbf{x}) - \frac{1}{2} \sum_j D[A_j(\mathbf{x})] A_j(\mathbf{x}) \in T_{\mathbf{x}}(\mathcal{M}), \quad A_j(\mathbf{x}) \in T_{\mathbf{x}}(\mathcal{M}).$$

Following the argument in Gu et al. [20], here we also only need to show that  $\mathbf{P}_{\perp}(\mathbf{x})\mu(\mathbf{x}) = 0$ , where  $\mathbf{P}_{\perp}(\mathbf{x}) := \mathbf{I}_d - \partial \Phi_{\mathbf{I}}(\mathbf{x})$ .  $\Phi_{\mathbf{I}}(\mathbf{x})$  is also the gradient flow projection at point  $\mathbf{x}$ .

$$\begin{aligned} \mathbf{P}_{\perp}(\mathbf{x}) \sum_j D[A_j(\mathbf{x})] A_j(\mathbf{x}) &= \mathbf{P}_{\perp}(\mathbf{x}) \sum_j D[\partial \Phi_{\mathbf{S}}(\mathbf{x}) \mathbf{S} \Sigma^{1/2}] \partial \Phi_{\mathbf{S}}(\mathbf{x}) \mathbf{S} \Sigma^{1/2} \\ &= \mathbf{P}_{\perp}(\mathbf{x}) \mathbf{S} \sum_j \partial^2 \Phi_{\mathbf{S}}(\mathbf{x}) [\Phi_{\mathbf{S}}(\mathbf{x}) \mathbf{S} \Sigma^{1/2}, \Sigma^{1/2}] \\ &= -\mathbf{P}_{\perp}(\mathbf{x}) \mathbf{S} \nabla^2 \mathcal{L}(\mathbf{x})^\dagger \partial^2(\nabla \mathcal{L})(\mathbf{x}) [\Sigma_{\parallel}(\mathbf{x}, \mathbf{S})]. \end{aligned}$$

The last equation uses [Lemma I.3](#). Again, applying [Lemma I.3](#) gives

$$\mathbf{P}_{\perp}(\mathbf{x}) \mathbf{b}(\mathbf{x}) = -\frac{1}{2} \mathbf{P}_{\perp}(\mathbf{x}) \mathbf{S} \nabla^2 \mathcal{L}(\mathbf{x})^\dagger \partial^2(\nabla \mathcal{L})(\mathbf{x}) [\Sigma_{\parallel}(\mathbf{x}, \mathbf{S})].$$

The above equation completes the proof.

To establish [Theorem 2.1](#), we give an equivalent theorem, which capture the closeness of  $\mathbf{X}(t)$  and  $\bar{\mathbf{X}}_t$  in a long horizon. Also, for the proof of [Theorem 2.1](#), it suffices to prove the following lemma, whose proof would be shown at the end of [Appendix I.4](#).

**Theorem I.2** (Adaption of Theorem I.4 in Gu et al. [20]). *If  $\|\boldsymbol{\theta}^{(0)} - \boldsymbol{\phi}^{(0)}\|_2 = \mathcal{O}(\sqrt{\eta \log \frac{1}{\eta}})$  and  $\zeta(0) = \boldsymbol{\phi}^{(0)}$ ,  $\mathbf{v}(0) = \mathbf{v}^{(0)}$ , then for a giant step  $R_{\text{grp}} = \lfloor \frac{1}{\eta^{0.25}} \rfloor$ , for every test function  $g \in \mathcal{C}^3$ ,*

$$\max_{0 \leq n \leq \lfloor \frac{T}{\eta^{0.75}} \rfloor} \left| \mathbb{E}[g(\bar{\mathbf{X}}^{(nR_{\text{grp}})})] - \mathbb{E}[g(\mathbf{X}(n\eta^{0.75}))] \right| = C_g \eta^{0.25} (\log \frac{1}{\eta})^b,$$

where  $C_g$  is a constant independent of  $\eta$  but depends on  $g(\cdot)$  and  $b > 0$  is a universal constant independent of  $g(\cdot)$  and  $\eta$ .

### I.5. Preliminary and Additional Notations

We first introduce some notations and preliminary background. We consider the following stochastic gradient algorithms (SGAs)

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \eta_e \mathbf{h}(\mathbf{x}_n, \boldsymbol{\xi}_n),$$

where  $\mathbf{x}_n \in \mathbb{R}^{2d}$  is the parameter vector,  $\eta_e$  is the effective learning rate,  $\mathbf{h}(\cdot, \cdot) : \mathbb{R}^{2d} \times \mathbb{R}^{2d} \rightarrow \mathbb{R}^{2d}$  depend on the current parameter vector  $\mathbf{x}_n$  and the noise vector  $\boldsymbol{\xi}_n$  sampled from some distribution  $\Xi(\mathbf{x}_n)$ .

We also consider the Stochastic Differential Equation (SDE) of the following form:

$$d\mathbf{X}_t = \mathbf{b}(\mathbf{X}_t, t)dt + \sigma(\mathbf{X}_t, t)d\mathbf{W}_t,$$

where  $\mathbf{b} : \mathbb{R}^{2d} \times \mathbb{R}^+ \rightarrow \mathbb{R}^{2d}$  is the drift vector function and  $\sigma : \mathbb{R}^{2d} \times \mathbb{R}^+ \rightarrow \mathbb{R}^{2d \times 2d}$  is the diffusion matrix function.

According to the moment calculations in [Corollary I.2](#), [Lemma I.10](#), [Lemma I.11](#), and [Lemma I.12](#), we set  $\eta_e = \eta^{1-\beta}$ , and

$$\begin{aligned} \mathbf{b}(\mathbf{X}_t, t) &= \left( \left( \frac{1}{2} \partial^2 \Phi_{\mathbf{S}(\mathbf{v})}(\boldsymbol{\zeta}) [\boldsymbol{\Sigma}(\boldsymbol{\zeta}, \mathbf{S}(\mathbf{v}))] \right)^\top, c(V(\boldsymbol{\Sigma}(\boldsymbol{\zeta})) - \mathbf{v})^\top \right)^\top, \\ \sigma(\mathbf{X}_t, t) &= \begin{pmatrix} \partial \Phi_{\mathbf{S}(\mathbf{v})}(\boldsymbol{\zeta}) \boldsymbol{\Sigma}^{1/2}(\boldsymbol{\zeta}, \mathbf{S}(\mathbf{v})), & \mathbf{0} \\ \mathbf{0}, & \mathbf{0} \end{pmatrix}. \end{aligned}$$

Next, we are going to define the one giant step change of the parameter, both for SGAs and SDE.

$$\begin{aligned} \hat{\mathbf{X}}^{(lR_{\text{grp}})} &:= (\Phi_{\hat{\mathbf{S}}^{(lR_{\text{grp}})}}(\hat{\boldsymbol{\theta}}, \mathbf{v}^{lR_{\text{grp}}}), \quad \Delta^{(n)} := \hat{\mathbf{X}}^{((n+1)R_{\text{grp}})} - \hat{\mathbf{X}}^{(nR_{\text{grp}})}, \\ \tilde{\Delta}^{(n)} &:= \mathbf{X}_{(n+1)\eta_e} - \hat{\mathbf{X}}^{(nR_{\text{grp}})}, \quad \mathbf{b}^{(n)} := \mathbf{b}(\hat{\mathbf{X}}^{(nR_{\text{grp}})}), \quad \sigma^{(n)} := \sigma(\hat{\mathbf{X}}^{(nR_{\text{grp}})}). \end{aligned}$$

We now give a lemma to give the approximation of the first, second, and higher-order moment change of the SDE.

**Lemma I.14.** *There exists a positive constant  $c_0$  independent of  $\eta_e$  and  $g$  such that for all  $\zeta \in \Gamma$ , it holds for all  $1 \leq i \leq d$  that*

$$\begin{aligned} \left| \mathbb{E}[\tilde{\Delta}_i(\zeta, n)] - \eta_e b_i(\zeta) \right| &\leq c_0 \eta_e^2, \\ \left| \mathbb{E}[\tilde{\Delta}_i(\zeta, n) \tilde{\Delta}_j(\zeta, n)] - \eta_e \sum_{l=1}^d \sigma_{i,l}(\zeta) \sigma_{l,j}(\zeta) \right| &\leq c_0 \eta_e^2 \\ \mathbb{E} \left[ \left| \prod_{s=1}^6 \tilde{\Delta}_{i_s}(\zeta, n) \right| \right] &\leq c_0 \eta_e^3. \end{aligned}$$

**Proof** (i) By Lemma I.13, the first half solution  $\zeta(t)$  in  $\mathbf{X}(t)$  of Equation (5) stays in the manifold almost surely when  $\zeta(0) \in \Gamma$ . (ii) We assume that  $\mathcal{L} \in \mathcal{C}^5$ , so  $\mathbf{b}, \sigma \in \mathcal{C}^4$ . (iii) We know that  $\Gamma$  is compact by Assumption 2.2. Then we can directly apply Lemma B.3 in Malladi et al. [41] and Lemma 26 in Li et al. [33].

**Lemma I.15** (Adaption of Lemma I.41 in Gu et al. [20]). *Given drift term and diffusion term  $\mathbf{b}, \sigma \in G^\alpha$  and Lipschitz. Let  $s \in [0, T]$  and  $g \in G^\alpha$ . Then for  $t \in [s, T]$ , we can define:*

$$u(\mathbf{x}, s, t) := \mathbb{E}_{\mathbf{X}_t \sim \mathcal{P}_X(\mathbf{x}, s, t)} [g(\mathbf{X}_t)].$$

where  $\mathcal{P}_X(\mathbf{x}, s, t)$  denotes the distribution of  $\mathbf{X}_t$  with the initial condition  $\mathbf{X}(s) = \mathbf{x}$ . Then  $u(\cdot, s, t) \in G^\alpha$  uniformly in  $s, t$ .

## I.6. Proof of the Approximation for Slow SDE of AGMs

For the giant step constant  $\beta \in (0, 0.5)$ , we define several quantities  $a_1 = \frac{1.5-2\beta}{1-\beta} \in (1, 1.5)$ ,  $a_2 = \frac{1}{1-\beta} \in (1, 2)$ ,  $a_3 = \frac{2-2\beta}{1-\beta} = 2$ ,  $a_4 = \frac{1.5-1.5\beta}{1-\beta} = 1.5$ . In this part, we will show that only  $a_1$  and  $a_2$  would impact the error bound in our approximation theorem.

The following lemma captures the difference between the SDEs' and the AGMs' first and second moment changes, as a key step to control the approximation error, utilizing the moment calculation results from the last section.

**Lemma I.16.** *If  $\|\theta^{(0)} - \phi^{(0)}\|_2 = \mathcal{O}(\sqrt{\eta_e \log \frac{1}{\eta_e}})$ , then it holds for all  $0 \leq n \leq \lfloor T/\eta_e \rfloor$  and  $1 \leq i \leq d$  that*

$$\begin{aligned} \left| \mathbb{E}[\Delta_i^{(n)} - \tilde{\Delta}_i^{(n)} \mid \mathcal{E}_0^{(nR_{\text{grp}})}] \right| &\leq c_1 \left( \eta_e^{a_1} \left( \log \frac{1}{\eta_e} \right)^b + \eta_e^{a_2} \left( \log \frac{1}{\eta_e} \right)^b \right), \\ \left| \mathbb{E}[\Delta_i^{(n)} \Delta_j^{(n)} - \tilde{\Delta}_i^{(n)} \tilde{\Delta}_j^{(n)} \mid \mathcal{E}_0^{(nR_{\text{grp}})}] \right| &\leq c_1 \left( \eta_e^{a_1} \left( \log \frac{1}{\eta_e} \right)^b + \eta_e^{a_2} \left( \log \frac{1}{\eta_e} \right)^b \right), \end{aligned}$$

$$\begin{aligned} \mathbb{E} \left[ \left| \prod_{s=1}^6 \Delta_{i_s}^{(n)} \mid \mathcal{E}^{(nR_{\text{grp}})} \right| \right] &\leq c_1^2 \eta_e^{2a_1} \left( \log \frac{1}{\eta_e} \right)^{2b}, \\ \mathbb{E} \left[ \left| \prod_{s=1}^6 \tilde{\Delta}_{i_s}^{(n)} \mid \mathcal{E}^{(nR_{\text{grp}})} \right| \right] &\leq c_1^2 \eta_e^{2a_1} \left( \log \frac{1}{\eta_e} \right)^{2b}, \end{aligned}$$

where  $c_1$  and  $b$  are constants independent of  $\eta_e$  and  $g$ .

**Proof** According to [Appendix I.2](#), we have that

$$\mathbb{E} \left[ \left| \prod_{s=1}^6 \Delta_{i_s}^{(n)} \mid \mathcal{E}^{(nR_{\text{grp}})} \right| \right] = \mathcal{O}(\eta^{3-3\beta}).$$

We can further use [Corollary I.2](#), [Lemma I.10](#), [Lemma I.11](#), and [Lemma I.12](#), which gives

$$\left| \mathbb{E}[\Delta_i^{(n)} - \eta_e b_i^{(n)}] \right| \leq c_2 \left( \eta_e^{a_1} \left( \log \frac{1}{\eta_e} \right)^b + \eta_e^{a_2} \left( \log \frac{1}{\eta_e} \right)^b \right), \quad (6)$$

$$\left| \mathbb{E}[\Delta_i^{(n)} \Delta_j^{(n)} - \eta_e \sum_{l=1}^d \sigma_{i,l}^{(n)} \sigma_{l,j}^{(n)}] \right| \leq c_2 \left( \eta_e^{a_1} \left( \log \frac{1}{\eta_e} \right)^b + \eta_e^{a_2} \left( \log \frac{1}{\eta_e} \right)^b \right) \quad (7)$$

$$\mathbb{E} \left[ \left| \prod_{s=1}^6 \Delta_{i_s}^{(n)} \right| \right] \leq c_2^2 \eta_e^{2a_1} \left( \log \frac{1}{\eta_e} \right)^{2b}. \quad (8)$$

Notice that the above equations uses  $a_1 < a_3$  and  $a_2 < a_4$  for all  $\beta \in (0, 0.5)$ . These three equations and [Lemma I.14](#) give the Lemma.

**Lemma I.17.** *For a test function  $g \in \mathcal{C}^3$ , and we define  $u_{l,n}(\mathbf{x}) := u(\mathbf{x}, l\eta_e, n\eta_e) = \mathbb{E}_{\mathbf{X}_t \sim \mathcal{P}(\mathbf{x}, l\eta_e, n\eta_e)}[g(\mathbf{X}_t)]$ . If  $\|\boldsymbol{\theta}^{(0)} - \phi^{(0)}\|_2 = \mathcal{O}(\sqrt{\eta \log \frac{1}{\eta}})$ , then for all  $0 \leq l \leq n-1$ , and  $1 \leq n \leq \lfloor T/\eta_e \rfloor$ , it holds that*

$$\left| \mathbb{E}[u_{l+1,n}(\bar{\mathbf{X}}^{(lR_{\text{grp}})} + \Delta^{(l)}) - u_{l+1,n}(\bar{\mathbf{X}}^{(lR_{\text{grp}})} + \tilde{\Delta}^{(l)}) \mid \bar{\mathbf{X}}^{(lR_{\text{grp}})}] \right| \leq C_{g,3}(\eta_e^{a_1} + \eta_e^{a_2}) \log\left(\frac{1}{\eta_e}\right)^b,$$

where  $C_{g,3}$  is some positive constant independent of  $\eta_e$  but can depend on  $g$ .

**Proof** Given  $g \in \mathcal{C}^3$ , by [Lemma I.15](#), we have  $u_{l,n}(\mathbf{x}) \in \mathcal{C}^3$  for all  $l$  and  $n$ . Which is to say that there exists a function  $Q(\cdot) \in G$ , such that the partial derivative of  $u_{l,n}(\mathbf{X})$  with respect to  $l, n, \mathbf{x}$  up to the third order is bounded by  $Q(\mathbf{x})$ . By the law of total expectation and triangle inequality,

$$\begin{aligned} & \left| \mathbb{E}[u_{l+1,n}(\hat{\mathbf{X}}^{(lR_{\text{grp}})} + \Delta^{(l)}) - u_{l+1,n}(\hat{\mathbf{X}}^{(lR_{\text{grp}})} + \tilde{\Delta}^{(l)}) \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}] \right| \\ & \leq \underbrace{\left| \mathbb{E}[u_{l+1,n}(\hat{\mathbf{X}}^{(lR_{\text{grp}})} + \Delta^{(l)}) - u_{l+1,n}(\hat{\mathbf{X}}^{(lR_{\text{grp}})} + \tilde{\Delta}^{(l)}) \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})}] \right|}_{I_1} \\ & \quad + \underbrace{\eta^{2025} \mathbb{E} \left[ \left| u_{l+1,n}(\hat{\mathbf{X}}^{(lR_{\text{grp}})} + \Delta^{(l)}) \right| \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})} \right]}_{I_2} \\ & \quad + \underbrace{\eta^{2025} \mathbb{E} \left[ \left| u_{l+1,n}(\hat{\mathbf{X}}^{(lR_{\text{grp}})} + \tilde{\Delta}^{(l)}) \right| \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})} \right]}_{I_3}. \end{aligned}$$

For  $I_2$  and  $I_3$ , due to the compactness of  $\Gamma$  and  $\mathbf{v} \preceq R_1$  from [Assumption F.3](#),  $Q(\mathbf{x})$  can be bounded for some constant  $C_{g,4}$  independent of  $\eta_e$  but could depend on test function  $g$ . Hence, we have that  $I_2 + I_3 \leq C_{g,4}\eta^{2025}$ .

Using the triangle inequality, we first decompose  $I_1$  into several terms as

$$\begin{aligned}
 I_1 &\leq \underbrace{\sum_{i=1}^d \left| \mathbb{E} \left[ \frac{\partial u_{l,n}}{\partial X_i}(\hat{\mathbf{X}}^{(R_{\text{grp}})}) \left( \Delta_i^{(l)} - \tilde{\Delta}_i^{(l)} \right) \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})} \right] \right|}_{I_{1,1}} \\
 &\quad + \underbrace{\frac{1}{2} \sum_{1 \leq i, j \leq d} \left| \mathbb{E} \left[ \frac{\partial^2 u_{l,n}}{\partial X_i \partial X_j}(\hat{\mathbf{X}}^{(R_{\text{grp}})}) \left( \Delta_j^{(l)} \Delta_i^{(l)} - \tilde{\Delta}_i^{(l)} \tilde{\Delta}_j^{(l)} \right) \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})} \right] \right|}_{I_{1,2}} \\
 &\quad + |\mathcal{R}| + |\tilde{\mathcal{R}}|,
 \end{aligned}$$

where the third order remainders  $\mathcal{R}$  and  $\tilde{\mathcal{R}}$  are

$$\begin{aligned}
 \mathcal{R} &= \frac{1}{6} \sum_{1 \leq i, j, k \leq d} \left| \mathbb{E} \left[ \frac{\partial^3 u_{l,n}}{\partial X_i \partial X_j \partial X_k}(\hat{\mathbf{X}}^{(R_{\text{grp}})} + \alpha \Delta^{(l)}) \left( \Delta_j^{(l)} \Delta_i^{(l)} \Delta_k^{(l)} \right) \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})} \right] \right| \\
 \tilde{\mathcal{R}} &= \frac{1}{6} \sum_{1 \leq i, j, k \leq d} \left| \mathbb{E} \left[ \frac{\partial^3 u_{l,n}}{\partial X_i \partial X_j \partial X_k}(\hat{\mathbf{X}}^{(R_{\text{grp}})} + \tilde{\alpha} \tilde{\Delta}^{(l)}) \left( \tilde{\Delta}_j^{(l)} \tilde{\Delta}_i^{(l)} \tilde{\Delta}_k^{(l)} \right) \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})} \right] \right|,
 \end{aligned}$$

where  $\alpha, \tilde{\alpha} \in (0, 1)$ . Again, notice that the  $\Gamma$  is compact and  $vv \preceq R_1$ , thus we can bound the derivatives of  $u_{l,n}(\mathbf{x})$  for any  $\mathbf{X}$  as

$$\left| \frac{\partial u_{l+1,n}}{\partial \mathbf{X}_i}(\mathbf{X}) \right| \leq C_{g,4}, \quad \left| \frac{\partial^2 u_{l+1,n}}{\partial \mathbf{X}_i \partial \mathbf{X}_j}(\mathbf{X}) \right| \leq C_{g,4}, \quad \left| \frac{\partial^3 u_{l+1,n}}{\partial \mathbf{X}_i \partial \mathbf{X}_j \partial \mathbf{X}_k}(\mathbf{X}) \right| \leq C_{g,4}. \quad (9)$$

For the term  $I_{1,1}$  and  $I_{1,2}$ , by applying [Lemma I.16](#), we have that

$$I_{1,1} \leq dc_1 C_{g,4} (\eta_e^{a_1} + \eta_e^{a_2}) (\log \frac{1}{\eta_e})^b, \quad I_{1,2} \leq \frac{d^2}{2} c_1 C_{g,4} (\eta_e^{a_1} + \eta_e^{a_2}) (\log \frac{1}{\eta_e})^b.$$

Next, we bound the remainders  $\mathcal{R}$  and  $\tilde{\mathcal{R}}$ . By Cauchy-Schwarz inequality,

$$\begin{aligned}
 |\mathcal{R}| &\leq \frac{1}{6} \sum_{1 \leq i, j, k \leq d} \sqrt{\mathbb{E} \left[ \left( \frac{\partial^3 u_{l,n}}{\partial X_i \partial X_j \partial X_k}(\hat{\mathbf{X}}^{(R_{\text{grp}})} + \alpha \Delta^{(l)}) \right)^2 \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})} \right]} \times \\
 &\quad \sqrt{\mathbb{E} \left[ \left( \Delta_j^{(l)} \Delta_i^{(l)} \Delta_k^{(l)} \right)^2 \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(lR_{\text{grp}})} \right]} \\
 &\leq \frac{d^3}{6} C_{g,4} c_1 \eta_e^{a_1} \log \left( \frac{1}{\eta_e} \right)^b,
 \end{aligned}$$

where the last inequality uses [Lemma I.16](#) and [Equation \(9\)](#).

Similarly, we can prove that there exists a positive constant  $C_{g,5}$  such that

$$|\tilde{\mathcal{R}}| \leq \frac{d^3}{6} C_{g,5} c_1 \eta_e^{a_1} \log \left( \frac{1}{\eta_e} \right)^b.$$

Combining the bounds for  $I_1$ ,  $I_2$ , and  $I_3$  gives the lemma.

Finally, we are ready to prove [Theorem I.2](#).

**Proof** [Proof of [Theorem I.2](#)] For  $0 \leq l \leq n = \lfloor \frac{T}{\eta^{0.75}} \rfloor$ , we denote the random variable by  $\hat{\mathbf{x}}_{l,n}$  such that follows a distribution  $\mathcal{P}_{\mathbf{X}}(\hat{\mathbf{X}}^{(lR_{\text{grp}})}, l\eta_e, n\eta_e)$ . When we set  $l = n$ ,  $\mathcal{P}(\hat{\mathbf{x}}_{n,n} = \hat{\mathbf{X}}^{(nR_{\text{grp}})})$  and setting  $l = 0$  gives  $\hat{\mathbf{x}}_{0,n} \sim \mathbf{X}(n\eta_e)$ . Recall the previous definition that  $u(\mathbf{x}, s, t) = \mathbb{E}_{\mathbf{X}_t \sim \mathcal{P}_{\mathbf{X}}(\mathbf{x}, s, t)}[g(\mathbf{X}_t)]$ , and we define that  $\mathcal{T}_{l+1,n} := u_{l+1,n}(\hat{\mathbf{X}}^{(lR_{\text{grp}})} + \Delta^{(l)}, (l+1)\eta_e, n\eta_e) - u_{l+1,n}(\hat{\mathbf{X}}^{(lR_{\text{grp}})} + \tilde{\Delta}^{(l)}, (l+1)\eta_e, n\eta_e)$ . Using the definition of  $\mathbf{x}_{l,n}$ , we can rewrite the distance between AGMs and SDE measured by a test function  $g$  as

$$\begin{aligned} & \left| \mathbb{E} \left[ g(\bar{\mathbf{X}}^{(nR_{\text{grp}})}) - g(\mathbf{X}(n\eta_e)) \right] \right| \\ & \leq \left| \mathbb{E} \left[ g(\mathbf{x}_{n,n}) - g(\mathbf{x}_{0,n}) \mid \mathcal{E}_0^{(nR_{\text{grp}})} \right] \right| + \mathcal{O}(\eta^{2025}). \end{aligned}$$

The above equation uses the law of total expectation and the definition of  $\delta$ -good event  $\mathcal{E}_0^{(nR_{\text{grp}})}$  in [Definition I.1](#). Then the Triangle inequality gives

$$\begin{aligned} \left| \mathbb{E} \left[ g(\mathbf{x}_{n,n}) - g(\mathbf{x}_{0,n}) \mid \mathcal{E}_0^{(nR_{\text{grp}})} \right] \right| & \leq \sum_{l=0}^{n-1} \left| \mathbb{E} \left[ g(\hat{\mathbf{x}}_{l+1,n}) - g(\hat{\mathbf{x}}_{l,n}) \mid \mathcal{E}_0^{(nR_{\text{grp}})} \right] \right| + \mathcal{O}(\eta^{2025}) \\ & = \sum_{l=0}^{n-1} \left| \mathbb{E} \left[ \mathcal{T}_{l+1,n} \mid \mathcal{E}_0^{(nR_{\text{grp}})} \right] \right| + \mathcal{O}(\eta^{2025}) \\ & = \sum_{l=0}^{n-1} \left| \mathbb{E} \left[ \mathbb{E} \left[ \mathcal{T}_{l+1,n} \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(nR_{\text{grp}})} \right] \mid \mathcal{E}_0^{(nR_{\text{grp}})} \right] \right| + \mathcal{O}(\eta^{2025}) \\ & \leq \sum_{l=0}^{n-1} \mathbb{E} \left[ \left| \mathbb{E} \left[ \mathcal{T}_{l+1,n} \mid \hat{\mathbf{X}}^{(lR_{\text{grp}})}, \mathcal{E}_0^{(nR_{\text{grp}})} \right] \right| \mid \mathcal{E}_0^{(nR_{\text{grp}})} \right] + \mathcal{O}(\eta^{2025}) \\ & \leq nC_{g,3}(\eta_e^{a_1} + \eta_e^{a_2}) \log\left(\frac{1}{\eta_e}\right)^b \\ & \leq TC_{g,3}(\eta_e^{a_1-1} + \eta_e^{a_2-1}) \log\left(\frac{1}{\eta_e}\right)^b. \end{aligned}$$

where the second last inequality uses [Lemma I.17](#). Recap that  $a_1 = \frac{1.5-2\beta}{1-\beta}$ ,  $a_2 = \frac{1}{1-\beta}$ ,  $\beta \in (0, 0.5)$ . Let  $\beta = 0.25$ , and we complete the proof.

## Appendix J. Proof of Theorems in [Appendix D](#)

### J.1. Proof of Adam and Adame- $\lambda$ 's Implicit Bias with Label Noise

In this part, we give the proof of [Theorem D.1](#) and [Theorem D.2](#).

**Proof** [Proof of [Theorem D.1](#)] With label noise, the gradient covariance matrix  $\Sigma(\zeta) = \alpha \nabla^2 \mathcal{L}(\zeta)$  for any  $\zeta \in \Gamma$ .

We recall the SDE formula in [Equation \(5\)](#) and [Lemma I.9](#)

$$\begin{cases} d\zeta(t) = \partial\Phi_{\mathbf{S}(\mathbf{v})}(\zeta) \mathbf{S}(\mathbf{v}) \Sigma^{1/2}(\zeta) d\mathbf{W}_t - \frac{1}{2} \mathbf{S}_t \partial^2(\nabla \mathcal{L})(\zeta) [\mathbf{P} \mathcal{V}_{\nabla^2 \mathcal{L}'(\phi'_{(0)})}(\mathbf{P} \Sigma_0 \mathbf{P}) \mathbf{P}] dt, \\ d\mathbf{v}(t) = c(V(\Sigma(\zeta)) - \mathbf{v}) dt. \end{cases}$$

And  $\mathbf{P} := \mathbf{S}_0^{1/2}$ . By applying [Lemma I.2](#), and replacing  $\Sigma$  with  $\alpha \nabla^2 \mathcal{L}$ , we can reduce the SDE formula as

$$\begin{cases} d\zeta(t) = -\frac{\alpha}{2} \mathbf{S}_t \partial^2(\nabla \mathcal{L})(\zeta)[\mathbf{S}_t] dt, \\ d\mathbf{v}(t) = c(V(\Sigma(\zeta)) - \mathbf{v}) dt. \end{cases}$$

Then we can write out the constraint for the fixed point  $(\zeta^*, \mathbf{v}^*)$  of this ODE as

$$-\frac{1}{2} \mathbf{S}(\mathbf{v}^*) \partial^2(\nabla \mathcal{L})(\zeta^*)[S(\mathbf{v}^*)] = 0, \quad (10)$$

$$V(\Sigma(\zeta^*)) - \mathbf{v}^* = 0. \quad (11)$$

Solving [Equation \(10\)](#) and [Equation \(11\)](#) gives

$$\partial^2(\nabla \mathcal{L})(\zeta^*)[S(V(\Sigma(\zeta^*)))] = 0. \quad (12)$$

Integrating by parts gives us

$$\partial^2(\nabla \mathcal{L})[\mathbf{S}] = \nabla [\langle \nabla^2 \mathcal{L}, \mathbf{S} \rangle] - \nabla(\mathbf{S}) [\nabla^2 \mathcal{L}]. \quad (13)$$

We use  $\mathbf{H}$  and  $\nabla^2 \mathcal{L}$  interchangeably to denote the Hessian matrix. For the first term, note that

$$\begin{aligned} \langle \mathbf{S}, \mathbf{H} \rangle &= \sum_{j,k} [\mathbf{S}]_{jk} \mathbf{H}_{jk} \\ &= \sum_{i,j,k} \mathbf{P}_{ji} \mathbf{H}_{jk} \mathbf{P}_{ki} \\ &= \sum_i [\mathbf{P} \mathbf{H} \mathbf{P}]_{ii} \\ &= \text{tr}(\mathbf{P} \mathbf{H} \mathbf{P}) \\ &= \text{tr}\left((\text{Diag} \mathbf{H})^{\frac{1}{2}}\right) + O(\epsilon \text{tr}(\mathbf{H})), \end{aligned}$$

where the last equality comes from the update rule of Adam:  $\mathbf{S} = (\text{Diag} \mathbf{H})^{-\frac{1}{2}} + O(\epsilon)$ ,  $\mathbf{P} = (\text{Diag} \mathbf{H})^{-\frac{1}{4}} + O(\sqrt{\epsilon})$ . For the second term, we also plug in the update rule of Adam, and use  $h_j$  to denote  $\mathbf{H}_{jj}$ , which turns out to be the gradient of the same thing:

$$\begin{aligned} \nabla(\mathbf{S}) [\nabla^2 \mathcal{L}] &= \sum_{j,k} \nabla([\mathbf{S}]_{jk}) \nabla_{jk}^2 \mathcal{L} \\ &= \sum_j \nabla([\mathbf{S}]_{jj}) \nabla_{jj}^2 \mathcal{L} \\ &= \sum_j \nabla\left(h_j^{-\frac{1}{2}}\right) \cdot h_j + \mathcal{O}\left(\epsilon \sum_j h_j\right) \\ &= \sum_j \nabla(h_j) \cdot -\frac{1}{2} h_j^{-\frac{1}{2}} + O(\epsilon \text{tr}(\mathbf{H})) \\ &= \sum_j \nabla\left(-h_j^{\frac{1}{2}}\right) + O(\epsilon \text{tr}(\mathbf{H})) \\ &= -\nabla \text{tr}\left((\text{Diag} \mathbf{H})^{\frac{1}{2}}\right) + O(\epsilon \text{tr}(\mathbf{H})). \end{aligned}$$

Summarizing, our drift term can be represented as a constant multiple of

$$\mathbf{S} \nabla \text{tr} \left( (\text{Diag} \mathbf{H})^{\frac{1}{2}} \right) + O(\mathbf{S} \epsilon \text{tr}(\mathbf{H})), \quad (14)$$

forming a preconditioned gradient flow that implicitly minimizes  $\text{tr} \left( (\text{Diag} \mathbf{H})^{\frac{1}{2}} \right)$  when  $\epsilon \rightarrow 0$ . Combining Equation (12) and Eq. (14) gives the result in Theorem D.1.

**Proof** [Proof of Theorem D.2] Now we consider the optimizer AdamE- $\lambda$ , the variant of Adam proposed as a verification case of our main results, whose update rule is

$$\begin{aligned} \mathbf{m}_{k+1} &:= \beta_1 \mathbf{m}_k + (1 - \beta_1) \nabla \ell_k(\boldsymbol{\theta}_k) \\ \mathbf{v}_{k+1} &:= \beta_2 \mathbf{v}_k + (1 - \beta_2) \nabla \ell_k(\boldsymbol{\theta}_k)^{\odot 2} \\ \theta_{k+1,i} &:= \theta_{k,i} - \eta \frac{m_{k+1,i}}{(v_{k+1,i})^\lambda + \epsilon} \quad \text{for all } i \in [d], \lambda \in (0, 1). \end{aligned}$$

Now for AdamE- $\lambda$ , the precondition matrix  $\mathbf{S} = (\text{Diag} \mathbf{H})^{-\lambda}$ , and  $\mathbf{P} = (\text{Diag} \mathbf{H})^{-\lambda/2}$ , which gives

$$\begin{aligned} \langle \mathbf{S}, \mathbf{H} \rangle &= \sum_{j,k} [\mathbf{S}]_{jk} \mathbf{H}_{jk} \\ &= \sum_{i,j,k} \mathbf{P}_{ji} \mathbf{H}_{jk} \mathbf{P}_{ki} \\ &= \sum_i [\mathbf{P} \mathbf{H} \mathbf{P}]_{ii} \\ &= \text{tr}(\mathbf{P} \mathbf{H} \mathbf{P}) \\ &= \text{tr} \left( (\text{Diag} \mathbf{H})^{1-\lambda} \right) + O(\epsilon \text{tr}(\mathbf{H})). \end{aligned}$$

Also, similar to the case for Adam, we have

$$\begin{aligned} \nabla(\mathbf{S}) [\nabla^2 \mathcal{L}] &= \sum_{j,k} \nabla \left( [\mathbf{S}]_{jk} \right) \nabla_{jk}^2 \mathcal{L} \\ &= \sum_j \nabla \left( [\mathbf{S}]_{jj} \right) \nabla_{jj}^2 \mathcal{L} \\ &= \sum_j \nabla \left( h_j^{-\lambda} \right) \cdot h_j + \mathcal{O} \left( \epsilon \sum_j h_j \right) \\ &= \sum_j \nabla(h_j) \cdot -\lambda h_j^{-\lambda} \\ &= -\frac{\lambda}{1-\lambda} \sum_j \nabla \left( -h_j^{1-\lambda} \right) + O(\epsilon \text{tr}(\mathbf{H})) \\ &= -\frac{\lambda}{1-\lambda} \nabla \text{tr} \left( (\text{Diag} \mathbf{H})^{1-\lambda} \right) + O(\epsilon \text{tr}(\mathbf{H})). \end{aligned}$$

Utilizaing Equation (13), we get the regularization term for AdamE- $\lambda$  as

$$\begin{aligned} \partial^2(\nabla \mathcal{L})[\mathbf{S}] &= \nabla [\langle \nabla^2 \mathcal{L}, \mathbf{S} \rangle] - \nabla(\mathbf{S}) [\nabla^2 \mathcal{L}] \\ &= \frac{1}{1-\lambda} \text{tr} \left( (\text{Diag} \mathbf{H})^{1-\lambda} \right) + O(\epsilon \text{tr}(\mathbf{H})), \end{aligned}$$

Evaluating the above equation when  $\epsilon \rightarrow 0$  completes the proof.

## J.2. Proof of Lemma D.1

**Proof** Recall that the manifold is defined as  $\Gamma = \{\boldsymbol{\theta} | \langle \mathbf{z}_i, \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2} \rangle = y_i, \forall i \in [n]\}$ . So if any  $\boldsymbol{\theta} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}$  belongs to  $\Gamma$ , and another  $\tilde{\boldsymbol{\theta}} = \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{v}} \end{pmatrix}$  satisfies that  $\tilde{u}_i^{\odot 2} - \tilde{v}_i^{\odot 2} = u_i^{\odot 2} - v_i^{\odot 2}$  for any  $i \in [d]$ , then  $\tilde{\boldsymbol{\theta}}$  also belongs to  $\Gamma$ .

Next, we derive the explicit expression of the Hessian matrix when  $\boldsymbol{\theta} \in \Gamma$ :

$$\begin{aligned} \nabla^2 \mathcal{L}(\boldsymbol{\theta}) &= \frac{2}{n} \sum_{i=1}^n 2 \begin{pmatrix} \mathbf{z}_i \odot \mathbf{u} \\ -\mathbf{z}_i \odot \mathbf{v} \end{pmatrix} \begin{pmatrix} \mathbf{z}_i \odot \mathbf{u} \\ -\mathbf{z}_i \odot \mathbf{v} \end{pmatrix}^\top + (\langle \mathbf{z}_i, \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2} \rangle - y_i) \begin{pmatrix} \text{diag}(\mathbf{z}) & 0 \\ 0 & -\text{diag}(\mathbf{z}) \end{pmatrix} \\ &= \frac{4}{n} \sum_{i=1}^n \begin{pmatrix} \mathbf{z}_i \odot \mathbf{u} \\ -\mathbf{z}_i \odot \mathbf{v} \end{pmatrix} \begin{pmatrix} \mathbf{z}_i \odot \mathbf{u} \\ -\mathbf{z}_i \odot \mathbf{v} \end{pmatrix}^\top. \end{aligned}$$

Hence, we have that

$$\text{tr}(\text{Diag}(\mathbf{H})^{e_0}) \propto \sum_{i=1}^d (|u_i|^{2e_0} + |v_i|^{2e_0}),$$

and  $\|\mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}\|_{e_0}^{e_0} = \sum_{i=1}^d |u_i^2 - v_i^2|^{e_0}$ . Let  $e_0 \in (0, 1]$ , and we assume that

$$\boldsymbol{\theta} \in \arg \min_{\boldsymbol{\theta}' \in \Gamma} \text{tr}(\text{Diag}(\mathbf{H})^{e_0}) = \arg \min_{\boldsymbol{\theta}' \in \Gamma} \sum_{i=1}^d (|u_i|^{2e_0} + |v_i|^{2e_0}).$$

First we prove by contradiction that  $u_i = 0$  or  $v_i = 0$  for any  $i \in [d]$ . If there exists some  $i$  such that  $u_i \neq 0$  and  $v_i \neq 0$ , then denote  $s = \min\{|u_i|, |v_i|\}$ , we construct  $\tilde{\boldsymbol{\theta}} = \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{v}} \end{pmatrix}$  by letting  $\tilde{u}_j = u_j, \tilde{v}_j = v_j$  for  $j \neq i$  and  $\tilde{u}_i = |u_i| - s, \tilde{v}_i = |v_i| - s$ . Then  $\tilde{\mathbf{u}}^{\odot 2} - \tilde{\mathbf{v}}^{\odot 2} = \mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}$ , so  $\tilde{\boldsymbol{\theta}} \in \Gamma$ , but  $|\tilde{u}_i|^{2e_0} + |\tilde{v}_i|^{2e_0} < |u_i|^{2e_0} + |v_i|^{2e_0}$ , a contradiction.

Now assume  $\boldsymbol{\theta} \notin \arg \min_{\boldsymbol{\theta}' \in \Gamma} \|\mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}\|_{e_0}$ . There must exist some  $\tilde{\boldsymbol{\theta}} \in \Gamma$  such that  $\|\tilde{\mathbf{u}}^{\odot 2} - \tilde{\mathbf{v}}^{\odot 2}\|_{e_0} < \|\mathbf{u}^{\odot 2} - \mathbf{v}^{\odot 2}\|_{e_0}$ . WLOG assume for any  $i \in [d]$ , either  $\tilde{u}_i = 0$  or  $\tilde{v}_i = 0$ , else we can construct another minimizer that preserves  $\|\tilde{\mathbf{u}}^{\odot 2} - \tilde{\mathbf{v}}^{\odot 2}\|_{e_0}$  as above. But now we have  $\sum_{i=1}^d |u_i^2 - v_i^2|^{e_0} = \sum_{i=1}^d |u_i|^{2e_0} + |v_i|^{2e_0}$ , and  $\sum_{i=1}^d |\tilde{u}_i^2 - \tilde{v}_i^2|^{e_0} = \sum_{i=1}^d |\tilde{u}_i|^{2e_0} + |\tilde{v}_i|^{2e_0}$ , which indicates that  $\sum_{i=1}^d |\tilde{u}_i|^{2e_0} + |\tilde{v}_i|^{2e_0} < \sum_{i=1}^d |u_i|^{2e_0} + |v_i|^{2e_0}$ , a contradiction.

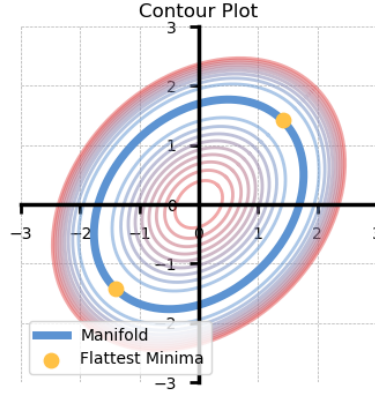


Figure 1

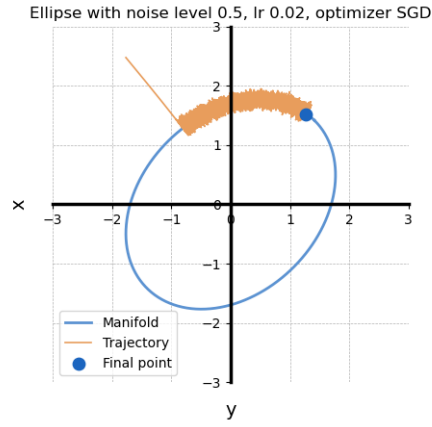


Figure 2

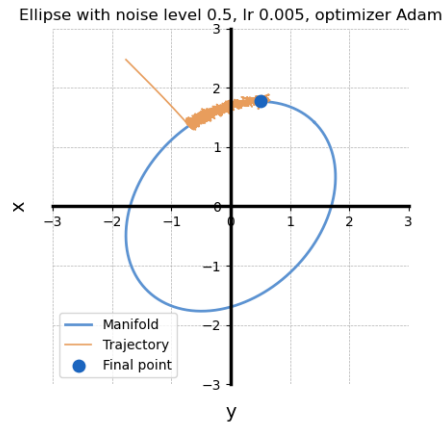


Figure 3

Figure 4: **(a)**: Coutour of the elliptical loss, from which we can see the two tips as the flattest minima. **(b)**: SGD implicitly minimizes  $\text{tr}(\mathbf{H})$  and converges to the flattest minima. **(c)**: Adam reduces sharpness too but converges to a different and sparser minimum.

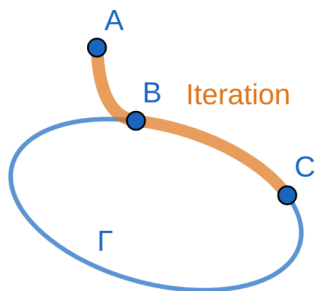


Figure 5

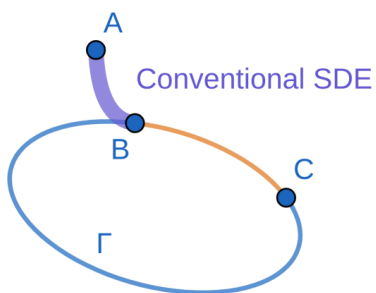


Figure 6

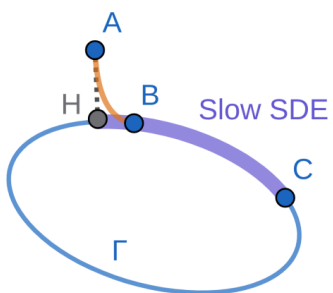


Figure 7

Figure 8: Comparison of conventional SDE and slow SDE.

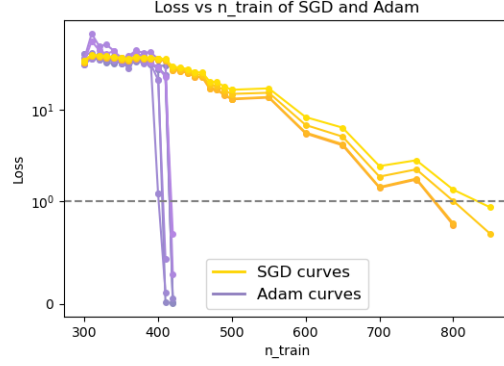


Figure 9

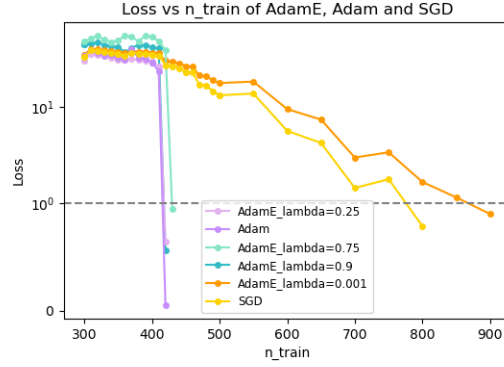


Figure 10

Figure 11: The curve of final test loss vs. scale of training data with  $d = 10000$ ,  $\kappa = 50$ . **(a)**: Loss comparison between SGD with different learning rates, and Adam with different learning rates and  $\beta_2$ 's. **(b)**: Loss comparison between AdamE with  $\lambda = 0.001, 0.25, 0.75, 0.9$ , Adam and SGD.

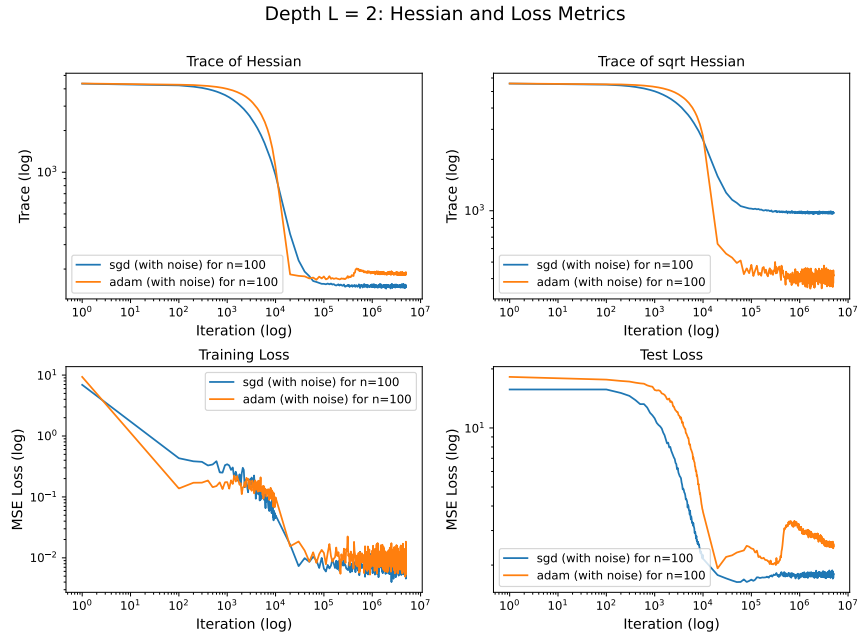


Figure 12: **Deep matrix factorization with label noise.** Adam and SGD are trained on identical data and noise realizations. *Top:* evolution of  $\text{tr}(\mathbf{H})$  and  $\text{tr}(\text{Diag}(\mathbf{H})^{1/2})$ . *Bottom:* training and test MSE. Adam converges to a point with larger overall curvature but smaller diagonal curvature, and exhibits higher test error.