
Analyzing Sharpness along GD Trajectory: Progressive Sharpening and Edge of Stability

Zhouzi Li*

IIIS, Tsinghua University
zhouzi188763@gmail.com

Zixuan Wang*

IIIS, Tsinghua University
wangzx2019012326@gmail.com

Jian Li†

IIIS, Tsinghua University
lapordge@gmail.com

Abstract

Recent findings demonstrate that modern neural networks trained by full-batch gradient descent typically enter a regime called Edge of Stability (EOS). In this regime, the sharpness, i.e., the maximum Hessian eigenvalue, first increases to the value $2/(\text{step size})$ (the progressive sharpening phase) and then oscillates around this value (the EOS phase). This paper aims to analyze the GD dynamics and the sharpness along the optimization trajectory. Our analysis naturally divides the GD trajectory into four phases depending on the change in the sharpness value. We empirically identify the norm of output layer weight as an interesting indicator of the sharpness dynamics. Based on this empirical observation, we attempt to theoretically and empirically explain the dynamics of various key quantities that lead to the change of the sharpness in each phase of EOS. Moreover, based on certain assumptions, we provide a theoretical proof of the sharpness behavior in the EOS regime in two-layer fully-connected linear neural networks. We also discuss some other empirical findings and the limitation of our theoretical results.

1 Introduction

Deep learning has achieved great success in a variety of machine learning applications, and gradient-based algorithms are the prevailing optimization methods for training deep neural networks. However, mathematically understanding the behavior of the optimization methods for deep learning is highly challenging, due to non-convexity, over-parameterization, and complicated architectures. In particular, some recent empirical findings in deep networks contradict the traditional understandings of gradient methods. For example, Wu et al. [30] observed that the solution found by gradient descent has sharpness approximately equal to $2/\eta$ instead of just being smaller than $2/\eta$. Also, Jastrzebski et al. [14] observed that there is a break-even point in the SGD trajectory, and after this point, there is a regularization effect on the loss curvature.

One recent well-known example is the phenomenon called “Edge of Stability” (EOS) (Cohen et al. [6]). Based on the classical optimization theory, the learning rate η of gradient-based method should be smaller than $2/\lambda$ so that the loss can decrease, where λ is the largest eigenvalue of the Hessian of the objective, also called “sharpness” in the literature. Otherwise, the loss diverges (even for simple quadratic functions). However, the empirical findings in Cohen et al. [6] show that under various

*Contributed equally, listed in alphabetical order.

†The authors are supported in part by the National Natural Science Foundation of China Grant 62161146004, Turing AI Institute of Nanjing and Xi’an Institute for Interdisciplinary Information Core Technology.

network settings, the EOS phenomena typically occurs along the gradient descent trajectory: (1) the sharpness first increases until it reaches $2/\eta$ (called “progressive sharpening”) (2) the sharpness starts hovering around $2/\eta$ (the EOS regime) and (3) the loss non-monotonically decreases without diverging.

Although (1) seems to be consistent with the traditional beliefs about optimization, a rigorous mathematical explanation for it is still open. Moreover, phenomena (2) and (3) are more mysterious because they violate the $\eta < 2/\lambda$ “rule” in traditional optimization theory, yet the training loss does not completely diverge. Instead, the loss may oscillate but still decrease in the long run, while the sharpness seems to be restrained from further increasing.

In this paper, we aim to provide a theoretical and empirical explanation for the mystery of EOS. Towards the goal, we focus on the dynamics of these key quantities when EOS happens and attempt to find out the main driving force to explain these phenomena along the gradient descent trajectory from both theoretical and empirical perspectives.

1.1 Our Contributions

Our contributions can be summarized as follows.

(Section 3.1) We analyze the typical sharpness behavior along the gradient descent trajectory when EOS happens, and propose a four-phase division of GD trajectory, based on the dynamics of some key quantities such as the loss and the sharpness, for further understanding this phenomenon.

(Section 3.2) We empirically identify the weight norm of the output layer as an effective indicator of the sharpness dynamics. We show that analyzing the dynamics of this surrogate can qualitatively explain the dynamics of sharpness. By assuming this relation, together with some additional simplifying assumptions and approximations, we can explain the dynamics of the sharpness, the loss, and the output layer norm in each phase of EOS (Section 3.3). In this context, we also offers an interesting explanation for the non-monotonic loss decrement (also observed in Cohen et al. [6], Xing et al. [32]) (Section 3.4).

(Section 4) Following similar ideas, we provide a more rigorous proof for the progressive sharpening and EOS phenomena in a two-layer fully-connected linear neural network setting based on certain assumptions. The assumptions made here are either weaker or arguably less restrictive.

1.2 Related work

The structure of Hessian The Hessian matrix carries the second order information of the loss landscape. Several prior works have empirically found that the spectrum of Hessian has several “outliers” and a continuous “bulk” (Sagun et al. [28, 29], Pappas [25, 26]). Typically, each outlier corresponds to one class in multi-class classification. As we consider the binary classification setting, there is typically one outlier (i.e., the largest eigenvalue) that is much larger than other eigenvalues. It is consistent with our Assumption 4.1. The Gauss-Newton decomposition of the Hessian was used in several prior works (Martens [23], Bottou et al. [4], Pappas [25, 26]). Pappas [25] empirically showed that the outliers of Hessian can be attributed to a “G component”, which is also known as Fisher Information Matrix (FIM) in Karakida et al. [15, 16]. Also, Wu et al. [31] analyzed the leading Hessian eigenspace by approximating the Hessian with Kronecker factorization and theoretically proved the outliers structure under some random setting assumption.

Neural Tangent Kernel A recent line of work studied the learning of over-parameterized neural networks in the so-called. “neural tangent kernel (NTK) regime or the lazy training regime (Jacot et al. [13], Lee et al. [18], Du et al. [8, 7], Arora et al. [2], Chizat et al. [5]). A main result in this regime is that if the neural network is wide enough, gradient flow can find the global optimal empirical minimizer very close to the initialization. Moreover, the Hessian does not change much in the NTK regime. Our findings go beyond NTK setting to analyze the change of sharpness.

Edge of Stability regime The Edge of Stability phenomena was first formalized by Cohen et al. [6]. Similar phenomena were also identified in Jastrzebski et al. [14] as the existence of the “break-even” point on SGD trajectory after which loss curvature gets regularized. Xing et al. [32] observed that gradient descent eventually enters a regime where the iterates oscillate on the leading curvature direction and the loss drops non-monotonically. Recently Ahn et al. [1] studied the non-monotonic

decreasing behavior of GD which they called unstable convergence, and discussed the possible causes of this phenomenon. Ma et al. [22] proposed a special subquadratic landscape property and proved that EOS occurs based on this assumption. Arora et al. [3] studied the implicit bias on the sharpness of deterministic gradient descent in the EOS regime. They proved in some specific settings with a varying learning rate (called normalized GD) or with a modified loss \sqrt{L} , gradient descent enters EOS and further reduces sharpness. They mainly focus on the analysis near the manifold of minimum loss, but our analysis also applies to the early stage of the training when the loss is not close to the minimum. In particular, our analysis provides an explanation of non-monotonic loss decrease that cannot be explained by their theory. Another difference is that they consider \sqrt{L} (for constant learning rate) where L is a fairly general MSE loss independent of any neural network structure, while our analysis is strongly tied with the MSE loss of a neural network. Very recently, Lyu et al. [21] explained how GD enters EOS for normalized loss (e.g., neural networks with normalization layers), and analyzed the sharpness reduction effect along the training trajectory. The notion of sharpness in their work is somewhat different due to normalization. In particular, they consider the so-called *spherical sharpness*, that is the sharpness of the normalized weight vector. They also mainly studied the regime where the parameter is close to the manifold of minimum loss as in [3] and proved that GD approximately tracks a continuous sharpness-reduction flow. Lewkowycz et al. [19] proposed a similar regime called “catapult phase” where loss does not diverge even if the largest Hessian eigenvalue is larger than $2/\eta$. Our work mainly considers training in this regime and assumes that the training is not in the “divergent phase” in Lewkowycz et al. [19]. Compared with Lewkowycz et al. [19], we provide a more detailed analysis in more general settings along gradient descent trajectory.

2 Preliminaries

Notations: We denote the training dataset as $\{\mathbf{x}_i, y_i\}_{i=1}^n \subset \mathbb{R}^d \times \{1, -1\}$ and the neural network as $f : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}$. The network $f(\boldsymbol{\theta}, \mathbf{x})$ maps the input $\mathbf{x} \in \mathbb{R}^d$ and parameter $\boldsymbol{\theta} \in \mathbb{R}^p$ to an output in \mathbb{R} . In this paper, we mainly consider the case of binary classification with mean square error (MSE) loss $\ell(z, y) = (z - y)^2$.

Denote the input matrix as $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) \in \mathbb{R}^{d \times n}$ and the label vector as $\mathbf{Y} = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n$. We let $\mathbf{F}(t) = (f(\boldsymbol{\theta}(t), \mathbf{x}_1), f(\boldsymbol{\theta}(t), \mathbf{x}_2), \dots, f(\boldsymbol{\theta}(t), \mathbf{x}_n)) \in \mathbb{R}^n$ and $\mathbf{D}(t) = \mathbf{F}(t) - \mathbf{Y}$ be the (output) prediction vector, and the residual vector respectively at time t . The training objective is: $\mathcal{L}(f(\boldsymbol{\theta})) = \frac{1}{n} \sum_{i=1}^n \ell(f(\boldsymbol{\theta}, \mathbf{x}_i), y_i) = \frac{1}{n} \sum_{i=1}^n (f(\boldsymbol{\theta}, \mathbf{x}_i) - y_i)^2$.

Hessian, Fisher information matrix and NTK: In this part, we apply previous works to show that the largest eigenvalue of Hessian is almost the same as the largest eigenvalue of NTK. We use the latter as the definition of **the sharpness** in this paper. Further details can be found in Appendix F.

As shown in Papyan [26], Martens [23], Bottou et al. [4], the Hessian can be decomposed into two components, where the term known as “Gauss-Newton matrix”, G-term or Fisher information matrix (FIM), dominates the second term in terms of the largest eigenvalue. Meanwhile, Karakida et al. [16] pointed out the duality between the FIM and a Gram matrix \mathbf{M} , defined as $\mathbf{M} = \frac{2}{n} \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}^\top$. It is also known as the neural tangent kernel NTK (Karakida et al. [16, 15]), which has been studied extensively in recent years (see e.g., [13],[8],[2],[5]). Note that in this paper, we do not assume the training is in NTK regime, in which the Hessian does not change much during training. It is not hard to see that \mathbf{M} and FIM share the same non-zero eigenvalues: if $\mathbf{G}\mathbf{u} = \lambda\mathbf{u}$ for some eigenvector $\mathbf{u} \in \mathbb{R}^p$, $\mathbf{M} \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{u} = \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{G}\mathbf{u} = \lambda \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{u}$, i.e., λ is also an eigenvalue of \mathbf{M} .

In this paper, we use $\boldsymbol{\theta}(t)$ to denote the parameter at iteration t (or time t) and the sharpness at time t as $\Lambda(t) = \Lambda(\boldsymbol{\theta}(t))$. We similarly define $\mathbf{M}(t)$, $\mathbf{F}(t)$, $\mathbf{D}(t)$, $\mathcal{L}(t)$.

Here we show the gradient flow dynamics of the residual vector $\mathbf{D}(t)$:

$$\frac{d\mathbf{D}(t)}{dt} = \frac{\partial \mathbf{D}(t)}{\partial \boldsymbol{\theta}} \frac{d\boldsymbol{\theta}(t)}{dt} = -\frac{\partial \mathbf{F}(t)}{\partial \boldsymbol{\theta}} \frac{\partial \mathcal{L}(t)}{\partial \boldsymbol{\theta}} = -\frac{2}{n} \frac{\partial \mathbf{F}(t)}{\partial \boldsymbol{\theta}} \frac{\partial \mathbf{F}(t)}{\partial \boldsymbol{\theta}}^\top \mathbf{D}(t) = -\mathbf{M}(t)\mathbf{D}(t) \quad (1)$$

3 A Four-phase Analysis of GD Dynamics

In this section, we study the dynamics of gradient descent and the change of sharpness along the optimization trajectory. We divide the whole training process into four phases, occurring repeatedly

in the EOS regime. In Section 3.1, we introduce the four phases. In Section 3.2, we show empirically that the change of the norm of the output layer weight vector almost coincides with the change of the sharpness. In Section 3.3, using this observation, we attempt to explain the dynamics of each phase and provide a mathematical explanation for the changes in the sharpness. In Section 3.4, we explain why the loss decreases but non-monotonically. We admit that a completely rigorous theoretical explanation is still beyond our reach and much of our argument is based on various simplifying assumptions and is somewhat heuristic at some points. Due to space limits, we defer all the proofs in this section to Appendix E.1.

3.1 A Four-phase Division

To further understand the properties along the trajectory when EOS happens, we study the behaviors of the loss and the sharpness during the training process. As illustrated in Figure 1, we train a shallow neural network by gradient descent on a subset of 1,000 samples from CIFAR-10 (Krizhevsky et al. [17]), using the MSE loss as the objective. Notice that the sharpness keeps increasing while the loss decreases until the sharpness reaches $2/\eta$. Then the sharpness begins to oscillate around $2/\eta$ while the loss decreases non-monotonically. This is a typical sharpness behavior in the EOS regime, and consistent with the experiments in [6].

We divide the training process into four phases according to the evolution of the loss, the sharpness, and their correlation, as shown in Figure 1. The four phases happen cyclically along the training trajectory. We first briefly describe the properties of each phase and explain the dynamics in more detail in Section 3.3.

Phase I: Sharpness $\Lambda < 2/\eta$. In this stage, all the eigenvalues of Gram matrix M are below the threshold $2/\eta$. In particular, using standard initialization, the training typically starts from this phase, and during this phase the loss keeps decreasing and the sharpness keeps growing along the trajectory. This initial phase is called *progressive sharpening (PS)* in prior work Cohen et al. [6]. Empirically, the behavior of GD trajectory (as well as the loss and the sharpness) is very similar to that of gradient flow, until the sharpness reaches $2/\eta$ (this phenomena is also observed in Cohen et al. [6]. See Figure 5 or Appendix J.1 in their paper). We note that GD may come back to this phase from Phase IV later.

Phase II: Sharpness $\Lambda > 2/\eta$. In this phase, the sharpness exceeds $2/\eta$ and may keep increasing. We will show shortly that the fact that $\Lambda > 2/\eta$ causes $|\mathbf{D}^\top \mathbf{v}_1|$ (where \mathbf{v}_1 the first eigenvector of M) to increase exponentially (Lemma 3.2). This would quickly lead $\|\mathbf{D}\|$ to exceed $\|\mathbf{Y}\|$ in a few iterations, which leads the sharpness to start decreasing by Proposition 3.1, hence the training process enters Phase III.

Phase III: Sharpness $\Lambda > 2/\eta$ yet begins to gradually drop. Before Λ drops below $2/\eta$, Lemma 3.2 still holds, so $|\mathbf{D}^\top \mathbf{v}_1|$ keeps increasing. Proposition 3.1 still holds and thus the sharpness keeps decreasing until it is below $2/\eta$, at which point we enter Phase IV. A distinctive feature of this phase is that the loss may increase due to the exponential increase of $|\mathbf{D}^\top \mathbf{v}_1|$.

Phase IV: Sharpness $\Lambda < 2/\eta$. When the sharpness is below $2/\eta$, $|\mathbf{D}^\top \mathbf{v}_1|$ begins to decrease quickly, leading the loss to decrease quickly. At the same time, the sharpness keeps oscillating and gradually decreasing for some iterations. This lasts until the loss decrease to a level that is around its value right before Phase III. The sharpness is still below $2/\eta$ and our training process gets back to Phase I.

3.2 The Norm of the Output Layer Weight

It is difficult to rigorously analyze the dynamics of the sharpness $\Lambda(t)$. In this subsection, we make an interesting observation, that the change of the norm of the output layer of the network (usually a fully-connected linear layer) is consistent with the change of the sharpness most of the time.

In particular, for a general neural network $f(\mathbf{x}) = \mathbf{A}^\top \mathbf{h}(\mathbf{W}, \mathbf{x})$, where $\mathbf{A} \in \mathbb{R}^m$ is the output layer weight and the feature extractor $\mathbf{h} : \mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathbb{R}^m$ outputs a m -dimensional feature vector (\mathbf{h} corresponds to all but the last layers). $\mathbf{W} \in \mathbb{R}^p$ is the parameter vector of the extractor \mathbf{h} .

Note that $M = \left(\frac{\partial \mathbf{F}}{\partial \boldsymbol{\theta}}\right)^\top \left(\frac{\partial \mathbf{F}}{\partial \boldsymbol{\theta}}\right)$ can be decomposed as follows:

$$M = \left(\frac{\partial \mathbf{F}}{\partial \boldsymbol{\theta}}\right) \left(\frac{\partial \mathbf{F}}{\partial \boldsymbol{\theta}}\right)^\top = \left(\frac{\partial \mathbf{F}}{\partial \mathbf{A}}\right) \left(\frac{\partial \mathbf{F}}{\partial \mathbf{A}}\right)^\top + \left(\frac{\partial \mathbf{F}}{\partial \mathbf{W}}\right) \left(\frac{\partial \mathbf{F}}{\partial \mathbf{W}}\right)^\top := M_{\mathbf{A}} + M_{\mathbf{W}}.$$

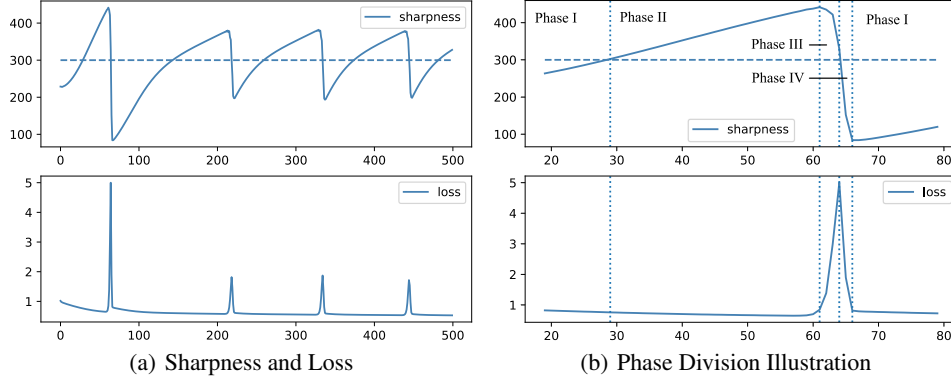


Figure 1: In Figure (a), we show the sharpness and the loss when training on a two-hidden-layer linear activated network by gradient descent. Shortly after the sharpness crosses $2/\eta$, it drops quickly back to some value below $2/\eta$. In Figure (b) we illustrate how the four phases are divided according to the evolution of the sharpness.

where the (i, j) -entry of $M_{\mathbf{W}}$ is $(M_{\mathbf{W}})_{ij} = \left\langle \frac{\partial f(\mathbf{x}_i)}{\partial \mathbf{W}}, \frac{\partial f(\mathbf{x}_j)}{\partial \mathbf{W}} \right\rangle = \mathbf{A}^\top \frac{\partial \mathbf{h}(\mathbf{W}, \mathbf{x}_i)}{\partial \mathbf{W}} \frac{\partial \mathbf{h}(\mathbf{W}, \mathbf{x}_j)}{\partial \mathbf{W}}^\top \mathbf{A}$.

In this expression, intuitively $\|\mathbf{A}\|$ should be positively related to $\|M_{\mathbf{W}}\|$. We empirically observe that the part $M_{\mathbf{A}} = \left(\frac{\partial \mathbf{F}}{\partial \mathbf{A}}\right) \left(\frac{\partial \mathbf{F}}{\partial \mathbf{A}}\right)^\top$ has a much smaller spectral norm compared to the whole Gram matrix M (see Figure 3(a) and Appendix D), which means $\|M_{\mathbf{W}}\|$ dominates $\|M_{\mathbf{A}}\|$. Therefore, $\|\mathbf{A}\|$ should be positively correlated with $\|M\|$.

The benefit of analyzing $\|\mathbf{A}\|^2$ is that the gradient flow of $\|\mathbf{A}\|^2$ enjoys the following clean formula:

$$\frac{d\|\mathbf{A}\|^2}{dt} = -2 \left(\frac{\partial \mathcal{L}}{\partial \mathbf{A}} \right)^\top \mathbf{A} = -\frac{4}{n} \mathbf{D}^\top \left(\frac{\partial \mathbf{F}}{\partial \mathbf{A}} \right) \mathbf{A} = -\frac{4}{n} \mathbf{D}^\top \mathbf{F}. \quad (2)$$

In this work, we do experiments on two-layer linear networks, fully connected deep neural networks, and Resnet18, and all of them have such output layer structures. From Figure 3(a), we can observe that **the output layer norm $\|\mathbf{A}\|^2$ and the sharpness Λ change in the same direction** most of the time along the gradient descent trajectory, i.e., they both increase or decrease at the same time. We note that they may change in different directions very occasionally around the time when $\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2$ changes its sign (see the experiments in Figure 2).

3.3 Detailed Analysis of Each Phase

In this section, we explain the dynamics of each phase in more detail. For clarity, we first list the assumptions we need in this section. For different phases, we may need some different assumptions to simplify the arguments. Most of the assumptions are consistent with the experiments or the findings in the literature. Some of them are somewhat stronger, and we also discuss how to relax them.

3.3.1 Assumptions Used in Section 3.3

Assumption 3.1. (*A-norm and sharpness*) Along the gradient descent training trajectory, for all time t , the norm $\|\mathbf{A}(t)\|$ of the output layer and the sharpness $\Lambda(t)$ moves in the same directions, i.e., $\text{sign}(\Lambda(t+1) - \Lambda(t)) = \text{sign}(\|\mathbf{A}(t+1)\| - \|\mathbf{A}(t)\|)$.

It is the key observation that we have discussed in Section 3.2. The following are two assumptions about the gradient descent trajectory. The first one assumes that $\mathbf{D}(t)$ and $\|\mathbf{A}\|^2$ are updated according to their first order approximations. Empirical justification of this approximation can be found in Appendix D.1.3.

Assumption 3.2. (*First Order Approximation of GD*) Along the gradient descent trajectory, the update rule is assumed as the first order approximation

$$\mathbf{D}(t+1) - \mathbf{D}(t) = -\eta \mathbf{M}(t) \mathbf{D}(t), \quad \|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2 = -\frac{4\eta}{n} \mathbf{D}(t)^\top \mathbf{F}(t) \quad (3)$$

Assumption 3.3. (*Gradient flow for the PS phase*) When $\Lambda(t) < 2/\eta$, $\mathbf{D}(t)$ follows the gradient flow trajectory: $\frac{d\mathbf{D}(t)}{dt} = -\mathbf{M}(t) \mathbf{D}(t)$.

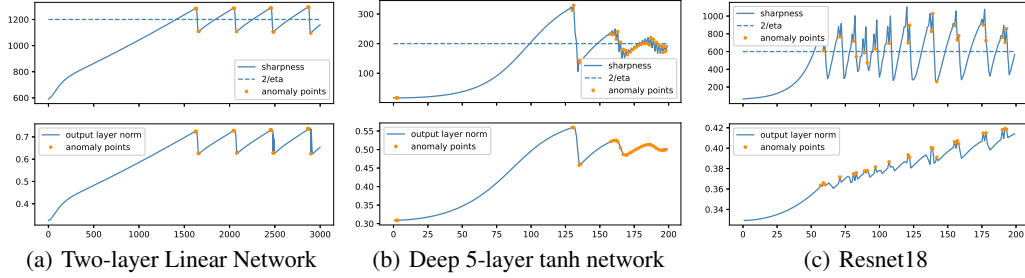


Figure 2: Sharpness and A norm are highly correlated; *Anomaly points* (orange points) in the figure correspond to the iterations when $\text{sign}(\Lambda(t+1) - \Lambda(t)) \neq \text{sign}(\|\mathbf{A}(t+1)\| - \|\mathbf{A}(t)\|)$. More discussions can be found in Section 5 and Appendix B.

Assumption 3.3 holds empirically, especially in the progressive sharpening phase (see Figure 5 or Appendix J.1 in Cohen et al. [6]) when the networks are continuously differentiable. We include these experimental details in Appendix D. See also (Theorems 4.3 and 4.5) in Arora et al. [3] for further theoretical justification. We need this assumption for the proof in the progressive sharpening phase.

Then we state an assumption on the upper bound of the sharpness to restrict the regime we discuss:

Assumption 3.4. (*Sharpness upper bound*) *If the training does not diverge, there exists some constant B_Λ , such that $0 < \Lambda(t) \leq \frac{B_\Lambda}{\eta}$ for all t .*

This assumption states that there is an upper bound of the sharpness throughout the optimization process. Actually, in Lewkowycz et al. [19], they proved that $4/\eta$ is an upper bound of the sharpness in a two-layer linear network with one datapoint, otherwise the training process (loss) would diverge. They empirically found that similar upper bounds exist also for nonlinear activations, albeit with somewhat larger constant B_Λ . In the work, We focus on the case when the loss does not diverge and hence we make Assumption 3.4.

The main set of assumptions we need is about the change of M 's eigendirections.

Assumption 3.5. *Denote $\{\mathbf{v}_i\}_{i=1}^n$ to be the set of eigenvectors of $M(t)$. We have three levels of assumptions on M 's eigenspace.*

- (i) (*fixed eigendirections*) *the set $\{\mathbf{v}_i\}_{i=1}^n$ is fixed throughout the phase under consideration;*
- (ii) (*eigendirections move slowly*) *at all time t and for any i , $\mathbf{F}(t)^\top \frac{d\mathbf{v}_i(t)}{dt} < \lambda_i(t) \mathbf{D}(t)^\top \mathbf{v}_i(t)$;*
- (iii) (*principal directions moves slowly*) *at all time t , there is a small constant $\epsilon_2 \geq 0$ such that $\langle \mathbf{v}_1(t), \mathbf{v}_1(t+1) \rangle \geq 1 - \epsilon_2$.*

Clearly, these three assumptions are increasingly weaker from (i) to (iii). Assumption 3.5 (i) on the eigenvectors is somewhat strong, and the eigenvectors corresponding to small eigenvalues may change notably in our experiments. We use it to illustrate a basic proof idea of the progressive sharpening phase, but later we relax this assumption to Assumption 3.5 (ii).

Moreover, for the proof in Phase II and III, Assumption 3.5 (iii), which only assume that the main direction changes slightly, is sufficient for our proof. Actually, we note that $\mathbf{v}_1(t)$ (the eigenvector corresponding to the largest eigenvalue) changes slowly and the inner product of its initial direction and its direction at the end of the phase is also large (see Appendix D for the empirical verification).

For the proof in Phase II, we need another small technical assumption:

Assumption 3.6. *Assume $\mathbf{D}(t)^\top \mathbf{v}_1(t) \geq c\epsilon_2 \|\mathbf{D}(t)\|$ for some $c > 1$ for some $t = t_0$ at the beginning of this phase. Here ϵ_2 is defined in Assumption 3.5 (iii).*

Assumption 3.6 says that $\mathbf{D}(t)$ has a non-negligible component in the direction of \mathbf{v}_1 . Since $\epsilon_2 > 0$ is a small constant, this is not a strong assumption as some small perturbation (due to discrete updates) would make the assumption hold for some $c > 1$.

3.3.2 Detailed Analysis

In each phase, we attempt to explain the main driving force of the change of the sharpness and the loss.

Phase I: In this phase, we show that $\mathbf{D}(t)^\top \mathbf{F}(t) < 0$ under certain assumptions (detailed shortly) on the spectral properties of $\mathbf{M}(t)$ (see Lemma 3.1 below). By Assumption 3.2, we have $\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2 > 0$, implying that the sharpness $\Lambda(t)$ also increases based on Assumption 3.1. This phase stops if $\Lambda(t)$ grows larger than $2/\eta$.

We assume the output vector $\mathbf{F}(t)$ is initialized to be small (this is true if we use very small initial weights). For simplicity, we assume $\mathbf{F}(0) = 0$ in the following argument.

Lemma 3.1. *For all t in Phase I, under Assumption 3.5 (i) and 3.3, it holds that $\mathbf{D}(t)^\top \mathbf{F}(t) < 0$.*

From this lemma, $\|\mathbf{A}\|$ keeps increasing by Assumption 3.2; hence the sharpness keeps increasing by Assumption 3.1 until it reaches $2/\eta$ or the loss converges to 0. In the former case, the training process enters Phase II, while the latter case is also possible when η is very small (e.g., even the largest possible sharpness value is less than $2/\eta$). We admit that Assumption 3.5 (i) is somewhat strong. In fact, the assumption can be relaxed significantly to Assumption 3.5 (ii). We show in Appendix E.2 that under Assumption 3.5 (ii) and Assumption 3.3, we can still guarantee $\mathbf{D}(t)^\top \mathbf{F}(t) < 0$. Moreover, we provide a dynamical system view of the dynamics of $\mathbf{D}(t)^\top \mathbf{F}(t)$ in that Appendix.

Phase II: When the training process just enters Phase II, the sharpness keeps increasing. We show shortly that $\mathbf{D}(t)^\top \mathbf{v}_1(t)$ starts to increase geometrically, and this causes the sharpness to stop increasing at some point, thus entering Phase III.

In this phase, we adopt a weaker assumption on the sharpness direction \mathbf{v}_1 : Assumption 3.5 (iii). This assumption holds in our experiments (See Figure 17). Also, Assumption 3.6 is necessary.

Lemma 3.2. *Suppose Assumption 3.5 (iii) and 3.6 hold during this phase (with constants $\epsilon_2 > 0$ and $c > 1$). If $\Lambda(t) = (2 + \tau)/\eta$ and $\tau > \frac{1}{1-\epsilon_2-1/c} - 1$, then $\mathbf{D}(t)^\top \mathbf{v}_1(t)$ increases geometrically with factor $(1 + \tau)(1 - \epsilon_2 - 1/c) > 1$ for $t \geq t_0$ in this phase.*

Since $\mathbf{D}(t)^\top \mathbf{v}_1(t)$ increases geometrically, $\|\mathbf{D}\| \geq \mathbf{D}(t)^\top \mathbf{v}_1(t)$ will exceed $\|\mathbf{Y}\|$ eventually. Next, the following proposition states that when this happens, $\mathbf{D}(t)^\top \mathbf{F}(t) > 0$. Consequently, $\|\mathbf{A}\|$ decreases by Assumption 3.2, leading to the decrement of the sharpness based on our Assumption 3.1.

Proposition 3.1. *If $\|\mathbf{D}(t)\| > \|\mathbf{Y}\|$, then $\mathbf{D}(t)^\top \mathbf{F}(t) > 0$.*

Phase III: The sharpness is still larger than $2/\eta$, but it starts decreasing. Meanwhile, the loss continues to increase rapidly due to Lemma 3.2. Eventually, the sharpness will fall below $2/\eta$ and then the training process enters phase IV.

By Lemma 3.2, if the sharpness stays above $2/\eta$, then we can have an arbitrarily large loss. According to Proposition 3.1, if the loss is large enough, the sharpness keeps decreasing.

Now we show that if the sharpness stays above $2/\eta$, $\|\mathbf{A}(t)\|^2$ will decrease by a significant amount. This partially explains that the sharpness should also decrease significantly until it drops below $2/\eta$ (instead of decreasingly converging to a value above $2/\eta$ without ever entering the next phase).

Proposition 3.2. *Under Assumption 3.2, if $\|\mathbf{D}(t)\| > \|\mathbf{Y}\|$, then $\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2 < -\frac{4\eta}{n}(\|\mathbf{D}(t)\| - \|\mathbf{Y}\|)^2$.*

From the above argument, we can see that if $\mathbf{D}(t)^\top \mathbf{v}_1(t)$ is larger than $\|\mathbf{Y}\|$, then $\mathbf{D}(t)^\top \mathbf{v}_1(t)$ does not decrease in Phase III, and according to Proposition 3.2, $\|\mathbf{A}(t)\|^2$ decreases significantly, implying the sharpness drops below $2/\eta$ eventually.

Remark: The fact that the sharpness can provably drop below $2/\eta$ in this phase can be proved more rigorously in Section 4 for the two-layer linear setting. See Theorem 2.

Phase IV: First, since the training process has just left phase III, $\mathbf{D}(t)^\top \mathbf{F}(t)$ is still positive and large, hence $\|\mathbf{A}(t)\|^2$ keeps decreasing and the sharpness decreases as well. Since the sharpness stays below $2/\eta$, the loss decreases due to the following descent lemma (with \mathbf{u} replaced by $\mathbf{D}(t)$).

Lemma 3.3. *If $\Lambda(t) < 2/\eta$, then for any vector $\mathbf{u} \in \mathbb{R}^n$, $\|\mathbf{u}^\top (\mathbf{I} - \eta \mathbf{M}(t))\| \leq (1 - \eta \alpha) \|\mathbf{u}\|$, where $\alpha = \min\{2/\eta - \Lambda(t), \lambda_{\min}(\mathbf{M}(t))\}$. In particular, replacing \mathbf{u} with $\mathbf{D}(t)$, we can see $\|\mathbf{D}(t+1)\| \leq (1 - \eta \alpha)^2 \|\mathbf{D}(t)\|$.*

Next we argue that $\mathbf{D}(t)^\top \mathbf{F}(t)$ will become negative eventually, which indicates that $\|\mathbf{A}(t)\|^2$ and hence the sharpness will grow again. Since the sharpness is below $2/\eta$, $\mathbf{D}(t)^\top \mathbf{v}_1(t)$ decreases geometrically due to Lemma 3.3 (replacing \mathbf{u} with $\mathbf{D}(t)\mathbf{v}_1(t)\mathbf{v}_1(t)^\top$).

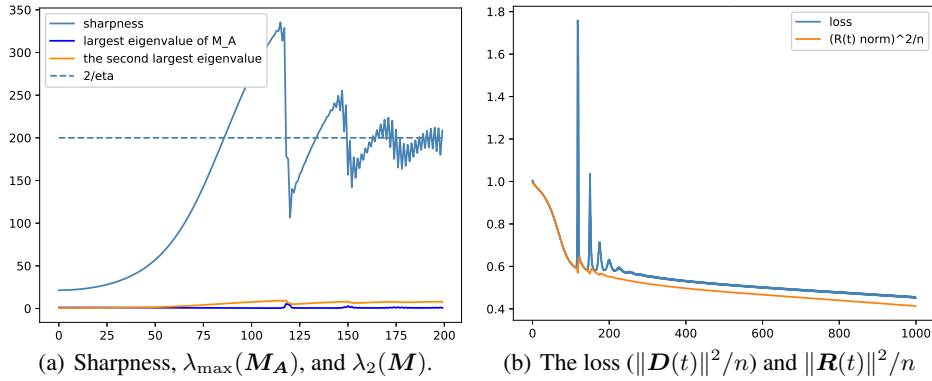


Figure 3: In Figure (a), we verify the assumptions in Section 3.2 and 3.3 empirically. The steel blue line represents the largest eigenvalue of M , i.e. the sharpness $\lambda_{\max}(M)$. The orange line: the second largest eigenvalue is far below $1/\eta$, illustrating Assumption 3.7. The blue line: the largest eigenvalue of M_A is negligible compared to the sharpness. In Figure (b), we show that even though the total loss $\mathcal{L}(t) = \|D(t)\|^2/n$ oscillates considerably, $\|R(t)\|^2/n$ decreases more steadily.

In fact, D can be decomposed into the v_1 -component $v_1 v_1^\top D$ and the remaining part R defined as $R(t) := (I - v_1(t)v_1(t)^\top)D(t)$. Then we have

$$D(t)^\top F(t) = (v_1(t)v_1(t)^\top D(t))^\top (v_1(t)v_1(t)^\top D(t) + Y) + R(t)^\top (R(t) + Y). \quad (4)$$

As shown in the next subsection, $R(t)$ almost follows a similar gradient descent trajectory $R'(t)$ (Lemma 3.5). More precisely, $R'(t)$ is defined as $R'(t+1) = (I - \eta M(t)(I - v_1(t)v_1(t)^\top))R'(t)$ (Lemma 3.4). While D 's dynamics is $D(t+1) = (I - \eta M(t))D(t)$, R' follows a similar dynamics $R'(t+1) = (I - \eta M'(t))R'(t)$, where $M'(t) = M(t)(I - v_1(t)v_1(t)^\top)$. Note that $M'(t)$ has eigenvalues smaller than $1/\eta$ for any time t (by Assumption 3.7), hence with an assumption similar to Assumption 3.5 (i) (or a similar version of our relaxed assumption in Appendix E.2 for M'), we can prove that $R'(t)^\top (R'(t) + Y) < 0$ for any time t (See Appendix E.1 for the rigorous proof).

Since $R(t) \approx R'(t)$, the second term in the decomposition (4) is always negative and the first term (v_1 direction term) is decreasing geometrically. Therefore, there are only two possible cases. The first possibility is that the first term decreases to a small value near 0 and the second term remains largely negative. Then their sum will be negative, which is $D(t)^\top F(t) < 0$, thus implying the training enters Phase I. The second possibility is that when the first term decreases to a small value near 0, the second term is also a small negative value. In this case both $R(t)$ and $D(t)^\top v_1(t)$ are small, implying the loss is almost 0, which is indeed the end of the training.

3.4 Explaining Non-monotonic Loss Decrement

In this subsection, we attempt to explain the non-monotonic decrement of the loss during the entire GD trajectory. See Figure 3(b). As defined in the last section, we decompose D into the v_1 -component $v_1 v_1^\top D$ and the remaining part R . Below, we prove that $R(t)$ is not affected much by the exponential growth of the loss (Proposition 3.5) in Phase II and III, and almost follows a converging trajectory (which is defined as $R'(t)$ later in this section).

The arguments in this subsection need Assumption 3.4 and Assumption 3.5 (iii), both very consistent with the experiments. We need an additional assumption on the spectrum of M .

Assumption 3.7. All $M(t)$'s eigenvalues except $\Lambda(t) = \lambda_{\max}(M(t))$ are smaller than $1/\eta$ for all t .

Recall that the largest eigenvalue is at most $\frac{B_\Lambda}{\eta}$ by Assumption 3.4. Empirically, the largest eigenvalue is an outlier in the spectrum, i.e., it is much larger than the other eigenvalues. Hence, we make Assumption 3.7 which states that all other eigenvalues are at most $1/\eta$, which is consistent with our experiments. See Figure 3(b). Similar fact is also mentioned in [28, 29].

First, we let B_D be an upper bound of $D(t)$, i.e., for all t , $\|D(t)\| \leq B_D$. In the two-layer linear network case, we can have an explicit form of B_D . (see Lemma C.9 in Appendix C.) Recall that in Assumption 3.4, B_Λ is the upper bound of $\eta\Lambda$.

Lemma 3.4. Suppose Assumption 3.5 (iii) holds. $R(t)$ satisfies the following:

$$R(t+1) = (I - \eta M(t))R(t) + e_1(t), \quad \text{where} \quad \|e_1(t)\| \leq 6\sqrt{\epsilon_2}\|D(t)\|(B_\Lambda - 1)$$

Lemma 3.5. Define an auxiliary sequence $\mathbf{R}'(t)$ by $\mathbf{R}'(0) = \mathbf{R}(0)$, and $\mathbf{R}'(t+1) = (\mathbf{I} - \eta \mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{v}_1(t) \mathbf{v}_1(t)^\top) \mathbf{R}'(t)$. If Assumption 3.4, Assumption 3.5 (iii), Assumption 3.7 hold, and for any time t there exists a quantity $\lambda_r > 0$, such that the smallest eigenvalue of $\mathbf{M}(t)$, i.e. $\lambda_{\min}(\mathbf{M}(t)) > \lambda_r$, then there exists a constant $c_r > 0$ such that $\|\mathbf{R}(t) - \mathbf{R}'(t)\| \leq c_r \frac{B_D(B_\Lambda - 1)\sqrt{\epsilon_2}}{\eta \lambda_r}$.

Now, in light of Lemma 3.2 and Lemma 3.5, we arrive at an interesting explanation of the phenomena of non-monotonic decrease of the loss. Basically, \mathbf{D} can be decomposed into the \mathbf{v}_1 -component $\mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}$ and the remaining part $\mathbf{R} = (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}$. The \mathbf{v}_1 -component may increase geometrically during the EOS (Lemma 3.2), but the behavior of the remaining part $\mathbf{R}(t)$ is close to $\mathbf{R}'(t)$, which follows the simple updating rule $\mathbf{R}'(t+1) = (\mathbf{I} - \eta \mathbf{M}(t)) \mathbf{R}'(t)$, so Lemma 3.3 implies that the \mathbf{R} part almost keeps decreasing during the entire trajectory (here Lemma 3.3 applies with \mathbf{u} replaced by $\mathbf{R}'(t)$, noticing that the eigenvalues except the first are well below $2/\eta$). Hence, the non-monotonicity of the loss is mainly due to the \mathbf{v}_1 -component of \mathbf{D} , and the rest part \mathbf{R} is optimized in the classical regime (step size well below $2/(\text{the operator norm})$) and hence steadily decreases. See Figure 3(b).

4 A Theoretical Analysis for 2-Layer Linear NN

In this section, we aim to provide a more rigorous explanation of the EOS phenomenon in two-layer linear networks. The proof ideas follow similar high-level intuition as the proofs in Section 3.3. In particular, we can remove or replace the assumptions in Section 3.3 with arguably weaker assumptions. Due to space limit, we state our main theoretical results and elaborate their relation with the proofs in Section 3.3. The detailed settings and proof are more tedious and can be found in Appendix C.

4.1 Setting and basic notations

Model: In this section, we study a two-layer neural network with linear activation, i.e. $f(x) = \sum_{q=1}^m \frac{1}{\sqrt{m}} a_q \mathbf{w}_q x = \frac{1}{\sqrt{m}} \mathbf{A}^\top \mathbf{W} x$ where $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_m]^\top \in \mathbb{R}^{m \times d}$, $\mathbf{A} = [a_1, \dots, a_m] \in \mathbb{R}^m$.

Dataset: For simplicity, we assume $y_i = \pm 1$ for all $i \in [n]$, and $\|\mathbf{X}^\top \mathbf{X}\|_2 = \Theta(n)$. We assume $\mathbf{X}^\top \mathbf{X}$ has rank r , and we decompose $\mathbf{X}^\top \mathbf{X}$ and \mathbf{Y} according to the orthonormal basis $\{\mathbf{v}_i\}$, the eigenvectors of $\mathbf{X}^\top \mathbf{X}$: $\mathbf{X}^\top \mathbf{X} = \sum_{i=1}^r \lambda_i \mathbf{v}_i \mathbf{v}_i^\top$, $\mathbf{Y} = \sum_{i=1}^r (\mathbf{Y}^\top \mathbf{v}_i) \mathbf{v}_i := \sum_{i=1}^r z_i \mathbf{v}_i$ where \mathbf{v}_i is the eigenvector corresponding to the i -th largest eigenvalue λ_i of $\mathbf{X}^\top \mathbf{X}$. $z_i = \mathbf{Y}^\top \mathbf{v}_i$ is the projection of \mathbf{Y} onto the direction \mathbf{v}_i . Here we suppose $n \gg r$ and the global minimum $(\mathbf{A}^*, \mathbf{W}^*)$ exists.

Update rule: We write explicitly the GD dynamics of $\mathbf{D}(t)$: $\mathbf{D}(t+1) = (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{D}(t)$, where $\mathbf{M}^*(t) = \frac{2}{mn} (\|\mathbf{A}(t)\|^2 \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{W}^\top(t) \mathbf{W}(t) \mathbf{X}) - \frac{4\eta}{n^2 m} (\mathbf{D}(t)^\top \mathbf{F}(t)) \mathbf{X}^\top \mathbf{X}$ is the Gram matrix combined with second order terms.

4.2 Main Theorem and The Proof Sketch

Phase I and Progressive Sharpening:

Assumption 4.1. There exists some constant $\chi > 1$, s.t. for all $i \in [r-1]$, $\lambda_i(\mathbf{X}^\top \mathbf{X}) \leq \chi \lambda_{i+1}(\mathbf{X}^\top \mathbf{X})$. Moreover, $\lambda_1(\mathbf{X}^\top \mathbf{X}) \geq 2\lambda_2(\mathbf{X}^\top \mathbf{X})$.

Assumption 4.2. There exists $\kappa = \Omega(r^{-1})$ such that $\min_{i \in [r]} \{z_i / \sqrt{n}\} \geq \kappa$.

The first assumption is about the eigenvalue spectrum of $\mathbf{X}^\top \mathbf{X}$.³ The second assumes that all component $z_i = \mathbf{Y}^\top \mathbf{v}_i$ are not too small.

Theorem 1 (Informal). Suppose Assumption 4.1, Assumption 4.2 hold, the smallest nonzero eigenvalue $\lambda_r = \lambda_r(\mathbf{X}^\top \mathbf{X}) > 0$ and $\lambda_1 = \lambda_{\max}(\mathbf{X}^\top \mathbf{X}) = c_1 n$. Then for any $\epsilon > 0$, if $m = \Omega(\frac{c_1 n^2}{\lambda_r^2})$, and $n = \Omega(\frac{\lambda_r^2}{\kappa^4 \epsilon^2})$, we have the progressive sharpening property: $\Lambda(t+1) - \Lambda(t) > 0$ for $t = 1, 2, \dots, t_0 - 1$ where t_0 is the time when $\|\mathbf{D}(t)\|^2 \leq O(\epsilon^2)$ or $\lambda_{\max}(\mathbf{M}^*(t)) > 1/\eta$ for the first time.

In the proof of this theorem, we show that the Gram matrix $\mathbf{M}(t) \approx \frac{2}{mn} (\|\mathbf{A}(t)\|^2 + \frac{m}{d}) \mathbf{X}^\top \mathbf{X}$, which serves as a justification of Assumption 3.5 we made in Section 3.3. That shows all $\mathbf{M}(t)$

³It guarantees the gap between two adjacent eigenvalues is not very large, and there is a gap between the largest and the second largest eigenvalue. Note the second part of the assumption is a relaxed version of Assumption 3.7. In our CIFAR-10 1k-subset with samples' mean subtracted, $\lambda_1/\lambda_2 = \chi \approx 3$ (See Figure 19).

approximately share the same set of eigenvectors as $\mathbf{X}^\top \mathbf{X}$. In our proof, we also prove more rigorously that $\|\mathbf{A}(t)\|^2$ is an indicator of the sharpness in this simpler setting.

Edge of Stability (Phase II - IV):

Assumption 4.3. *There exists some constant $c_2 > 0$, such that $\|\mathbf{\Gamma}(t)\| \leq \frac{c_2}{m}$.*

This assumption is based on Theorem 1. In Theorem 1, we state that in the progressive sharpening phase, $\|\mathbf{\Gamma}(t)\|$ has an upper bound of $O(1/m)$. Now in the EOS phase, we assume that $\|\mathbf{\Gamma}(t)\|$ grows larger by at most a constant factor. Further discussions refer to Appendix D.2.2.

Assumption 4.4. *There exists some constant $\beta > 0$, such that $\Lambda \leq \frac{4}{\eta}(1 - \beta)$.*

This assumption is consistent with Assumption 3.4, which assumes an upper bound of the sharpness.

Assumption 4.5. *There exist some constant c_3 such that $|\mathbf{D}(t)^\top \mathbf{v}_1| > c_3 \sqrt{n}/m$ for some $t = t_0$ at the beginning of phase II.*

This assumption is in the same spirit of Assumption 3.6 with the only change of the bound in terms of m and n . Now, we are ready to state our theorem in this stage.

Theorem 2. *Denote the smallest nonzero eigenvalue as $\lambda_r \triangleq \lambda_r(\mathbf{X}^\top \mathbf{X}) > 0$ and the largest eigenvalue as $\lambda_1 \triangleq \lambda_1(\mathbf{X}^\top \mathbf{X})$. Under Assumption 4.3, 4.4, 4.5, and $\lambda_1(\mathbf{X}^\top \mathbf{X}) \geq 2\lambda_2(\mathbf{X}^\top \mathbf{X})$ in Assumption 4.1, there exist constants c_4, c_5, c_6 such that if $n > c_6 \lambda_r \eta$, $m > \max\{\frac{c_4 d^2 n^2}{\lambda_r^2}, c_5 \eta\}$, then*

- *There exists $\rho = O(1)$ which depends on c_3 such that if $\Lambda(t_0) > \frac{2}{\eta}(1 + \rho)$ for some t_0 , there must exist some $t_1 > t_0$ such that $\Lambda(t_1) < \frac{2}{\eta}(1 + \rho)$.*
- *If $\Lambda(t), \Lambda(t + 1) > \frac{2}{\eta}(1 + \rho)$, then there is a constant $c_7 > 0$ (depending on c_3) such that $|\mathbf{D}(t + 1)^\top \mathbf{v}_1| > |\mathbf{D}(t)^\top \mathbf{v}_1|(1 + c_7)$.*
- *Define $\mathbf{R}(t) := (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t)$, and $\mathbf{R}'(t) := (\mathbf{I} - \eta \mathbf{M}^*(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top)) \mathbf{R}'(t - 1)$. It holds that $\|\mathbf{R}(t) - \mathbf{R}'(t)\| = O(\frac{\sqrt{n^3 d}}{\lambda_r \sqrt{m}})$.*

We can conclude the following from Theorem 2: (1) The first statement of the theorem states that if the progressive sharpening phase causes the sharpness to grow over $2/\eta$, then the sharpness eventually goes below $2/\eta$. This illustrates the regularization effect of gradient descent on the sharpness (this is consistent with the analysis of Phase III in Section 3.3). (2) The second states that $|\mathbf{D}(t)^\top \mathbf{v}_1|$ geometrically increases in Phase II and III. Note that we proved a similar Lemma 3.2 for Phase II in the more general setting in Section 3.3. (3) The third conclusion gives an upper bound for the distance between $\mathbf{R}(t)$'s trajectory and $\mathbf{R}'(t)$'s. This bound helps illustrate why $\mathbf{R}(t)$'s trajectory is similar with $\mathbf{R}'(t)$ in Phase IV of Section 3.3.

5 Discussions and Open Problems

In this section, we discuss the limitation of our theory and some related findings. First, our argument crucially relies on the assumption that $\|\mathbf{A}\|$ changes in the same direction as Λ does most of the time. Here, we elaborate more on this point. Seeing from a longer time scale, $\|\mathbf{A}\|^2$ and the sharpness may have very different overall trends (See Figure (c) in 2), i.e., the sharpness oscillates around $2/\eta$ but $\|\mathbf{A}\|^2$ increases. Moreover, the sharpness may oscillate more frequently than $\|\mathbf{A}\|^2$, while the low-frequency trends seem to match well (See the late training phases in Figure (b) in 2). Currently, our theory cannot explain the high-frequency oscillation of the sharpness in Figure (b).

While we still believe the change of $\|\mathbf{A}\|$ is a major driving force of the change of the sharpness, other factors (such as other layers) must be taken into consideration for a complete understanding and explanation of the sharpness dynamics. We also carry out some experiments that reveal some interesting relation between the inner layers and the sharpness, which is not yet reflected in our theory. Due to space limit, we defer it to Appendix D.3.

We conclude with some open problems. It would be very interesting to remove some of our assumptions or replace them (especially those related to the spectrum of \mathbf{M}) by weaker or more natural assumptions on the data or architectures, or make some of the heuristic argument more rigorous (e.g., first order approximation of the dynamics (3)). Extending our results in Section 4 to deeper neural networks with nonlinear activation function is an intriguing and challenging open problem.

References

- [1] Kwangjun Ahn, Jingzhao Zhang, and Suvrit Sra. Understanding the unstable convergence of gradient descent. *arXiv preprint arXiv:2204.01050*, 2022.
- [2] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR, 2019.
- [3] Sanjeev Arora, Zhiyuan Li, and Abhishek Panigrahi. Understanding gradient descent on edge of stability in deep learning. *arXiv preprint arXiv:2205.09745*, 2022.
- [4] Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- [5] Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in Neural Information Processing Systems*, 32, 2019.
- [6] Jeremy M Cohen, Simran Kaur, Yuanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. *arXiv preprint arXiv:2103.00065*, 2021.
- [7] Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International conference on machine learning*, pages 1675–1685. PMLR, 2019.
- [8] Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- [9] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*, 2020.
- [10] Noah Golmant, Zhewei Yao, Amir Gholami, Michael Mahoney, and Joseph Gonzalez. pytorchhessian-eigentings: efficient pytorch hessian eigendecomposition, october 2018. URL <https://github.com/noahgolmant/pytorch-hessian-eigentings>.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [12] Wei Hu, Lechao Xiao, Ben Adlam, and Jeffrey Pennington. The surprising simplicity of the early-time learning dynamics of neural networks. *Advances in Neural Information Processing Systems*, 33:17116–17128, 2020.
- [13] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [14] Stanislaw Jastrzebski, Maciej Szymczak, Stanislav Fort, Devansh Arpit, Jacek Tabor, Kyunghyun Cho, and Krzysztof Geras. The break-even point on optimization trajectories of deep neural networks. *arXiv preprint arXiv:2002.09572*, 2020.
- [15] Ryo Karakida, Shotaro Akaho, and Shun-ichi Amari. The normalization method for alleviating pathological sharpness in wide neural networks. *Advances in neural information processing systems*, 32, 2019.
- [16] Ryo Karakida, Shotaro Akaho, and Shun-ichi Amari. Pathological spectra of the fisher information metric and its variants in deep neural networks. *arXiv preprint arXiv:1910.05992*, 2019.
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [18] Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32, 2019.
- [19] Aitor Lewkowycz, Yasaman Bahri, Ethan Dyer, Jascha Sohl-Dickstein, and Guy Gur-Ari. The large learning rate phase of deep learning: the catapult mechanism. *arXiv preprint arXiv:2003.02218*, 2020.

- [20] Zhiyuan Li, Tianhao Wang, and Sanjeev Arora. What happens after sgd reaches zero loss?—a mathematical framework. *arXiv preprint arXiv:2110.06914*, 2021.
- [21] Kaifeng Lyu, Zhiyuan Li, and Sanjeev Arora. Understanding the generalization benefit of normalization layers: Sharpness reduction. *arXiv preprint arXiv:2206.07085*, 2022.
- [22] Chao Ma, Lei Wu, and Lexing Ying. The multiscale structure of neural network loss functions: The effect on optimization and origin. *arXiv preprint arXiv:2204.11326*, 2022.
- [23] James Martens. *Second-order optimization for neural networks*. University of Toronto (Canada), 2016.
- [24] Rotem Mulayoff, Tomer Michaeli, and Daniel Soudry. The implicit bias of minima stability: A view from function space. *Advances in Neural Information Processing Systems*, 34, 2021.
- [25] Vardan Papyan. The full spectrum of deepnet hessians at scale: Dynamics with sgd training and sample size. *arXiv preprint arXiv:1811.07062*, 2018.
- [26] Vardan Papyan. Measurements of three-level hierarchical structure in the outliers in the spectrum of deepnet hessians. *arXiv preprint arXiv:1901.08244*, 2019.
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [28] Levent Sagun, Leon Bottou, and Yann LeCun. Eigenvalues of the hessian in deep learning: Singularity and beyond. *arXiv preprint arXiv:1611.07476*, 2016.
- [29] Levent Sagun, Utku Evci, V Ugur Guney, Yann Dauphin, and Leon Bottou. Empirical analysis of the hessian of over-parametrized neural networks. *arXiv preprint arXiv:1706.04454*, 2017.
- [30] Lei Wu, Chao Ma, et al. How sgd selects the global minima in over-parameterized learning: A dynamical stability perspective. *Advances in Neural Information Processing Systems*, 31, 2018.
- [31] Yikai Wu, Xingyu Zhu, Chenwei Wu, Annie Wang, and Rong Ge. Dissecting hessian: Understanding common structure of hessian in neural networks. *arXiv preprint arXiv:2010.04261*, 2020.
- [32] Chen Xing, Devansh Arpit, Christos Tsirigotis, and Yoshua Bengio. A walk with sgd. *arXiv preprint arXiv:1802.08770*, 2018.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]**
 - (b) Did you describe the limitations of your work? **[Yes]**
 - (c) Did you discuss any potential negative societal impacts of your work? **[N/A]**
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** See Section 3.3 and 4, we included all assumptions.
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** See Appendix C.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[No]** For anonymity we will include the URL after the review.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** See Appendix A.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[N/A]**

- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix A.
- 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 3.3 and Appendix A.
 - (b) Did you mention the license of the assets? [Yes] See Appendix A.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] The data is downloaded through torchvision.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
- 5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Experimental Setup

We provide detailed experimental setup in this section.

A.1 Dataset

Under GPU memory constraints when computing Gram matrix of our models, we limit the size of dataset we use. The dataset is a 1,000-sample subset from CIFAR-10 (Krizhevsky et al. [17]) (<https://www.cs.toronto.edu/~kriz/cifar.html>). To make it a binary dataset, we constructed it by selecting the first 500 samples of class 0 and 1, respectively.⁴ Then we label the samples by ± 1 . In the experiments in Appendix B and D, we fix the objective as training on the 1000-example subset of CIFAR-10.

A.2 Network Architecture

In general settings, we experiment with three architectures from simple models to more complicated models: one-hidden-layer linear neural network, four-hidden-layer fully-connected network, convolutional network and a ResNet18 (He et al. [11]) model. The initialization of each layer follows the default initialization of PyTorch (Paszke et al. [27]).

Linear Network We first use a simple two-layer linear neural network. The hidden layer has 200 neurons. We empirically show that even a simple linear network can enter the EOS regime.

Table 1: Linear network

Layer #	Name	Layer	In shape	Out shape
1		Flatten()	(3, 32, 32)	3072
2	fc1	nn.Linear(3072, 200)	3072	200
3	fc2	nn.Linear(200, 1)	200	1

Fully-connected Network We conduct further experiments on several different fully-connected networks with 4 hidden layers with various activation functions. We consider tanh, ReLU, ELU activations. For example, the structure of a fully-connected tanh network is shown in Table 2.

Table 2: Fully-connected network

Layer #	Name	Layer	In shape	Out shape
1		Flatten()	(3, 32, 32)	3072
2	fc1	nn.Linear(3072, 200, bias=False)	3072	200
3		nn.tanh()	200	200
4	fc2	nn.Linear(200, 200, bias=False)	200	200
5		nn.tanh()	200	200
6	fc3	nn.Linear(200, 200, bias=False)	200	200
7		nn.tanh()	200	200
8	fc4	nn.Linear(200, 200, bias=False)	200	200
9		nn.tanh()	200	200
10	fc5	nn.Linear(200, 1, bias=False)	200	1

Convolutional Network We also conduct experiments on several different convolutional networks with two convolutional layers and two max-pooling layers. Like the fully-connected network experiments, we consider tanh, ReLU and ELU activations. For example, the structure of a convolutional tanh network is shown in Table 3.

⁴Cohen et al. [6] selects the first 5,000 examples from CIFAR-10. Our subset is smaller because of the computation limitation when calculating the Gram matrix. Experiments show that the properties along GD trajectory (e.g. the loss, the sharpness) is similar on both datasets.

Table 3: Convolutional network

#	Name	Layer	In shape	Out shape
1	conv1	nn.Conv2d(3,32, kernel_size=3, padding=1)	(3, 32, 32)	(32,32,32)
2		nn.tanh()	(32,32,32)	(32,32,32)
3		nn.MaxPool2d(2)	(32,32,32)	(32,16,16)
4	conv2	nn.Conv2d(32,32, kernel_size=3, padding=1)	(32,16,16)	(32,16,16)
5		nn.tanh()	(32,16,16)	(32,16,16)
6		nn.MaxPool2d(2)	(32,16,16)	(32,8,8)
7		Flatten()	(32,8,8)	2048
8	fc1	nn.Linear(2048,1, bias=False)	2048	1

ResNet18 We also conduct experiment on the ResNet18 architecture proposed by He et al. [11]. We use the default architecture implemented in PyTorch (Paszke et al. [27]). When calculating the sharpness, we use the numerical methods in the package (Golmant et al. [10]) to calculate the top eigenvalue of the Hessian matrix.

B Further Experiments and Discussions

In this appendix, we use the 1000-sized subset of CIFAR-10 introduced in Appendix A to conduct further experiments on various architectures. We verify our main observation about the correlation between the sharpness and the weight norm of the output layer (A-norm) of the neural network through the following experiments.

We consider simple linear networks, fully-connected networks, convolutional networks in this appendix. For nonlinear networks, we choose tanh, ReLU, ELU as activation functions. We train the networks with MSE loss. Here we exclude ResNet18 experiment since it is not feasible to compute the Gram matrix or the leading Gram matrix eigenvector due to GPU memory limitation. The sharpness and A-norm correlation of ResNet18 are included in Figure 2 in Section 3.2.

Here we run full-batch gradient descent with a selected step size η such that the sharpness at initialization is smaller than $2/\eta$.

B.1 Further Experiments

B.1.1 Linear Networks

We first verify our four-phase division of EOS phenomena in a simple linear network. This experiment shows that even linear networks can also enter EOS regime and the four-phase division of the gradient descent trajectory is quite apparent in this setting. The following Figure 4 illustrates the positive correlation between the sharpness and the A-norm, and the relationship between the loss $\|\mathbf{D}(t)\|^2$ and $\|\mathbf{R}(t)\|^2$ along the trajectory.

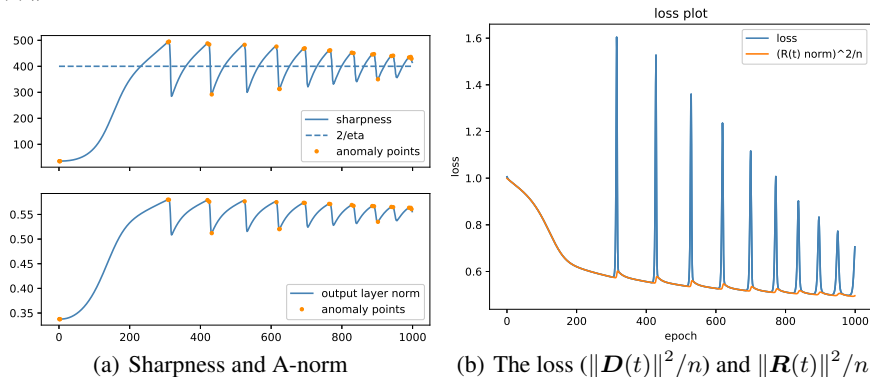


Figure 4: Fully-connected linear network. In Figure (a), we verify the correlation between sharpness and A-norm in convolutional network setting. In Figure (b), we show that in this case, even though the total loss $\mathcal{L}(t) = \|\mathbf{D}(t)\|^2/n$ oscillates considerably, $\|\mathbf{R}(t)\|^2/n$ decreases more steadily.

B.1.2 Fully-connected networks

We train three 5-layer fully-connected networks with different activation functions: tanh, ReLU, and ELU activations. In Figure 5, 6, and 7, we verify that in these fully-connected networks, the sharpness is positively correlated to the dynamics output layer norm (A-norm) most of the time in the progressive sharpening stage and the first few oscillations.

Note that anomaly points appear much more frequently after a few oscillations. Meanwhile, the sharpness oscillates more frequently around $2/\eta$ in every few iterations. We further discuss the phenomenon in Appendix B.2, in which we elaborate the complicated relationship of the sharpness, A-norm and parameters of other layers.

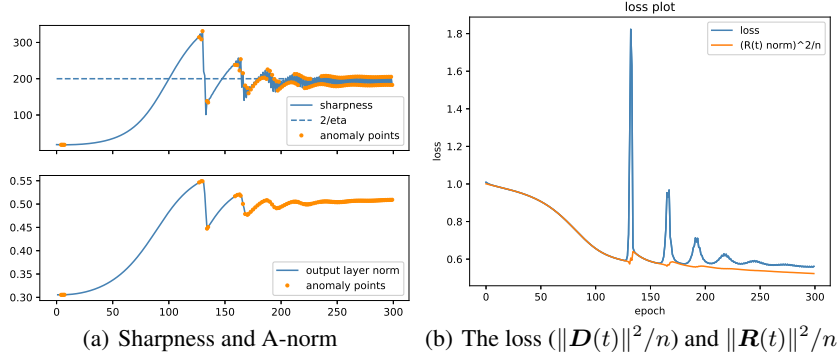


Figure 5: Fully-connected tanh Network. Refer to Figure 4 for more information.

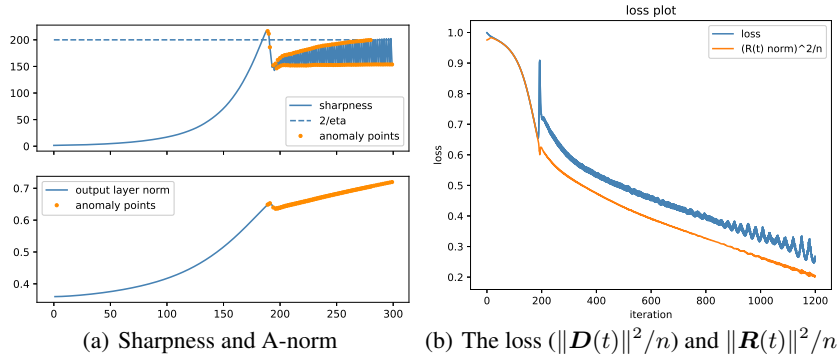


Figure 6: Fully-connected ReLU Network. Refer to Figure 4 for more information.

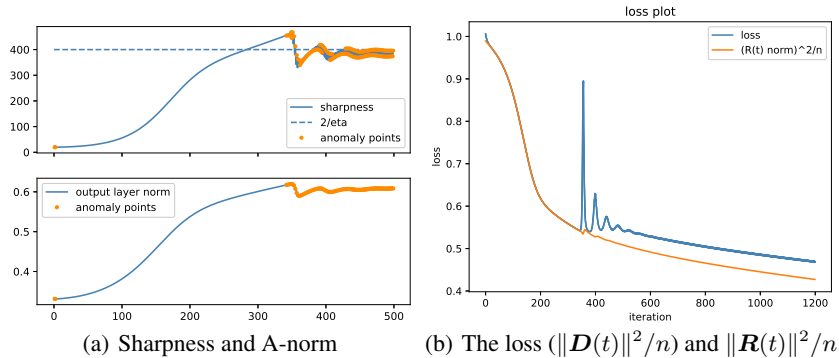


Figure 7: Fully-connected ELU Network. Refer to Figure 4 for more information.

B.1.3 Convolutional Networks

We train three convolutional networks with different activation functions: ReLU, tanh, and ELU activations. In Figure 8, 9, and 10, we verify that in convolutional networks, the positive correlation between the sharpness and the output layer norm (A norm) is still correct in the training process. In the first few oscillations of the sharpness, the four-phase division is also valid. On the other hand, we notice that the same anomaly appears in this convolutional setting as in the fully-connected examples. We defer the discussion to Section B.2.

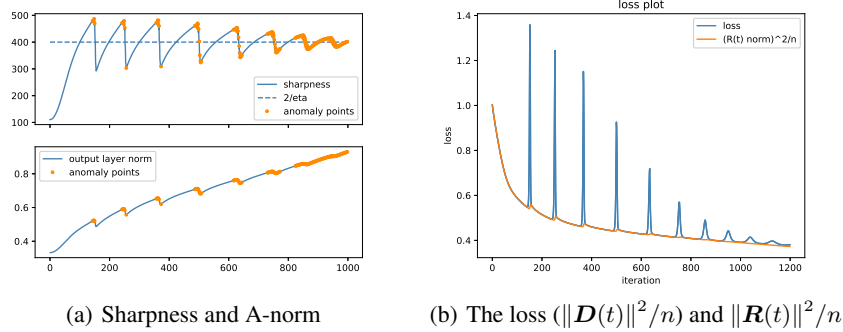


Figure 8: Conv. tanh Network. In Figure (a), we verify the correlation between sharpness and A-norm in convolutional network setting. In Figure (b), we show that in this case, even though the total loss $\mathcal{L}(t) = \|\mathbf{D}(t)\|^2/n$ oscillates considerably, $\|\mathbf{R}(t)\|^2/n$ decreases more steadily.

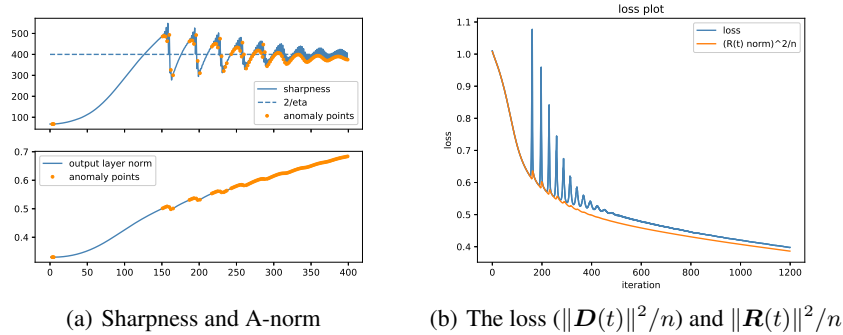


Figure 9: Conv. ReLU Network. Refer to Figure 8 for more information.




Figure 10: Conv. ELU Network. Refer to Figure 8 for more information.

B.1.4 Gaussian Data on Linear Networks

In this section, we train two-layer fully-connected linear networks with datapoints sampled from Gaussian distribution and different label vectors.

In Figure 11, we train two-layer fully-connected linear networks with datapoints sampled from Gaussian distribution. In particular, the width of the network is 200. The data \mathbf{X} is 1000 datapoints sampled from $\mathcal{N}(0, \mathbb{I}_{3072})$, where $\mathbb{I}_{3072} \in \mathbb{R}^{3072 \times 3072}$ is the identity matrix. The label \mathbf{Y} is uniformly sampled from $\text{Unif}\{-1, 1\}^n$.

In Figure 11, we verify that even with simple two-layer linear network and Gaussian data, progressive sharpening and EOS can still be observed. However, because Gaussian data is easy for the network to learn, the convergence is so fast that within tens of epochs the training loss converges to zero. In this case, the EOS phenomenon is not quite typical.

To further explore the effect of different factors on the degree of progressive sharpening, we train the network with data points sampled from Gaussian distribution and different label vectors. In particular, the width of the network is 400. The input \mathbf{X} consists of 500 data points sampled from $\mathcal{N}(\mu, \Sigma)$, where the mean $\mu \in \mathbb{R}^{3072}$ and the covariance $\Sigma \in \mathbb{R}^{3072 \times 3072}$ are the mean vector and the covariance matrix of 5000 CIFAR-10 data points, respectively. We note that to illustrate the degree of progressive sharpening, we choose a very small learning rate $\eta = 2/20000$ (so that the training converges before EOS can happen).

In Figure 12(a), the label \mathbf{Y} is uniformly sampled from $\text{Unif}\{-1, 1\}^n$. Comparing the experimental results with those obtained using the same network trained with 500 CIFAR-10 data points, we find out that both have a similar degree of progressive sharpening. In Figure 12(b) and Figure 12(c), we let the label \mathbf{Y} be $\sqrt{n}v_1$ and $\sqrt{n}v_{400}$ respectively. The result shows that when the label vector is aligned with the top eigenvector v_1 , the degree of progressive sharpening is relatively small, and when the label vector is aligned with a bottom eigenvector, the degree of progressive sharpening is relatively large.

This phenomenon is consistent with our intuition. Empirically we found that a dataset that is easier to learn leads to faster convergence rate, which then leads to a smaller degree of progressive sharpening. For example, training standard Gaussian distributed dataset converges faster than that with CIFAR-10 dataset, hence the degree of progressive sharpening of the former is much smaller; as another example (Figure 12(b) and Figure 12(c)), if the label vector is aligned with the top eigenvector, the convergence in the first phase (the PS phase) is faster, which leads to a shorter first phase and thus a smaller degree of progressive sharpening compared to the other case.

B.2 Further Discussions on the Relation between A-norm and Sharpness

In our paper, we use the observation that $\|\mathbf{A}\|^2$ shares the same trend as the sharpness to explain the dynamics of the sharpness. However, as shown in the experiments in this section (see Figure 5, 6, 7, 8, 9, 10), anomaly points exist during the training process. Here we briefly discuss these anomaly points further. We divide them into three kinds.

At the time that $\|\mathbf{A}\|^2$ changes its trend: This kind of anomaly points appear when $\|\mathbf{A}\|^2$ changes its trend. When $\|\mathbf{A}\|^2$ changes its trend, its changing rate (or differential, the $\mathbf{D}^\top \mathbf{F}$ term) changes its sign, hence the changing rate’s absolute value is small. Now because the dynamics of the sharpness is effected by both $\|\mathbf{A}\|^2$ and the inner layers, when changing rate of $\|\mathbf{A}\|^2$ is small, the inner layers may play a larger role in the direction of sharpness, which may cause the anomaly.

When the sharpness oscillates more frequently: We notice that in some cases, $\|\mathbf{A}(t)\|^2$ and the sharpness $\Lambda(t)$ have very similar overall trend, but the sharpness oscillates more frequently but the magnitude is small (i.e., the sharpness curve has higher frequency oscillations. See Figure 6 or Figure 10). In this case, while we believe the change of $\|\mathbf{A}\|^2$ is a major driving force of the change


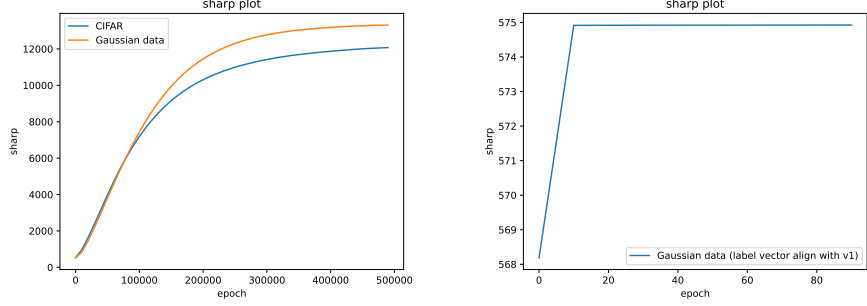
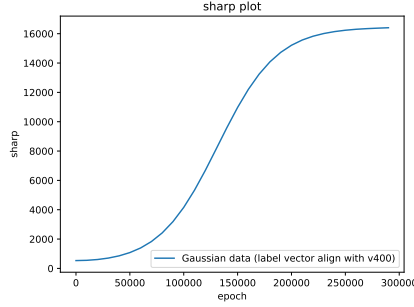


Figure 11: Two-layer Linear Network. Trained with data sampled from the standard Gaussian distribution. Both progressive sharpening and EOS are observed.



(a) Gaussian data with random label vector, compared with CIFAR (b) Gaussian data with label vector align with v_1



(c) Gaussian data with label vector align with v_{400}

Figure 12: progressive sharpening of two-layer fully-connected linear networks trained with data points sampled from anisotropic Gaussian distribution. In Figure (a), the label \mathbf{Y} is uniformly sampled from $\text{Unif}\{-1, 1\}^n$ and the result is compared with that trained with CIFAR data points. In Figure (b) and (c), the label \mathbf{Y} takes $\sqrt{n}v_1$ and $\sqrt{n}v_{400}$ respectively.

of the sharpness, other layers must be taken into consideration for understanding such small and frequent oscillations of sharpness.

In late training phases: We notice that in the late training phases in most settings, sharpness oscillates and crosses $2/\eta$ more frequently (it changes directions in a few iterations). In this case, our four-phase division does not strictly apply. At the same time, $\|\mathbf{A}\|^2$ may also change direction more frequently, resulting more anomaly points during this period of time. Hence, understanding the behavior of the late stages is beyond our current analysis and requires new insights.

C Proof for Section 4

C.1 Detailed Settings

Model: In this section, we study a two-layer neural network with linear activation, i.e.

$$f(\mathbf{x}) = \sum_{q=1}^m \frac{1}{\sqrt{m}} a_q \mathbf{w}_q \mathbf{x} = \frac{1}{\sqrt{m}} \mathbf{A}^\top \mathbf{W} \mathbf{x}$$

where $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_m]^\top \in \mathbb{R}^{m \times d}$ is the hidden layer’s weight matrix, and $\mathbf{A} = [a_1, a_2, \dots, a_m] \in \mathbb{R}^m$ is the weight vector of the output layer. $\mathbf{x} \in \mathbb{R}^d$ is the input vector.

Input distribution: Denote by $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$ the training data matrix and by $\mathbf{Y} = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n$ the label vector. We assume $y_i = \pm 1$ for all $i \in [n]$, and $\|\mathbf{X}^\top \mathbf{X}\| = \Theta(n)$.⁵ We

⁵This property is mentioned in Hu et al. [12]. Empirically, for a randomly selected k -sample subset from CIFAR-10, $\|\mathbf{X}^\top \mathbf{X}\|/k$ is nearly constant.

assume $\mathbf{X}^\top \mathbf{X}$ has rank r , and we decompose $\mathbf{X}^\top \mathbf{X}$ and \mathbf{Y} according to the orthonormal basis $\{\mathbf{v}_i\}$, the eigenvectors of $\mathbf{X}^\top \mathbf{X}$: $\mathbf{X}^\top \mathbf{X} = \sum_{i=1}^r \lambda_i \mathbf{v}_i \mathbf{v}_i^\top$, $\mathbf{Y} = \sum_{i=1}^r (\mathbf{Y}^\top \mathbf{v}_i) \mathbf{v}_i := \sum_{i=1}^r z_i \mathbf{v}_i$ where \mathbf{v}_i is the eigenvector corresponding to the i -th largest eigenvalue λ_i of $\mathbf{X}^\top \mathbf{X}$. $z_i = \mathbf{Y}^\top \mathbf{v}_i$ is the projection of \mathbf{Y} onto the direction \mathbf{v}_i . Here we assume that $n \gg r$.

Also, we suppose there exists some parameter \mathbf{A}^* , \mathbf{W}^* , s.t.

$$\mathbf{Y}^\top = \mathbf{A}^{*\top} \mathbf{W}^* \mathbf{X}$$

Note when $n \gg d$, the matrix $\mathbf{X}^\top \mathbf{X}$ cannot be full rank. This condition guarantees that Gradient Descent (GD) only travels in the column space of $\mathbf{X}^\top \mathbf{X}$ according to Lemma E.10 at Appendix E.

Define the output vector as \mathbf{F} and the residual vector as \mathbf{D} in the training process.

$$\mathbf{F} = (f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_n)) \in \mathbb{R}^n, \mathbf{D} = \mathbf{F} - \mathbf{Y}$$

(Note that they are functions of time t . So we use $\mathbf{D}(t)$, $\mathbf{F}(t)$, $\mathbf{A}(t)$, $\mathbf{W}(t)$ to denote \mathbf{D} , \mathbf{F} , \mathbf{A} , \mathbf{W} at time t .)

Consider the mean square error MSE loss during the training process:

$$\mathcal{L}(\mathbf{A}, \mathbf{W}) = \frac{1}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2 = \frac{1}{n} \|\mathbf{D}\|^2$$

Initialization: We run GD on the loss and start from *symmetric initialization* for the weights, which guarantees $\mathbf{F}(0) = 0$.

$$\begin{aligned} a_q &\sim \text{Unif}(\{-1, 1\}), \quad a_{q+m/2} = -a_q, \quad q = 1, 2, \dots, m/2 \\ \sum_{q=1}^{m/2} \mathbf{w}_q \mathbf{w}_q^\top &= \frac{m}{2d} \mathbf{I}_d, \quad \mathbf{w}_{q+m/2} = \mathbf{w}_q, \quad q = 1, 2, \dots, m/2 \end{aligned}$$

where $\mathbf{I}_d \in \mathbb{R}^{d \times d}$ is the identity matrix of d -dimension.

Sharpness: Recall the definition of the sharpness: $\Lambda(t) = \lambda_{\max}(\mathbf{M}(t))$, where $\mathbf{M}(t)$ is the Gram matrix (8).

Learning rate selection: We select a learning rate η such that $\Lambda(0) < 2/\eta$. Specially, based on our method of initialization, $\Lambda(0) = \frac{m\lambda_1(d+1)}{d} \cdot \frac{2}{mn}$, hence we have

$$\eta < \frac{nd}{(d+1)\lambda_1} \quad (5)$$

Gradient Descent Update Rule: Following GD, the training dynamics is as follows:

$$a_q(t+1) - a_q(t) = -\eta \frac{\partial \mathcal{L}}{\partial a_q} = -\frac{2\eta}{n\sqrt{m}} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i) \mathbf{w}_q \mathbf{x}_i, \quad \forall q \in [m]$$

Similarly, we have

$$\mathbf{w}_q(t+1) - \mathbf{w}_q(t) = -\frac{2\eta}{n\sqrt{m}} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i) a_q \mathbf{x}_i, \quad \forall q \in [m]$$

Write them into matrix forms and we have:

$$\begin{aligned} \mathbf{A}(t+1) - \mathbf{A}(t) &= -\frac{2\eta}{n\sqrt{m}} \mathbf{W}(t) \mathbf{X} \mathbf{D}(t) \\ \mathbf{W}(t+1) - \mathbf{W}(t) &= -\frac{2\eta}{n\sqrt{m}} \mathbf{A}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \end{aligned} \quad (6)$$

Then we have the $\|\mathbf{A}(t)\|$'s dynamics according to (6):

$$\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2 = -2\eta \mathbf{A}(t)^\top \frac{\partial \mathcal{L}}{\partial \mathbf{A}} + \eta^2 \left\| \frac{\partial \mathcal{L}}{\partial \mathbf{A}} \right\|^2 = -\frac{4\eta}{n} \mathbf{F}(t)^\top \mathbf{D}(t) + \eta^2 \left\| \frac{\partial \mathcal{L}}{\partial \mathbf{A}} \right\|^2 \quad (7)$$

where $\mathbf{F}(t)^\top \mathbf{D}(t) = \mathbf{D}(t)^\top \mathbf{F}(t) \in \mathbb{R}$ is a real number.

Define the Gram matrix as:

$$\mathbf{M}(t) = \frac{2}{mn} (\|\mathbf{A}(t)\|^2 \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{W}(t)^\top \mathbf{W}(t) \mathbf{X}) \quad (8)$$

We can also show the dynamics of $\mathbf{D}(t)$ according to (6) and (8). See Appendix E for a proof.

Lemma C.1. (See Lemma E.8) *The update rule of the residual vector of $\mathbf{D}(t)$:*

$$\mathbf{D}(t+1)^\top - \mathbf{D}(t)^\top = -\eta \mathbf{D}(t)^\top \mathbf{M}(t) + \frac{4\eta^2}{n^2 m} \mathbf{D}(t)^\top (\mathbf{F}(t)^\top \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X} \quad (9)$$

We define some extra notations for preparation.

$$\begin{aligned} \mathbf{M}^*(t) &= \mathbf{M}(t) - \frac{4\eta}{n^2 m} (\mathbf{D}(t)^\top \mathbf{F}(t)) \mathbf{X}^\top \mathbf{X} \\ \mathbf{\Gamma}(t) &= \frac{2}{mn} (\mathbf{X}^\top \mathbf{W}(t)^\top \mathbf{W}(t) \mathbf{X} - \frac{m}{d} \mathbf{X}^\top \mathbf{X}) \\ \widetilde{\mathbf{M}}(t) &= \mathbf{M}^*(t) - \mathbf{\Gamma}(t) \end{aligned} \quad (10)$$

Remark. Here we explain the notations above and what they are for.

- $\mathbf{M}^*(t)$ is the Gram matrix $\mathbf{M}(t)$ plus a second order term $(-\frac{4\eta}{n^2 m} (\mathbf{D}(t)^\top \mathbf{F}(t)) \mathbf{X}^\top \mathbf{X})$. Since the second order term is small, we have $\mathbf{M}(t) \approx \mathbf{M}^*(t)$. Also, in Section C.3 we prove $\mathbf{M}^*(t)$ is almost a linear interpolation of $\mathbf{M}(t)$ and $\mathbf{M}(t+1)$, which corroborates this argument.
- $\mathbf{\Gamma}(t)$ is the difference between $\frac{2}{mn} \mathbf{X}^\top \mathbf{W}(t)^\top \mathbf{W}(t) \mathbf{X}$ and $\frac{2}{mn} \mathbf{X}^\top \mathbf{W}(0)^\top \mathbf{W}(0) \mathbf{X}$. Note that $\mathbf{\Gamma}(0) = 0$ because of the initialization condition. In the following proof, we will show a small upper bound of the norm of $\mathbf{\Gamma}(t)$. Hence, we can show that the eigenvectors of $\mathbf{M}(t)$ are approximately aligned with the eigenvectors of $\mathbf{X}^\top \mathbf{X}$.
- $\widetilde{\mathbf{M}}(t)$ is the Gram matrix with the second order term, yet with $\mathbf{\Gamma}(t)$ excluded. $\widetilde{\mathbf{M}}(t)$ is used for bounding the main part of $\mathbf{\Gamma}(t)$ (i.e., the terms without the noise $\mathbf{e}(t)$, e.g. (A1),(A5)) which is defined in Theorem 3) in the proof of Theorem 3. Note that there exists some number γ s.t. $\widetilde{\mathbf{M}}(t) = \gamma \mathbf{X}^\top \mathbf{X}$. This is because

$$\widetilde{\mathbf{M}}(t) = \mathbf{M}^*(t) - \mathbf{\Gamma}(t) = \frac{2}{mn} (\|\mathbf{A}(t)\|^2 + \frac{m}{d} - \frac{2\eta}{n} (\mathbf{D}(t)^\top \mathbf{F}(t))) \mathbf{X}^\top \mathbf{X} \quad (11)$$

Hence for any eigenvector \mathbf{v}_i of $\mathbf{X}^\top \mathbf{X}$, $\lambda_i(\widetilde{\mathbf{M}}(t)) = \mathbf{v}_i^\top \widetilde{\mathbf{M}}(t) \mathbf{v}_i$.

Then the dynamics of $\mathbf{D}(t)$ can be written as:

$$\mathbf{D}(t+1) = (\mathbf{I}_n - \eta \mathbf{M}^*(t)) \mathbf{D}(t) \quad (12)$$

Update Rule of $\mathbf{M}(t)$: Using the update rule of $\mathbf{A}(t)$ and $\mathbf{W}(t)$, we can also derive the update rule of the Gram matrix $\mathbf{M}(t)$ (see Appendix E for the proof detail):

Lemma C.2. (See Lemma E.9) *The update rule of the Gram matrix $\mathbf{M}(t)$ is,*

$$\begin{aligned} \mathbf{M}(t+1) - \mathbf{M}(t) &= -\frac{4\eta}{n^2 m} (2(\mathbf{F}(t)^\top \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X} + \mathbf{F}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{F}(t)^\top) \\ &\quad + \frac{8\eta^2}{n^3 m^2} (\mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}(t)^\top \mathbf{W}(t) \mathbf{X} \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X} \\ &\quad + \frac{8\eta^2}{n^3 m^2} \|\mathbf{A}(t)\|^2 \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} \end{aligned}$$

C.2 Phase I and Progressive Sharpening

We suppose $\mathbf{X}^\top \mathbf{X}$ has rank r , and we decompose $\mathbf{X}^\top \mathbf{X}$ and \mathbf{Y} into a basis $\{\mathbf{v}_i\}$ composed of orthonormal eigenvectors of $\mathbf{X}^\top \mathbf{X}$.

$$\mathbf{X}^\top \mathbf{X} = \sum_{i=1}^r \lambda_i \mathbf{v}_i \mathbf{v}_i^\top, \mathbf{Y} = \sum_{i=1}^r (\mathbf{Y}^\top \mathbf{v}_i) \mathbf{v}_i := \sum_{i=1}^r z_i \mathbf{v}_i$$

where \mathbf{v}_i is the eigenvector corresponding to the i -th largest eigenvalue λ_i of $\mathbf{X}^\top \mathbf{X}$. $z_i = \mathbf{Y}^\top \mathbf{v}_i$ is the \mathbf{Y} 's projection onto the direction \mathbf{v}_i . In particular, \mathbf{v}_1 is the eigenvector corresponding to the largest eigenvalue of $\mathbf{X}^\top \mathbf{X}$.

Recall that the sharpness $\Lambda(t) = \lambda_{\max}(\mathbf{M}(t))$. Here we propose an approximation of sharpness Λ by the eigenvector \mathbf{v}_1 :

$$\Lambda^*(t) = \mathbf{v}_1^\top \mathbf{M}(t) \mathbf{v}_1 \quad (13)$$

The following corollary rationalizes this approximation. The proof is deferred to Corollary C.2 in Section C.3.

Corollary C.1. *There exists some constant c_2, c_4 , s.t. for all time t $\langle \mathbf{v}_1, \mathbf{u} \rangle^2 \geq 1 - c_4 \sqrt{\frac{1}{nm}}$ and*

$$\Lambda(t) \geq \Lambda^*(t) \geq \Lambda(t) - \frac{2c_2}{m}$$

where \mathbf{u} is the top eigenvector of $\mathbf{M}(t)$.

By this corollary, we can see that the top eigenvector of $\mathbf{M}(t)$ is approximately \mathbf{v}_1 , and the approximation $\Lambda^*(t)$ is very close to the real sharpness $\Lambda(t)$. Thus, we consider $\Lambda^*(t)$ as an approximation of the sharpness and analyze its dynamics.

We first show the main theorem of progressive sharpening under Assumption C.1 and C.2 below.

Assumption C.1. *There exists some constant $\chi > 1$, s.t. for all $i \in [r-1]$, $\lambda_i(\mathbf{X}^\top \mathbf{X}) \leq \chi \lambda_{i+1}(\mathbf{X}^\top \mathbf{X})$. Moreover, $\lambda_1(\mathbf{X}^\top \mathbf{X}) \geq 2\lambda_2(\mathbf{X}^\top \mathbf{X})$.*

Assumption C.2. *There exists $\kappa = \Omega(r^{-1})$ such that $\min_{i \in [r]} \{z_i / \sqrt{n}\} \geq \kappa$.*

The first assumption is about the eigenvalue spectrum of $\mathbf{X}^\top \mathbf{X}$. It guarantees the gap between two adjacent eigenvalues is not very large, and there is a gap between the largest and the second largest eigenvalue. Note the second part of the assumption is a relaxed version of Assumption 3.7. In our CIFAR-10 1k-subset with samples' mean subtracted, $\lambda_1/\lambda_2 = \chi \approx 3$ (See Appendix D.2.1, Figure 19). The second assumes that all components $z_i = \mathbf{Y}^\top \mathbf{v}_i$ are not too small.

Under the two assumptions, we have the following theorem:

Theorem 3 (Progressive Sharpening). *Suppose Assumption C.1, C.2 hold. Suppose $\lambda_r = \lambda_{\min}(\mathbf{X}^\top \mathbf{X}) > 0$, $\lambda_1 = \lambda_{\max}(\mathbf{X}^\top \mathbf{X}) = \Theta(n)$. For any $\epsilon > 0$, if $m = \Omega(\frac{n^2}{\chi^2})$, and $n = \max\{\Omega(\frac{\lambda_r}{\kappa^2}), \Omega(\frac{\lambda_r^2}{\kappa^2}), \Omega(\frac{\lambda_r^2}{\kappa^4 \epsilon^2}), \Omega(\frac{\lambda_r^4}{\kappa^4 \epsilon^2})\}$, we have the following properties for $t = 1, 2, \dots, t_0 - 1$ where t_0 is the time when $\|\mathbf{D}(t)\|^2 = O(\epsilon^2)$ or $\lambda_{\max}(\mathbf{M}^*(t)) > 1/\eta$ for the first time.*

- (Progressive Sharpening) (Lemma C.6) $\Lambda^*(t+1) - \Lambda^*(t) > 0$.
- (Lemma C.3) $\|\Gamma(t)\| = O(1/m)$.

We first prove Lemma C.3, which implies that $\mathbf{D}(t)$ approximately converges independently in each direction of $\mathbf{X}^\top \mathbf{X}$'s non-zero eigenspace. Lemma C.3 also proves the second point $\|\Gamma(t)\| \leq O(1/m)$. (Recall that $\|\Gamma(t)\|$ is defined in (10).)

Then, based on the conclusion in Lemma C.3, Assumption C.1 and C.2, we prove that $\|\mathbf{A}(t)\|^2$ grows in Lemma C.4 and Lemma C.5. Specifically, here we prove a sufficient condition of $-\mathbf{D}(t)^\top \mathbf{F}(t) > 0$.

Finally, we use the dynamics of $\Lambda^*(t)$ and its dependence on $\|\mathbf{A}(t)\|^2$ dynamics to prove the sharpness grows (Lemma C.6), thus finishing the proof.

Remark 1. In the theorem, $\lambda_{\max}(\mathbf{M}^*(t)) > 1/\eta$ is the termination condition for this theorem. Here, $\mathbf{M}^*(t)$ is the Gram matrix $\mathbf{M}(t)$ plus the second order term $(-\frac{4\eta}{n^2m}(\mathbf{D}(t)^\top \mathbf{F}(t))\mathbf{X}^\top \mathbf{X})$. Since the second order term is small, we have $\mathbf{M}(t) \approx \mathbf{M}^*(t)$. Also, in Section C.3, we prove $\mathbf{M}^*(t)$ is almost a linear interpolation of $\mathbf{M}(t)$ and $\mathbf{M}(t+1)$. Thus, $\lambda_{\max}(\mathbf{M}^*(t))$ can be seen as an approximation of the sharpness.

Remark 2. From the experiments, one can see that gradient descent is in progressive sharpening phase until the sharpness crosses the threshold $2/\eta$. Right now, our proof only works till $\lambda_{\max}(\mathbf{M}^*(t))$ reaches $1/\eta$. It would be interesting to extend our result to $2/\eta$.

Remark 3. In this theorem, we require mild over-parameterization ($m = \Omega(n^2)$), assuming $\lambda_r = \Theta(1)$, to prove the direction guarantee of the Gram matrix $\mathbf{M}(t)$ and the monotone increment of $\|\mathbf{A}(t)\|^2$. Astute readers may find it similar to the NTK regime, where the parameters do not move far from the initialization. However, we stress that our analysis does not necessarily require NTK regime, and can go beyond NTK. We defer the discussion on the difference between our results and the NTK regime to Appendix C.4.

Then we break the proof of theorem into lemmas. The first lemma (Lemma C.3) proves that the Gram matrix $\mathbf{M}(t) \approx \frac{2}{mn}(\|\mathbf{A}(t)\|^2 + \frac{m}{d})\mathbf{X}^\top \mathbf{X}$ by bounding the movement of $\|\mathbf{\Gamma}(t)\|$. In this way, $\mathbf{D}(t) \approx -\prod_{j=0}^{t-1}(\mathbf{I}_n - \eta\widetilde{\mathbf{M}}(j))\mathbf{Y}$ can approximately descend in each eigenvector \mathbf{v}_i of $\mathbf{X}^\top \mathbf{X}$ independently. Also, by proving $\mathbf{M}(t) \approx \frac{2}{mn}(\|\mathbf{A}(t)\|^2 + \frac{m}{d})\mathbf{X}^\top \mathbf{X}$, we justify $\|\mathbf{A}(t)\|^2$ as an indicator of the sharpness.

Lemma C.3 (Direction Guarantee). *Along GD training trajectory, if $m \geq \frac{112c_1n^2}{\lambda_r^2}$, we have the following properties until $\lambda_{\max}(\mathbf{M}^*(t)) > 1/\eta$:*

1. $\|\mathbf{A}(t)\|^2 \geq m/2$ and $\|\mathbf{\Gamma}(t)\| \leq R_w := 40/m$.
2. $\mathbf{D}(t) = -\prod_{j=0}^{t-1}(\mathbf{I}_n - \eta\widetilde{\mathbf{M}}(j))\mathbf{Y} + \mathbf{e}(t)$, where $\|\mathbf{e}(t)\|_2 \leq \frac{40n^{3/2}}{\lambda_r m}$
3. $\lambda_r(\widetilde{\mathbf{M}}(t)) \geq \lambda_r/n$, $\lambda_r(\mathbf{M}^*(t)) \geq \lambda_r/n$.

Proof. We prove the theorem by induction.

We first consider the base case. For property 1, $\|\mathbf{A}(0)\|^2 = m$ and $\|\mathbf{\Gamma}(0)\| = 0$. For property 2, with symmetric initialization, $\mathbf{F}(0) = 0$, hence $\mathbf{D}(0) = -\mathbf{Y}$, $\|\mathbf{e}(0)\| = 0$. For property 3,

$$\mathbf{M}^*(0) = \widetilde{\mathbf{M}}(0) = \mathbf{M}(0) \succeq \frac{2}{nm}\|\mathbf{A}(0)\|^2\mathbf{X}^\top \mathbf{X}$$

The minimal eigenvalue of $\frac{2}{nm}\|\mathbf{A}(0)\|^2\mathbf{X}^\top \mathbf{X}$ is $2\lambda_r/n > \lambda_r/n$. Thus all the properties hold at iteration $t = 0$.

Suppose for all $k \leq t$, these properties hold. Then we consider the case in iteration $t + 1$.

We first show a worst-case upper bound for $\|\mathbf{D}(t)\|$.

$$\begin{aligned} \|\mathbf{D}(t)\| &= \left\| \prod_{k=0}^{t-1} (\mathbf{I} - \eta\mathbf{M}^*(k))\mathbf{D}(0) \right\| \\ &\leq \left\| \prod_{k=0}^{t-1} (\mathbf{I} - \eta\mathbf{M}^*(k)) \right\| \|\mathbf{D}(0)\| \\ &\leq (1 - \frac{\eta\lambda_r}{n})^t \|\mathbf{Y}\| = \sqrt{n}(1 - \frac{\eta\lambda_r}{n})^t \end{aligned} \tag{D}$$

The second inequality uses property 3 in $k \leq t$. Note $\|\mathbf{Y}\| = \sqrt{n}$ in our setting, so the last equality holds. That means $\|\mathbf{D}(t)\| \leq \sqrt{n}$ and $\|\mathbf{F}(t)\| = \|\mathbf{D}(t) + \mathbf{Y}\| \leq 2\sqrt{n}$.

Now, we show the error $\mathbf{e}(t)$ is bounded. We have

$$\begin{aligned}
\mathbf{D}(t+1) &= (\mathbf{I} - \eta \widetilde{\mathbf{M}}(t))\mathbf{D}(t) - \eta \boldsymbol{\Gamma}(t)\mathbf{D}(t) \\
&= - \prod_{j=0}^t (\mathbf{I} - \eta \widetilde{\mathbf{M}}(j))\mathbf{Y} + (\mathbf{I} - \eta \widetilde{\mathbf{M}}(t))\mathbf{e}(t) - \eta \boldsymbol{\Gamma}(t)\mathbf{D}(t) \\
&= - \prod_{j=0}^t (\mathbf{I} - \eta \widetilde{\mathbf{M}}(j))\mathbf{Y} + (\mathbf{I} - \eta \mathbf{M}^*(t))\mathbf{e}(t) - \eta \boldsymbol{\Gamma}(t) \prod_{j=0}^{t-1} (\mathbf{I} - \eta \widetilde{\mathbf{M}}(j))\mathbf{D}(0).
\end{aligned}$$

Hence we have $\mathbf{e}(t+1) = \mathbf{e}(t)(\mathbf{I} - \eta \mathbf{M}^*(t)) - \eta \boldsymbol{\Gamma}(t) \prod_{j=0}^{t-1} (\mathbf{I} - \eta \widetilde{\mathbf{M}}(j))\mathbf{D}(0)$.

Then we can have the following bound by this recursion:

$$\begin{aligned}
\|\mathbf{e}(t+1)\|_2 &\leq \|\mathbf{e}(t)\|_2 \|\mathbf{I} - \eta \mathbf{M}^*(t)\|_2 + \eta \|\mathbf{D}(0)\|_2 \left\| \prod_{j=0}^{t-1} (\mathbf{I} - \eta \widetilde{\mathbf{M}}(j)) \right\|_2 \|\boldsymbol{\Gamma}(t)\|_2 \\
&\leq \|\mathbf{e}(t)\|_2 \left(1 - \frac{\eta \lambda_r}{n}\right) + \eta \sqrt{n} \left(1 - \frac{\eta \lambda_r}{n}\right)^t R_w \\
&\leq \sum_{k=0}^t \eta \sqrt{n} \left(1 - \frac{\eta \lambda_r}{n}\right)^k R_w \tag{E1} \\
&\leq \frac{\sqrt{n}}{\frac{1}{n} \lambda_r} \cdot R_w \\
&\leq \frac{40n^{3/2}}{m \lambda_r}
\end{aligned}$$

Here the third inequality holds because $\|\mathbf{e}(0)\| = 0$. Thus the second property holds for $t+1$.

Then we consider the lower bound of $\|\mathbf{A}(t)\|^2$.

Consider the dynamics of $\|\mathbf{A}(t)\|^2$ in (7) and sum the difference up from 0 to t . Recall that we proved $\|\mathbf{D}(t)\| \leq \sqrt{n}$ by (D) above.

$$\begin{aligned}
\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(0)\|^2 &= \frac{4\eta}{n} \sum_{k=0}^t (-\mathbf{D}(k)^\top \mathbf{F}(k)) + \eta^2 \sum_{k=0}^t \left\| \frac{\partial \mathcal{L}(k)}{\partial \mathbf{A}(k)} \right\|^2 \\
&\geq -\frac{4\eta}{n} \sum_{k=0}^t \mathbf{D}(k)^\top (\mathbf{D}(k) + \mathbf{Y}) \quad \left[\left\| \frac{\partial \mathcal{L}(k)}{\partial \mathbf{A}(k)} \right\|^2 > 0 \right] \\
&\geq -\frac{4\eta}{n} \sum_{k=0}^t (\|\mathbf{D}(k)\|^2 + \|\mathbf{D}(k)\| \|\mathbf{Y}\|) \quad [\ell_2\text{-norm inequality}] \\
&\geq -\frac{4\eta}{n} \sum_{k=0}^t \|\mathbf{D}(k)\| (\sqrt{n} + \sqrt{n}) \quad [\|\mathbf{D}(t)\| \leq \sqrt{n}] \\
&\geq -\frac{8\eta}{n} \sum_{k=0}^t n \left(1 - \frac{\eta \lambda_r}{n}\right)^k \quad [\text{by (D)}] \\
&\geq -\frac{8n}{\lambda_r} \geq -\frac{m}{2} \tag{\#}
\end{aligned}$$

Since $\|\mathbf{A}(0)\|^2 = m$, we have the lower bound $\|\mathbf{A}(t)\|^2 \geq m/2$.

Next, we lower bound the minimal non-zero eigenvalue (i.e. the r -th largest eigenvalue since $\mathbf{X}^\top \mathbf{X}$ is rank r) of $\mathbf{M}^*(t+1)$ and $\widetilde{\mathbf{M}}(t+1)$.

For $\mathbf{M}^*(t+1)$ and $\widetilde{\mathbf{M}}(t+1)$, since the $\mathbf{X}^\top \mathbf{W}(t)^\top \mathbf{W}(t) \mathbf{X}$ part is PSD, we have

$$\begin{aligned} \mathbf{M}^*(t+1) &\succeq \frac{2}{mn} \|\mathbf{A}(t+1)\|^2 \mathbf{X}^\top \mathbf{X} - \frac{4\eta}{n^2 m} (\mathbf{D}(t+1)^\top \mathbf{F}(t+1)) \mathbf{X}^\top \mathbf{X} \\ \widetilde{\mathbf{M}}(t+1) &\succeq \frac{2}{mn} \|\mathbf{A}(t+1)\|^2 \mathbf{X}^\top \mathbf{X} - \frac{4\eta}{n^2 m} (\mathbf{D}(t+1)^\top \mathbf{F}(t+1)) \mathbf{X}^\top \mathbf{X} \end{aligned}$$

From the inequality in (#), and from (D) $\|\mathbf{D}(t)\| \leq \sqrt{n}$, $\|\mathbf{F}(t)\| \leq 2\sqrt{n}$, we know

$$\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(0)\|^2 - \frac{2\eta}{n} (\mathbf{D}(t+1)^\top \mathbf{F}(t+1)) \geq -\frac{8n}{\lambda_r} - \frac{2\eta}{n} \sqrt{n} \cdot 2\sqrt{n} \geq -m/2$$

$$\mathbf{M}^*(t+1) \succeq \frac{2}{mn} (\|\mathbf{A}(t+1)\|^2 - \frac{2\eta}{n} (\mathbf{D}(t+1)^\top \mathbf{F}(t+1))) \mathbf{X}^\top \mathbf{X} \succeq \frac{1}{n} \mathbf{X}^\top \mathbf{X}$$

Thus $\mathbf{M}^*(t+1)$ (and also $\widetilde{\mathbf{M}}(t+1)$) has its smallest eigenvalue larger than λ_r/n . Property 3 holds. To extend this conclusion, we actually have $\lambda_i(\widetilde{\mathbf{M}}(t+1)) \geq \lambda_i/n$ for all $i \in [r]$ from the argument above.

Finally we bound $\|\Gamma(t+1)\|$. We first write down the dynamics of $\Gamma(t)$ according to Lemma C.2.

$$\begin{aligned} \Gamma(t+1) - \Gamma(t) &= \frac{2}{nm} (\mathbf{X}^\top (\mathbf{W}(t+1)^\top \mathbf{W}(t+1) - \mathbf{W}(t)^\top \mathbf{W}(t)) \mathbf{X}) \\ &= \frac{2}{nm} (\mathbf{X}^\top ((\mathbf{W}(t+1)^\top - \mathbf{W}(t)^\top) \mathbf{W}(t) + \mathbf{W}(t)^\top (\mathbf{W}(t+1) - \mathbf{W}(t))) \mathbf{X} \\ &\quad + \mathbf{X}^\top (\mathbf{W}(t+1)^\top - \mathbf{W}(t)^\top) (\mathbf{W}(t+1) - \mathbf{W}(t)) \mathbf{X}) \quad (\text{Use Equation (6)}) \\ &= -\frac{4}{n^2 m} \eta (\mathbf{F}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{F}(t)^\top) \\ &\quad + \frac{8\eta^2}{n^3 m^2} \|\mathbf{A}(t)\|^2 \cdot \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} \end{aligned}$$

Hence we sum it up and get:

$$\begin{aligned} \|\Gamma(t+1) - \Gamma(0)\|_2 &= \left\| \frac{4}{n^2 m} \eta \sum_{k=0}^t (\mathbf{F}(k) \mathbf{D}(k)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(k) \mathbf{F}(k)^\top) \right. \\ &\quad \left. + \frac{8\eta^2}{n^3 m^2} \sum_{k=0}^t \|\mathbf{A}(k)\|^2 \cdot \mathbf{X}^\top \mathbf{X} \mathbf{D}(k) \mathbf{D}(k)^\top \mathbf{X}^\top \mathbf{X} \right\|_2 \quad (\text{Triangle Inequality}) \\ &\leq \frac{4}{n^2 m} \eta \left\| \sum_{k=0}^t (\mathbf{D}(k) \mathbf{D}(k)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(k) \mathbf{D}(k)^\top) \right\|_2 \quad (\text{G1}) \\ &\quad + \frac{4}{n^2 m} \eta \left\| \sum_{k=0}^t (\mathbf{Y} \mathbf{D}(k)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(k) \mathbf{Y}^\top) \right\|_2 \quad (\text{G2}) \\ &\quad + \frac{8\eta^2}{n^3 m^2} \left\| \sum_{k=0}^t \|\mathbf{A}(k)\|^2 \cdot \mathbf{X}^\top \mathbf{X} \mathbf{D}(k) \mathbf{D}(k)^\top \mathbf{X}^\top \mathbf{X} \right\|_2 \quad (\text{G3}) \end{aligned}$$

Then we bound these three terms one by one.

$$\text{Term (G1): } \frac{4}{n^2 m} \eta \left\| \sum_{k=0}^t (\mathbf{D}(k) \mathbf{D}(k)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(k) \mathbf{D}(k)^\top) \right\|_2$$

By symmetry, we just consider the first term $\mathbf{D}^\top(t) \mathbf{D}(t) \mathbf{X}^\top \mathbf{X}$.

$$\mathbf{D}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} = \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{Y} \mathbf{Y}^\top \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{X}^\top \mathbf{X} \quad (\text{A1})$$

$$+ \mathbf{e}(t) \mathbf{Y}^\top \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{X}^\top \mathbf{X} \quad (\text{A2})$$

$$+ \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{Y} \mathbf{e}(t)^\top \mathbf{X}^\top \mathbf{X} \quad (\text{A3})$$

$$+ \mathbf{e}(t) \mathbf{e}(t)^\top \mathbf{X}^\top \mathbf{X} \quad (\text{A4})$$

We bound each term in ℓ_2 -norm and then add them up.

Term (A1)

$$\begin{aligned} & \frac{4\eta}{n^2 m} \left\| \sum_{k=0}^t \prod_{j=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{Y} \mathbf{Y}^\top \prod_{s=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(s)) \mathbf{X}^\top \mathbf{X} \right\|_2 \\ &= \frac{4\eta}{n^2 m} \left\| \sum_{k=0}^t \prod_{j=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \sum_{i=1}^r z_i \mathbf{v}_i \sum_{j'=1}^r z_{j'} \mathbf{v}_{j'}^\top \prod_{s=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(s)) \mathbf{X}^\top \mathbf{X} \right\|_2 \\ &= \frac{4\eta}{n^2 m} \left\| \sum_{k=0}^t \left(\sum_{i=1}^r \prod_{j=0}^{k-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) z_i \mathbf{v}_i \right) \left(\sum_{j'=1}^r \mathbf{v}_{j'}^\top z_{j'} \prod_{s=0}^{k-1} (1 - \eta \lambda_{j'}(\widetilde{\mathbf{M}}(s))) \lambda_{j'} \right) \right\|_2 \\ & \quad \text{where } \lambda_j(\widetilde{\mathbf{M}}) = \mathbf{v}_j^\top \widetilde{\mathbf{M}} \mathbf{v}_j \text{ (See (11)). We know since } \|\mathbf{A}(t)\|^2 \geq m/2, \lambda_j(\widetilde{\mathbf{M}}(t)) \geq \frac{1}{n} \lambda_j, \forall t, j \\ &= \frac{4\eta}{n^2 m} \left\| \sum_{i=1}^r \sum_{j'=1}^r \sum_{k=0}^t \left(\prod_{j=0}^{k-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) z_i z_{j'} (1 - \eta \lambda_{j'}(\widetilde{\mathbf{M}}(j))) \lambda_{j'} \right) \mathbf{v}_i \mathbf{v}_{j'}^\top \right\|_2 \\ &\leq \frac{4\eta}{n^2 m} \left\| \sum_{i=1}^r \sum_{j'=1}^r \sum_{k=0}^t \left(\prod_{j=0}^{k-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) z_i z_{j'} (1 - \eta \lambda_{j'}(\widetilde{\mathbf{M}}(j))) \lambda_{j'} \right) \mathbf{v}_i \mathbf{v}_{j'}^\top \right\|_F \\ &= \frac{4\eta}{n^2 m} \sqrt{\sum_{i=1}^r \sum_{j'=1}^r \left(\sum_{k=0}^t \prod_{j=0}^{k-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) z_i z_{j'} (1 - \eta \lambda_{j'}(\widetilde{\mathbf{M}}(j))) \lambda_{j'} \right)^2} \\ &\leq \frac{4\eta}{n^2 m} \sqrt{\sum_{i=1}^r \sum_{j'=1}^r \left(\sum_{k=0}^t (1 - \eta \frac{1}{n} \lambda_{j'})^k z_i z_{j'} (1 - \eta \frac{1}{n} \lambda_i)^k \lambda_{j'} \right)^2} \\ &\leq \frac{4}{n^2 m} \sqrt{\sum_{i=1}^r \sum_{j'=1}^r z_i^2 z_{j'}^2 \lambda_{j'}^2 \frac{\eta^2}{(\frac{1}{n} \eta (\lambda_i + \lambda_{j'}) - \frac{\eta^2}{n^2} \lambda_i \lambda_{j'})^2}} \\ &\leq \frac{4}{n^2 m} \sqrt{\|\mathbf{Y}\|_2^4 \cdot n^2} \\ &= \frac{4}{m} \end{aligned}$$

The first inequality comes from $\|\cdot\|_2 \leq \|\cdot\|_F$. The last three inequalities require that $\eta \lambda_{\max}(\mathbf{M}^*) < 1$.

Combined with its symmetric counterpart in $\mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{D}(t)^\top$, the sum of (A1) is smaller than $\frac{8}{m}$.

Term (A2) and (A3):

$$\frac{4\eta}{n^2 m} \left\| \sum_{k=0}^t \mathbf{e}(k) \mathbf{Y}^\top \prod_{j=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{X}^\top \mathbf{X} + \prod_{j=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{Y} \mathbf{e}(k)^\top \mathbf{X}^\top \mathbf{X} \right\|_2$$

$$\begin{aligned}
&\leq \frac{8\eta}{n^2 m} \left\| \sum_{k=0}^t \mathbf{e}(k) \mathbf{Y}^\top \prod_{j=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{X}^\top \mathbf{X} \right\|_2 && \text{(Triangle Inequality)} \\
&\leq \frac{8\eta}{n^2 m} \sum_{k=0}^t \|\mathbf{e}(k)\|_2 \left\| \mathbf{Y}^\top \prod_{j=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \right\|_2 \|\mathbf{X}^\top \mathbf{X}\|_2 && \text{(Cauchy-Schwarz)} \\
&\leq \frac{8}{n^2 m} \left(\sum_{k=0}^t \eta (1 - \frac{\eta}{n} \lambda_r)^k \sqrt{n} \right) \cdot \frac{n^{3/2}}{\lambda_r} R_w \cdot c_1 n && \text{((E1) and Algebra)} \\
&\leq \frac{8c_1 n^2}{m \lambda_r^2} R_w
\end{aligned}$$

Similarly, we have their symmetric counterpart added and get a bound of $\frac{16c_1 n^2}{m \lambda_r^2} R_w$.

Term (A4):

$$\begin{aligned}
&\frac{4\eta}{n^2 m} \left\| \sum_{k=0}^t \mathbf{e}(k) \mathbf{e}(k)^\top \mathbf{X}^\top \mathbf{X} \right\|_2 \\
&\leq \frac{4}{n^2 m} \eta \sum_{k=0}^t \|\mathbf{e}(k)\|_2^2 \|\mathbf{X}^\top \mathbf{X}\|_2 && \text{(Cauchy-Schwarz)} \\
&\quad (\|\mathbf{e}(k)\|_2 \leq \|\mathbf{D}(k)\| + \left\| \prod_{j=0}^{k-1} (\mathbf{I} - \eta \widetilde{\mathbf{M}}(j)) \mathbf{Y} \right\| \leq 2\sqrt{n} (1 - \frac{\eta}{n} \lambda_r)^k) \\
&\leq \frac{4\eta}{N^2 m} \sum_{k=0}^t 2(1 - \frac{\eta}{n} \lambda_r)^k \sqrt{n} \cdot \frac{n^{3/2}}{\lambda_r} R_w \cdot c_1 n && \text{((E1) and Algebra)} \\
&\leq \frac{8n^2 c_1}{m \lambda_r^2} R_w
\end{aligned}$$

Similarly, we can combine its symmetric counterpart and get a bound of $\frac{16c_1 n^2}{m \lambda_r^2} R_w$.

Add them up, and we can get the bound of the first part.

$$\|(1)\|_2 \leq \frac{8}{m} + \frac{32c_1 n^2}{m \lambda_r^2} R_w$$

Term (G2):

$$\begin{aligned}
&\frac{4}{n^2 m} \eta \left\| \sum_{k=0}^t (\mathbf{Y} \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{Y}^\top) \right\|_2 && \text{(Triangle Inequality)} \\
&\leq \frac{8}{N^2 m} \eta \left\| \sum_{k=0}^t \mathbf{Y} \mathbf{Y}^\top \prod_{j=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{X}^\top \mathbf{X} \right\|_2 && \text{(A5)}
\end{aligned}$$

$$+ \frac{8}{N^2 m} \left\| \sum_{k=0}^t \mathbf{Y} \mathbf{e}(k)^\top \mathbf{X}^\top \mathbf{X} \right\|_2 \quad \text{(A6)}$$

We bound the two parts separately.

Term (A5):

$$\frac{8}{N^2 m} \eta \left\| \sum_{k=0}^t \mathbf{Y} \mathbf{Y}^\top \prod_{j=0}^{k-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{X}^\top \mathbf{X} \right\|_2 \quad \text{(Cauchy-Schwarz)}$$

we have $-\mathbf{D}(t)^\top \mathbf{F}(t) > 0$.

Proof. We use property 2 of Lemma C.3 to obtain the following expression:

$$\begin{aligned}
-\mathbf{D}(t)^\top \mathbf{F}(t) &= \mathbf{Y}^\top \prod_{j=0}^t (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) (\mathbf{I}_n - \prod_{j=0}^t (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j))) \mathbf{Y} \\
&\quad + \mathbf{Y}^\top (2 \prod_{j=0}^t (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) - \mathbf{I}_n) \mathbf{e}(t) - \|\mathbf{e}(t)\|^2 \quad (\text{Use (E1) and } m \geq \frac{112c_1 n^2}{\lambda_r^2}) \\
&\geq \sum_{i=1}^r z_i^2 \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) (1 - \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j)))) - \frac{40n^2}{\lambda_r m} - \frac{1600n^3}{\lambda_r^2 m^2} \\
&\geq \kappa^2 n \max_i \left\{ \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) (1 - \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j)))) \right\} - \frac{5\lambda_r}{14} - \frac{25\lambda_r^2}{196}
\end{aligned} \tag{C.4}$$

Notice that all inequalities hold since $\lambda_i(\mathbf{M}^*(t)) < 1/\eta$ for all i . In the first inequality, we use the $\|\mathbf{e}(t)\|$ bound in (E1), and in the second we just replace m with its lower bound $\frac{112c_1 n^2}{\lambda_r^2}$. Then we use $n > (70\lambda_r + 25\lambda_r^2)/(196\kappa^2 \min\{a_1 - a_1^2, a_2 - a_2^2\})$ and complete the proof.

$$-\mathbf{D}(t)^\top \mathbf{F}(t) \geq \kappa^2 n \min\{a_1 - a_1^2, a_2 - a_2^2\} - \frac{5\lambda_r}{14} - \frac{25\lambda_r^2}{196} > 0$$

□

Then, we use Assumption C.1 to prove the next lemma. It tells that condition (14) is satisfied under this assumption in a time interval. That means $-\mathbf{D}(t)^\top \mathbf{F}(t) > 0$ during this period.

Lemma C.5. *Under Assumption C.1, we have condition (14) satisfied with $a_2 = e^{(a_1-1)/\chi}$ in the time interval $[t_1, t_2]$. Here, t_1 is the iteration that $\prod_{j=0}^{t-1} (1 - \eta \lambda_1(\widetilde{\mathbf{M}}(j))) < a_2$ for the first time, and t_2 is the iteration when $\prod_{j=0}^{t-1} (1 - \eta \lambda_r(\widetilde{\mathbf{M}}(j))) < a_1$ for the first time.*

Proof. We prove for all $i \in [r-1]$,

$$\text{If } \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) < a_1, \text{ then } \prod_{j=0}^{t-1} (1 - \eta \lambda_{i+1}(\widetilde{\mathbf{M}}(j))) < e^{(a_1-1)/\chi}. \tag{15}$$

In this way, the only two possibility that all $i \in [r]$ doesn't satisfy the condition (14) is: (1) the $\prod_{j=0}^{t-1} (1 - \eta \lambda_1(\widetilde{\mathbf{M}}(j))) > a_2$; (2) $\prod_{j=0}^{t-1} (1 - \eta \lambda_r(\widetilde{\mathbf{M}}(j))) < a_1$. Otherwise, there must be some i s.t. $\prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) \in [a_1, a_2]$. Thus, if condition (15) is satisfied, we have this lemma proved.

Now suppose $\prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) < a_1$. By Bernoulli's inequality,

$$(1 + x_1)(1 + x_2) \dots (1 + x_n) \geq 1 + x_1 + x_2 + \dots + x_n, \text{ if } x_i \geq -1, \forall i \in [n]$$

since $-\eta \lambda_i(\widetilde{\mathbf{M}}(j)) > -1$ for all i , we have

$$\sum_{j=0}^{t-1} \eta \lambda_i(\widetilde{\mathbf{M}}(j)) > 1 - a_1$$

By Jensen inequality, we have

$$\sum_{j=0}^{t-1} \log(1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))/\chi) \leq t \log \left(1 - \frac{\sum_{j=0}^{t-1} \eta \lambda_i(\widetilde{\mathbf{M}}(j))}{t\chi} \right)$$

Hence we have the following inequalities:

$$\begin{aligned}
\prod_{j=0}^{t-1} (1 - \eta \lambda_{i+1}(\widetilde{\mathbf{M}}(j))) &\leq \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))/\chi) \\
&\leq \exp \left\{ t \log \left(1 - \frac{\sum_{j=0}^{t-1} \eta \lambda_i(\widetilde{\mathbf{M}}(j))}{t\chi} \right) \right\} \\
&\leq \left(1 + \frac{a_1 - 1}{t\chi} \right)^t \\
&\leq e^{(a_1 - 1)/\chi}
\end{aligned}$$

So we complete the proof. \square

Then we pay attention back to sharpness. We have the sharpness's dynamics by Lemma C.2.

$$\begin{aligned}
\Lambda^*(t+1) - \Lambda^*(t) &= \mathbf{v}_1^\top (\mathbf{M}(t+1) - \mathbf{M}(t)) \mathbf{v}_1 \quad (\text{Definition (13)}) \\
&= -\frac{8\eta\lambda_1}{n^2m} (\mathbf{D}(t)^\top \mathbf{F}(t) + (\mathbf{D}(t)^\top \mathbf{v}_1)(\mathbf{F}(t)^\top \mathbf{v}_1)) \\
&\quad + \frac{8\eta^2\lambda_1}{m^2n^3} (\mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}^\top(t) \mathbf{W}(t) \mathbf{X} \mathbf{D}(t) + \|\mathbf{A}(t)\|^2 \lambda_1 (\mathbf{D}(t)^\top \mathbf{v}_1)^2)
\end{aligned}$$

This equation shows that the dynamics of sharpness is closely related to the dynamics of $\|\mathbf{A}(t)\|^2$, i.e. highly dependent on $-\mathbf{D}(t)^\top \mathbf{F}(t)$. Based on the lemmas above, we can prove progressive sharpening happens almost along the whole training trajectory (Lemma C.6) until the loss $\mathcal{L} = \frac{1}{n} \|\mathbf{D}(t)\|_2^2$ converges to $O(n^{-1})$.

Lemma C.6. Under Assumption C.1 and C.2, if $m > \frac{112c_1n^2}{\lambda_r^2}$ and

$$n > \max\{(70\lambda_r + 25\lambda_r^2)/(98\kappa^2 \min\{e^{-1/\chi} - e^{-2/\chi}, \eta c_1(1 - \eta c_1)\}), (\frac{70\lambda_r + 25\lambda_r^2}{98\kappa^2\epsilon})^2\}$$

we have $\Lambda^*(t+1) - \Lambda^*(t) > 0$ until the time t when $\|\mathbf{D}(t)\|^2 \leq \epsilon^2 + \frac{5\epsilon\lambda_r}{7\sqrt{n}} + \frac{25\lambda_r^2}{196n}$.

Remark. When $a_1 = \epsilon/\sqrt{n}$, $a_2 = e^{(a_1-1)/\chi} < e^{-1/\chi}$, the lower bound of n guarantee that

$$\kappa^2 n \min\{a_1 - a_1^2, a_2 - a_2^2, (1 - \eta c_1)\eta c_1\} > \frac{5\lambda_r}{7} - \frac{25\lambda_r^2}{98} \quad (16)$$

Proof. Note that the second order term

$$\frac{8\eta^2\lambda_1}{m^2n^3} (\mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}^\top(t) \mathbf{W}(t) \mathbf{X} \mathbf{D}(t) + \|\mathbf{A}(t)\|^2 \lambda_1 (\mathbf{D}(t)^\top \mathbf{v}_1)^2)$$

is larger than 0. So as long as the first order term

$$-\frac{8\eta\lambda_1}{n^2m} (\mathbf{D}(t)^\top \mathbf{F}(t) + (\mathbf{D}(t)^\top \mathbf{v}_1)(\mathbf{F}(t)^\top \mathbf{v}_1)) > 0$$

the approximate sharpness will grow.

First, we give a lower bound for the number $-(\mathbf{D}(t)^\top \mathbf{v}_1)(\mathbf{F}(t)^\top \mathbf{v}_1)$:

$$\begin{aligned}
-(\mathbf{D}(t)^\top \mathbf{v}_1)(\mathbf{F}(t)^\top \mathbf{v}_1) &= z_1^2 \prod_{j=0}^{t-1} (1 - \eta \lambda_1(\widetilde{\mathbf{M}}(j))) (1 - \prod_{j=0}^{t-1} (1 - \eta \lambda_1(\widetilde{\mathbf{M}}(j)))) \\
&\quad + z_1 (2 \prod_{j=0}^{t-1} (1 - \eta \lambda_1(\widetilde{\mathbf{M}}(j))) - 1) (\mathbf{e}(t)^\top \mathbf{v}_1) - (\mathbf{e}(t)^\top \mathbf{v}_1)^2 \\
&\geq -|z_1| \|\mathbf{e}(t)\| - \|\mathbf{e}(t)\|^2 \quad (\lambda_1(\mathbf{M}^*(t)) < 1/\eta)
\end{aligned}$$

$$\begin{aligned}
&\geq -\frac{40n^2}{\lambda_r m} - \frac{1600n^3}{\lambda_r^2 m^2} && \text{(Use (E1))} \\
&\geq -\frac{5\lambda_r}{14} - \frac{25\lambda_r^2}{196} && (m \geq \frac{112c_1 n^2}{\lambda_r^2})
\end{aligned}$$

where the first equation holds due to Property 2 of Lemma C.3.

With the lower bound of $-(\mathbf{D}(t)^\top \mathbf{v}_1)(\mathbf{F}(t)^\top \mathbf{v}_1)$, we show the dynamics of the first order term by similar technique in the expression (C.4) :

$$\begin{aligned}
&-\mathbf{D}(t)^\top \mathbf{F}(t) - (\mathbf{D}(t)^\top \mathbf{v}_1)(\mathbf{F}(t)^\top \mathbf{v}_1) && \text{(Property 2 of Lemma C.3)} \\
&= \mathbf{Y}^\top \prod_{j=0}^t (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) (\mathbf{I}_n - \prod_{j=0}^t (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j))) \mathbf{Y} + \mathbf{Y}^\top (2 \prod_{j=0}^t (\eta \widetilde{\mathbf{M}}(j) - \mathbf{I}_n) - \mathbf{I}_n) \mathbf{e}(t) \\
&\quad - \|\mathbf{e}(t)\|^2 - (\mathbf{D}(t)^\top \mathbf{v}_1)(\mathbf{F}(t)^\top \mathbf{v}_1) && \text{(Use (E1) and } m \geq \frac{112c_1 n^2}{\lambda_r^2} \text{)} \\
&\geq \kappa^2 n \max_i \left\{ \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) (1 - \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j)))) \right\} - \frac{5\lambda_r}{14} - \frac{25\lambda_r^2}{196} \\
&\quad + \left(-\frac{5\lambda_r}{14} - \frac{25\lambda_r^2}{196} \right) && \text{(Property 2 of Lemma C.3)} \\
&\geq \kappa^2 n \max_i \left\{ \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) (1 - \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j)))) \right\} - \frac{5\lambda_r}{7} - \frac{25\lambda_r^2}{98} \\
&\geq \kappa^2 n (1 - \eta c_1) \eta c_1 - \frac{5\lambda_r}{7} - \frac{25\lambda_r^2}{98} > 0
\end{aligned}$$

The last inequality holds due to the lower bound of n (16) assumed in Lemma C.6.

Now, $\prod_{j=0}^{t-1} (1 - \eta \lambda_1(\widetilde{\mathbf{M}}(j)))$ begins to decrease each iteration. Before the time when $\prod_{j=0}^{t-1} (1 - \eta \lambda_1(\widetilde{\mathbf{M}}(j)))$ becomes smaller than $\frac{\epsilon}{\sqrt{n}}$, $-\mathbf{D}(t)^\top \mathbf{F}(t) > 0$ always holds because of the lower bound of n (16) assumed in Lemma C.6.

Then, after the time t_1 when $\prod_{j=0}^{t-1} (1 - \eta \lambda_1(\widetilde{\mathbf{M}}(j))) < \frac{\epsilon}{\sqrt{n}}$ and before the time t_2 when $\prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) < \frac{\epsilon}{\sqrt{n}}$ for all i , we enter the time interval $[t_1, t_2)$ where Lemma C.5 begins to hold. We use Lemma C.5 to show that there exists some i to make

$$\prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) \in \left[\frac{\epsilon}{\sqrt{n}}, e^{-1/\chi} \right]$$

Before the time t_2 , we have this inequality always hold. Thus

$$-\mathbf{D}(t)^\top \mathbf{F}(t) - (\mathbf{D}(t)^\top \mathbf{v}_1)(\mathbf{F}(t)^\top \mathbf{v}_1) > 0$$

holds until the iteration t_2 . Thus during this period, $\Lambda^*(t)$ keeps increasing.

At this iteration, $\left\| \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \right\| \leq \frac{\epsilon}{\sqrt{n}}$, and we can bound the norm of the residual $\mathbf{D}(t)$ with the inequality below. We have

$$\begin{aligned}
\|\mathbf{D}(t)\|^2 &= \left\| -\prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{Y} + \mathbf{e}(t) \right\|^2 && \text{(Triangle Inequality)} \\
&\leq \left\| \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{Y} \right\|^2 + 2 \left\| \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \mathbf{Y} \right\| \|\mathbf{e}(t)\| + \|\mathbf{e}(t)\|^2 \\
&&& \text{(Cauchy-Schwarz)} \\
&\leq \left\| \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \right\|^2 \|\mathbf{Y}\|^2 + 2 \left\| \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \right\| \|\mathbf{Y}\| \|\mathbf{e}(t)\| + \|\mathbf{e}(t)\|^2 \\
&&& (\|\mathbf{Y}\| = \sqrt{n})
\end{aligned}$$

$$\begin{aligned}
&\leq n \max_i \left\{ \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) \right\}^2 + \frac{80n^2}{\lambda_r m} \max_i \left\{ \prod_{j=0}^{t-1} (1 - \eta \lambda_i(\widetilde{\mathbf{M}}(j))) \right\} + \frac{1600n^3}{\lambda_r^2 m^2} \\
&< \epsilon^2 + \frac{5\epsilon\lambda_r}{7\sqrt{n}} + \frac{25\lambda_r^2}{196n} \quad \left(\left\| \prod_{j=0}^{t-1} (\mathbf{I}_n - \eta \widetilde{\mathbf{M}}(j)) \right\| \leq \frac{\epsilon}{\sqrt{n}} \right)
\end{aligned}$$

Thus, before the norm of the residual $\mathbf{D}(t)$ decreases to this value, $\Lambda^*(t)$ keeps increasing. \square

C.3 Edge of Stability (Phase II-IV)

In the edge of stability regime, we focus on the largest eigenvalue Λ and its corresponding eigenvector \mathbf{u} . Since $\mathbf{M}(t)$ has a large similarity with $\mathbf{X}^\top \mathbf{X}$ in progressive sharpening phase, we consider the eigenvector \mathbf{v}_1 corresponding to the largest eigenvalue λ_1 of $\mathbf{X}^\top \mathbf{X}$.

After $\Lambda \geq 2/\eta$, the proof in Section C.2 does not extend to this phase. However, the bound of $\|\mathbf{\Gamma}(t)\|$ (Lemma C.3) still holds up to a constant factor empirically (See Figure 20). Hence, we make this bound an assumption as follows.

Assumption C.3. *There exists some constant $c_2 > 0$, such that for any time t ,*

$$\|\mathbf{\Gamma}(t)\| \leq \frac{c_2}{m}$$

Note that in above progressive sharpening stage, the assumption holds by Theorem 3. We propose this assumption to keep the gram matrix from deviating too far from the original trajectory even in other phases. The verification of this assumption can be found in Appendix D.2.2.

Corollary C.2. *Recall that \mathbf{v}_1 is the largest eigenvector of $\mathbf{X}^\top \mathbf{X}$ and \mathbf{u} is the largest eigenvector of \mathbf{M} . There exists a constant c_4 , such that $\langle \mathbf{v}_1, \mathbf{u} \rangle^2 \geq 1 - c_4 \sqrt{\frac{1}{nm}}$, and $\Lambda(t) \geq \Lambda^*(t) \geq \Lambda(t) - \frac{2c_2}{m}$.*

Proof. By definition of $\mathbf{\Gamma}(t)$ in 10, $\mathbf{M} = \gamma(t)\mathbf{X}^\top \mathbf{X} + \mathbf{\Gamma}(t)$, here $\gamma(t) = \frac{2}{mn}(\|\mathbf{A}(t)\|^2 + \frac{m}{d}) \geq \frac{2}{nd}$.

Let $\langle \mathbf{v}_1, \mathbf{u} \rangle = \cos \theta$ and decompose \mathbf{u} by $\mathbf{u} = \mathbf{v} \cos \theta + \mathbf{v}_\perp \sin \theta$. Then we have $\Lambda = \mathbf{u}^\top \mathbf{M} \mathbf{u} \leq \lambda_1 \gamma(t) \cos^2 \theta + \frac{\lambda_1}{2} \sin^2 \theta \gamma(t) + \frac{c_2}{m}$.

Also we have $\Lambda \geq \mathbf{v}_1^\top \mathbf{M} \mathbf{v}_1 \geq \gamma(t) \lambda_1 - \frac{c_2}{m}$. So $\lambda_1 \gamma(t) \cos^2 \theta + \frac{\lambda_1}{2} \sin^2 \theta \gamma(t) + \frac{c_2}{m} \geq \gamma(t) \lambda_1 - \frac{c_2}{m}$, which induces $\sin^2 \theta \leq \frac{2c_2}{m\gamma(t)} \cdot \frac{2}{\lambda_1} \leq \frac{2c_2^2 d}{m\lambda_1}$. Because in our setting $\lambda_1 = \Theta(n)$, there exists a constant c_4 such that $\cos \theta \geq \sqrt{1 - \frac{c_4^2}{nm}} \geq 1 - c_4 \sqrt{\frac{1}{nm}}$.

The inequality $\Lambda(t) \geq \Lambda^*(t)$ is because $\mathbf{u}^\top \mathbf{M} \mathbf{u} \geq \mathbf{v}_1^\top \mathbf{M} \mathbf{v}_1$ by definition of \mathbf{u} . The other side can be proved as the following:

$$\begin{aligned}
\mathbf{v}_1^\top \mathbf{M} \mathbf{v}_1 &= \mathbf{v}_1^\top (\mathbf{M} - \mathbf{\Gamma}) \mathbf{v}_1 + \mathbf{v}_1^\top \mathbf{\Gamma} \mathbf{v}_1 \\
&\geq \mathbf{u}^\top (\mathbf{M} - \mathbf{\Gamma}) \mathbf{u} - c_2/m \\
&\geq \mathbf{u}^\top \mathbf{M} \mathbf{u} - 2c_2/m.
\end{aligned}$$

\square

To prove the main theorem (Theorem 4), we need two more assumptions.

Assumption C.4. *There exists some constant c_3 such that $|\mathbf{D}(t)^\top \mathbf{v}_1| > c_3 \sqrt{n}/m$ for some $t = t_0$ at the beginning of phase II.*

Assumption C.5. *There exists some constant $\beta > 0$, such that*

$$\frac{4}{\eta}(1 - \beta) \geq \Lambda$$

The above assumption is consistent with Assumption 3.4, in which we assume an upper bound of the sharpness. Lewkowycz et al. [19] showed that $4/\eta$ is a upper bound of the sharpness in two-layer linear network with one datapoint, otherwise the training process would diverge. Here we make it an assumption.

Theorem 4. *Suppose the smallest nonzero eigenvalue $\lambda_r = \lambda_r(\mathbf{X}^\top \mathbf{X}) > 0$, $\lambda_1 = \lambda_{\max}(\mathbf{X}^\top \mathbf{X})$. Under Assumption C.3, C.5, C.4, and $\lambda_1(\mathbf{X}^\top \mathbf{X}) \geq 2\lambda_2(\mathbf{X}^\top \mathbf{X})$ in Assumption C.1, there exists constants C_1, C_2, C_3 , such that if $n > c_1 \lambda_r \eta$, $m > \max\{\frac{C_2 d^2 n^2}{\lambda_r^2}, C_3 \eta\}$, we have*

- (Lemma C.13) *There exists $\rho = O(1)$ which is related to c_3 such that if $\Lambda(t_0) > \frac{2}{\eta}(1 + \rho)$ for some t_0 , there must exist some $t_1 > t_0$ such that $\Lambda(t_1) < \frac{2}{\eta}(1 + \rho)$.*
- (Lemma C.11) *If $\Lambda(t), \Lambda(t+1) > \frac{2}{\eta}(1 + \rho)$, then there is a constant $c_7 > 0$ (related to c_3) such that $|\mathbf{D}(t+1)^\top \mathbf{v}_1| > |\mathbf{D}(t)^\top \mathbf{v}_1|(1 + c_7)$.*
- (Lemma C.10) *Define $\mathbf{R}(t) := (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t)$, and $\mathbf{R}'(t) := (\mathbf{I} - \eta \mathbf{M}^*(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top)) \mathbf{R}'(t-1)$. We have $\|\mathbf{R}(t) - \mathbf{R}'(t)\| = O(\frac{\sqrt{n^3 d}}{\lambda_r \sqrt{m}})$.*

Next, we prove this theorem in three parts: first we prove the third statement, which gives $\mathbf{R}(t)$ an upper bound (Lemma C.10), then we prove the second statement, in which we use Assumption C.4 to prove that when sharpness is above $2/\eta$, $\mathbf{D}^\top \mathbf{v}_1$ increases geometrically (Lemma C.11), and lastly we prove that the sharpness eventually drops below $2/\eta$ (Lemma C.13), which is the first statement in our theorem.

Now we first give a key equation:

Lemma C.7. *The dynamics of the approximation on sharpness is:*

$$\begin{aligned} & \Lambda^*(t+1) - \Lambda^*(t) \\ &= -\frac{8\eta\lambda_1}{mn^2} \left(\mathbf{F}(t)^\top \mathbf{D}(t) + (\mathbf{F}(t)^\top \mathbf{v}_1)(\mathbf{D}(t)^\top \mathbf{v}_1) - \frac{\eta}{2}(\mathbf{D}(t)^\top \mathbf{v}_1)^2 \Lambda^*(t) \right. \\ & \quad \left. - \frac{\eta}{2} \mathbf{R}(t)^\top \Gamma(t) \mathbf{R}(t) - \eta \mathbf{R}(t)^\top \Gamma(t) (\mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t)) - \frac{\eta}{mn} \mathbf{R}(t)^\top \left(\frac{m}{d} \mathbf{X}^\top \mathbf{X} \right) \mathbf{R}(t) \right). \end{aligned} \quad (17)$$

The proof of the equation is long and tedious, so we leave that in Lemma E.11.

Next we deal with the gap between $\mathbf{M}(t)$ and $\mathbf{M}^*(t)$. Note that \mathbf{M} is the Gram Matrix, but in gradient descent trajectory, $\mathbf{D}(t+1) - \mathbf{D}(t) = -\mathbf{M}^*(t) \mathbf{D}(t)$, so \mathbf{M}^* is the one that truly controls \mathbf{D} 's dynamics. From the following lemma, we can see that $\mathbf{M}^*(t)$ is a smoothed version of $\mathbf{M}(t)$ and $\mathbf{M}(t+1)$ plus a small perturbation.

Lemma C.8. *If $m \geq \eta c_2(d+1)$ (recall that c_2 is defined in Assumption C.3), then there exists $k_s \in [0, 1)$ and a constant c_6 such that $\|\mathbf{M}^*(t) - (1 - k_s)\mathbf{M}(t) - k_s \mathbf{M}(t+1)\| \leq \frac{c_6}{m}$.*

Proof. Recall that $\mathbf{M}^*(t) = \mathbf{M}(t) - \frac{4\eta}{N^2 m} (\mathbf{D}(t)^\top \mathbf{F}(t)) \mathbf{X}^\top \mathbf{X}$. Then we consider two different cases.

Case 1: $|\mathbf{D}(t)^\top \mathbf{F}(t)| \leq \frac{d+1}{d-2} (1 + \frac{d+1}{d-2}) n$.

In this case $\|\mathbf{M}^* - \mathbf{M}\|_2 \leq \frac{4\eta\lambda_1}{mn^2} \frac{d+1}{d-2} (1 + \frac{d+1}{d-2}) n = O(\frac{\eta\lambda_1}{mn}) = O(\frac{1}{m})$. Here the last inequality follows from the inequality (5).

Case 2: $|\mathbf{D}(t)^\top \mathbf{F}(t)| > \frac{d+1}{d-2} (1 + \frac{d+1}{d-2}) n$.

We claim that in this case we have $\mathbf{F}(t)^\top \mathbf{D}(t) > (\frac{2\eta}{mn} \frac{m}{d} \lambda_1 + \frac{c_2 \eta}{m}) \|\mathbf{D}(t)\|^2$. If it does not hold, then $R.H.S. \geq \mathbf{F}(t)^\top \mathbf{D}(t) \geq \|\mathbf{D}(t)\|^2 - \|\mathbf{D}(t)\| \|\mathbf{Y}\|$. Hence we can have

$$\begin{aligned} \|\mathbf{Y}\| &\geq \|\mathbf{D}(t)\| \left(1 - \frac{2\eta}{mn} \frac{m}{d} \lambda_1 - \frac{c_2 \eta}{m}\right) \\ &= \|\mathbf{D}(t)\| \left(1 - \frac{2\eta\lambda_1}{nd} - \frac{\eta c_2}{m}\right) \\ &\geq \|\mathbf{D}(t)\| \left(1 - \frac{2}{d+1} - \frac{1}{d+1}\right) \end{aligned}$$

The last inequality uses the restriction of m in this lemma and the inequality (5). Hence we have $|\mathbf{F}(t)^\top \mathbf{D}(t)| \leq \|\mathbf{D}(t)\|(\|\mathbf{D}(t)\| + \|\mathbf{Y}\|) \leq \frac{d+1}{d-2}(1 + \frac{d+1}{d-2})n$, which leads to a contradiction.

Hence now $\mathbf{F}(t)^\top \mathbf{D}(t) > (\frac{2\eta}{mn} \frac{m}{d} \lambda_1 + \frac{c_2\eta}{m})\|\mathbf{D}(t)\|^2$ holds. Since $\mathbf{X}^\top \mathbf{W}^\top \mathbf{W} \mathbf{X} = \frac{m}{d} \mathbf{X}^\top \mathbf{X} + \mathbf{\Gamma}(t)$ and $\|\mathbf{X}^\top \mathbf{X}\| = \lambda_1, \|\mathbf{\Gamma}(t)\| \leq \frac{c_2}{m}$, we can have:

$$\mathbf{F}(t)^\top \mathbf{D}(t) > \frac{2\eta}{mn} \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}^\top \mathbf{W} \mathbf{X} \mathbf{D}(t).$$

Now let $k_s := \frac{\mathbf{F}(t)^\top \mathbf{D}(t)}{2\mathbf{F}(t)^\top \mathbf{D}(t) - \frac{2\eta}{mn} \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}^\top \mathbf{W} \mathbf{X} \mathbf{D}(t)}$, by the inequality above we have $k_s < 1$.

By the equation

$$\begin{aligned} \mathbf{M}(t+1) - \mathbf{M}(t) &= \frac{2}{mn} (\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2) \mathbf{X}^\top \mathbf{X} + (\mathbf{\Gamma}(t+1) - \mathbf{\Gamma}(t)) \\ \|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2 &= -\frac{4\eta}{n} \mathbf{F}(t)^\top \mathbf{D}(t) + \frac{4\eta^2}{n^2} \cdot \frac{1}{m} \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}^\top \mathbf{W} \mathbf{X} \mathbf{D}(t) \end{aligned}$$

we have $\mathbf{M}^* = \mathbf{M}(t) + k_s(\mathbf{M}(t+1) - \mathbf{M}(t) - (\mathbf{\Gamma}(t+1) - \mathbf{\Gamma}(t)))$. Hence $\|\mathbf{M}^* - (1 - k_s)\mathbf{M}(t) - k_s\mathbf{M}(t+1)\| \leq \frac{2c_2}{m}$. Now combining the conclusions in both cases together, we finish the proof. \square

Then we consider a corollary of this lemma. Basically, since $\mathbf{M}^*(t)$ is a weighted sum of $\mathbf{M}(t)$ and $\mathbf{M}(t+1)$ adding a small perturbation and $\mathbf{M}(t)$ has a decomposition to the $\mathbf{X}^\top \mathbf{X}$ component and a small noise $\mathbf{\Gamma}(t)$, $\mathbf{M}^*(t)$ can also be decomposed into the $\mathbf{X}^\top \mathbf{X}$ component and a small noise.

Corollary C.3. \mathbf{M}^* can be decomposed to $\mathbf{M}^*(t) = \gamma^*(t) \mathbf{X}^\top \mathbf{X} + \mathbf{\Gamma}^*(t)$, where $\|\mathbf{\Gamma}^*(t)\| < \frac{c_2 + c_6}{m}$ and for any eigenvector \mathbf{u} of $\mathbf{X}^\top \mathbf{X}$ except \mathbf{v}_1 , and if let $\tau = \frac{2\lambda_i}{nd}$,

$$\tau \leq \mathbf{u}^\top (\mathbf{M}^*(t) - \mathbf{\Gamma}^*(t)) \mathbf{u} \leq \frac{2}{\eta} - \tau$$

Proof. By Lemma C.8, if we denote $\mathbf{\Gamma}'(t) = \mathbf{M}^*(t) - k_s\mathbf{M}(t+1) - (1 - k_s)\mathbf{M}(t)$, then $\|\mathbf{\Gamma}'(t)\| \leq \frac{c_6}{m}$.

Hence, we can see

$$\begin{aligned} \mathbf{M}^*(t) &= k_s\mathbf{M}(t+1) + (1 - k_s)\mathbf{M}(t) + \mathbf{\Gamma}' \\ &= \frac{2}{mn} (k_s\|\mathbf{A}(t+1)\|^2 + (1 - k_s)\|\mathbf{A}(t)\|^2 + \frac{2m}{d}) \mathbf{X}^\top \mathbf{X} \\ &\quad + k_s\mathbf{\Gamma}(t+1) + (1 - k_s)\mathbf{\Gamma}(t) + \mathbf{\Gamma}'(t) \end{aligned}$$

Denote $k_s\mathbf{\Gamma}(t+1) + (1 - k_s)\mathbf{\Gamma}(t) + \mathbf{\Gamma}'(t)$ by $\mathbf{\Gamma}^*(t)$. By Assumption C.3 and Lemma C.8, $\|\mathbf{\Gamma}^*(t)\| < \frac{c_2 + c_6}{m}$.

Note that by Assumption C.1, $\lambda_i(\mathbf{X}^\top \mathbf{X}) \leq \frac{1}{2}\lambda_1$ for any $i > 1$. Then, we can see that

$$\mathbf{v}_1^\top (\mathbf{M}(t) - \mathbf{\Gamma}(t)) \mathbf{v}_1 = \frac{2\lambda_1}{mn} (\|\mathbf{A}(t)\|^2 + \frac{m}{d}) \geq \frac{4\lambda_i}{mn} (\|\mathbf{A}(t)\|^2 + \frac{m}{d})$$

Hence, we can see

$$\begin{aligned} \frac{2\lambda_i}{mn} (\|\mathbf{A}(t)\|^2 + \frac{m}{d}) &\leq \frac{1}{2} \mathbf{v}_1^\top (\mathbf{M}(t) - \mathbf{\Gamma}(t)) \mathbf{v}_1 \\ &\leq \frac{1}{2} (\Lambda + \frac{c_2}{m}) \\ &\leq \frac{2}{\eta} (1 - \beta) + \frac{c_2}{2m}. \end{aligned}$$

Here the last inequality holds due to Assumption C.5.

Now the inequality above shows that for any eigenvector \mathbf{u} of $\mathbf{X}^\top \mathbf{X}$ except \mathbf{v}_1 ,

$$\mathbf{u}^\top (\mathbf{M}^*(t) - \mathbf{\Gamma}^*(t)) \mathbf{u} \leq \frac{2}{\eta} - \left(\frac{2\beta}{\eta} - \frac{c_2}{2m} \right).$$

On the other side, $\frac{2\lambda_i}{mn}(\|\mathbf{A}(t)\|^2 + \frac{m}{d}) \geq \frac{2\lambda_r}{nd}$.

Take $\tau = \min\{\frac{2\lambda_r}{nd}, \frac{2\beta}{\eta} - \frac{c_2}{2m}\}$. Now because $m = \Omega(\eta)$ and β is some constant, we can see $\tau = \min\{\frac{2\lambda_r}{nd}, \frac{2\beta}{\eta} - \frac{c_2}{2m}\} = \frac{2\lambda_r}{nd}$. \square

Before using this corollary to derive the dynamics of $\mathbf{R}(t)$ (and thus gives an upper bound of $\|\mathbf{R}(t)\|$), we need an upper bound for $\|\mathbf{D}(t)\|$.

Lemma C.9. *For any constant $c_5 < \min\{2\beta, \frac{d}{d+1}\}$, if $m > \frac{3\eta c_2}{\frac{2d}{d+1} - 2c_5}$ and $m > \frac{\eta c_2}{2\beta - c_5}$, there exists a constant c_3 such that $\|\mathbf{D}(t)\| \leq c_3\sqrt{nm}$.*

Proof. First we analyze the right hand side of equation (17). We can get

$$\begin{aligned}
& \text{R.H.S.} \\
&= -\frac{8\eta\lambda_1}{mn^2} \left(\mathbf{F}(t)^\top \mathbf{D}(t) + (\mathbf{F}(t)^\top \mathbf{v}_1)(\mathbf{D}(t)^\top \mathbf{v}_1) - \frac{\eta}{2}(\mathbf{D}(t)^\top \mathbf{v}_1)^2 \lambda_1 \right. \\
&\quad \left. - \frac{\eta}{2} \mathbf{R}(t)^\top \mathbf{\Gamma}(t) \mathbf{R}(t) - \eta \mathbf{R}(t)^\top \mathbf{\Gamma}(t) (\mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t)) - \frac{\eta}{mn} \mathbf{R}(t)^\top \left(\frac{m}{d} \mathbf{X}^\top \mathbf{X} \right) \mathbf{R}(t) \right) \\
&\leq -\frac{8\eta\lambda_1}{mn^2} \left(\|\mathbf{D}(t)\|^2 + (\mathbf{D}(t)^\top \mathbf{v}_1)^2 - \frac{\eta}{2}(\mathbf{D}(t)^\top \mathbf{v}_1)^2 \lambda_1 - \frac{\eta}{2} \|\mathbf{R}(t)\|^2 \|\mathbf{\Gamma}(t)\|_2 \right. \\
&\quad \left. - \eta \|\mathbf{R}(t)\| \cdot \|\mathbf{D}(t)\| \cdot \|\mathbf{\Gamma}(t)\|_2 - \frac{\eta\lambda_1}{nd} \|\mathbf{R}(t)\|^2 + \mathbf{Y}^\top \mathbf{D}(t) + (\mathbf{Y}^\top \mathbf{v}_1) (\mathbf{D}(t)^\top \mathbf{v}_1) \right) \\
&\leq -\frac{8\eta\lambda_1}{mn^2} \left(2\beta (\mathbf{D}(t)^\top \mathbf{v}_1)^2 + \|\mathbf{R}(t)\|^2 - \frac{c_2\eta}{2m} \|\mathbf{R}(t)\|^2 \right. \\
&\quad \left. - \frac{c_2\eta}{m} \|\mathbf{R}(t)\| \|\mathbf{D}(t)\| - \frac{1}{d+1} \|\mathbf{R}(t)\|^2 - \|\mathbf{Y}\| \cdot \|\mathbf{D}(t)\| - \|\mathbf{Y}\| (\mathbf{D}(t)^\top \mathbf{v}_1) \right) \\
&\leq -\frac{8\eta\lambda_1}{mn^2} \left(2\beta (\mathbf{D}(t)^\top \mathbf{v}_1)^2 + \left(\frac{d}{d+1} - \frac{\eta c_2 n}{2m} \right) \|\mathbf{R}(t)\|^2 - \frac{c_2\eta}{m} \|\mathbf{D}(t)\|^2 - 2\|\mathbf{Y}\| \cdot \|\mathbf{D}(t)\| \right) \\
&\leq -\frac{8\eta\lambda_1}{mn^2} (c_5 \|\mathbf{D}(t)\|^2 - 2\|\mathbf{Y}\| \cdot \|\mathbf{D}(t)\|)
\end{aligned}$$

Here the third inequality follows from inequality (5) which is $\frac{\eta\lambda_1}{nd} < \frac{1}{d+1}$, and the last inequality holds because $2\beta - \frac{c_2\eta}{m} > c_5$ and $\frac{d}{d+1} - \frac{3\eta c_2}{2m} > c_5$ by the restriction of m in this lemma.

On the other hand, the left hand side of equation (17) is $\Lambda^*(t+1) - \Lambda^*(t) \geq 0 - \Lambda(t) > -\frac{4}{\eta}$. Here the first inequality follows because of $\Lambda^*(t) \geq 0$ and Corollary C.2, and the second inequality holds because of Assumption C.5.

Hence we have $\frac{4}{\eta} > \frac{8}{n^2} \frac{\eta\lambda_1}{m} (c_5 \|\mathbf{D}(t)\|^2 - 2\|\mathbf{Y}\| \|\mathbf{D}(t)\|)$. Now if $\|\mathbf{D}(t)\| > \frac{2\|\mathbf{Y}\|}{c_5} + n\sqrt{m} \frac{1}{\eta\sqrt{2\lambda_1}}$, the inequality cannot hold. Note that $\|\mathbf{Y}\| = \sqrt{n}$ and $\lambda_1 \in \Theta(n)$, and also from inequality (5), $\eta < \frac{\lambda_1}{n} = O(1)$, we finish the proof. \square

Now we can give a lemma on $\mathbf{R}(t) := (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t)$. The proof idea is similar to Lemma 3.5 in Section 3.

Lemma C.10. *Define $\mathbf{R}'(t) := (\mathbf{I} - \eta \mathbf{M}^*(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top)) \mathbf{R}'(t-1)$. We have $\|\mathbf{R}(t) - \mathbf{R}'(t)\| = O(\frac{\sqrt{n^3 d}}{\lambda_r \sqrt{m}})$.*

Proof. We consider the update rule for $\mathbf{R}(t)$, whose proof is in Lemma E.12:

$$\mathbf{R}(t+1) = (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{R}(t) + \eta (\mathbf{v}_1 \mathbf{v}_1^\top \mathbf{\Gamma}(t) - \mathbf{\Gamma}(t) \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t) \quad (18)$$

Hence if we denote $\mathbf{e}_1(t) = \mathbf{R}(t+1) - (\mathbf{I} - \eta \mathbf{M}^*(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top)) \mathbf{R}(t)$, then we have

$$\mathbf{e}_1(t) = -\eta \mathbf{M}^* \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{R}(t) + \eta (\mathbf{v}_1 \mathbf{v}_1^\top \Gamma(t) - \Gamma(t) \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t) = \eta (\mathbf{v}_1 \mathbf{v}_1^\top \Gamma(t) - \Gamma(t) \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t).$$

Using the upper bound of $\|\mathbf{D}(t)\|$ in Lemma C.9, and Assumption C.3 we have

$$\|\mathbf{e}_1(t)\| \leq 2\eta \|\Gamma(t)\| \|\mathbf{D}(t)\| \leq 2\eta \frac{c_2}{m} c_3 \sqrt{nm} = 2c_2 c_3 \eta \sqrt{\frac{n}{m}}.$$

Now we consider the sequence $\mathbf{R}'(t)$.

First by Corollary C.3, we can see that the eigenvalues of $(\mathbf{M}^*(t) - \Gamma^*(t))(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top)$ are all the eigenvalues of $\mathbf{M}^*(t) - \Gamma^*(t)$ except the largest one, hence are in $(\tau, 2/\eta - \tau)$. Hence because $m = \Omega(n^2)$, if we let $\frac{\lambda_r}{nd} < \frac{c_2 + c_6}{m}$, then by Corollary C.3 we have all eigenvalues of $\mathbf{M}^*(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top)$ are in $(\tau', 2/\eta - \tau')$ where $\tau' = \frac{\lambda_r}{nd}$. Hence $\|\mathbf{I} - \eta \mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t) \mathbf{v}_1(t)^\top)\| \leq 1 - \eta \tau'$.

By the calculations above we can get

$$\begin{aligned} \|\mathbf{R}(t+1) - \mathbf{R}'(t+1)\| &= \|(\mathbf{I} - \eta \mathbf{M}^*(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top))(\mathbf{R}(t) - \mathbf{R}'(t)) + \mathbf{e}_1(t)\| \\ &\leq \|(\mathbf{I} - \eta \mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top))(\mathbf{R}(t) - \mathbf{R}'(t))\| + \|\mathbf{e}_1(t)\| \\ &\leq \|\mathbf{I} - \eta \mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top)\| \|\mathbf{R}(t) - \mathbf{R}'(t)\| + \|\mathbf{e}_1(t)\| \\ &\leq \|\mathbf{R}(t) - \mathbf{R}'(t)\| (1 - \eta \tau') + 2c_2 c_3 \eta \sqrt{\frac{n}{m}}. \end{aligned}$$

Thus if we denote $\|\mathbf{R}(t) - \mathbf{R}'(t)\| - \frac{2c_2 c_3 \eta \sqrt{\frac{n}{m}}}{\eta \tau'}$ by $p(t)$, and replace $\|\mathbf{R}(t) - \mathbf{R}'(t)\|$ in the inequality above by $p(t) + \frac{2c_2 c_3 \eta \sqrt{\frac{n}{m}}}{\eta \tau'}$, we can get $|p(t+1)| \leq |p(t)|(1 - \eta \tau')$. Hence, we can have $|p(t)| < |p(0)| = \frac{2c_2 c_3 \eta \sqrt{\frac{n}{m}}}{\eta \tau'}$ for any time t . Therefore, we can obtain $\|\mathbf{R}(t) - \mathbf{R}'(t)\| < \frac{2c_2 c_3 \eta \sqrt{\frac{n}{m}}}{\eta \tau'} + |p(0)| = \frac{4c_2 c_3 \sqrt{\frac{n}{m}}}{\tau'} = O(d \sqrt{\frac{n^3}{m}} / \lambda_r)$.

□

Now based on Theorem C.10, we can give an upper bound of $\|\mathbf{R}(t)\|$. Because $\|\mathbf{R}'(t)\|$ is always decreasing and $\|\mathbf{R}'(0)\| = \sqrt{n}$, hence if $m = \Omega(n^2 d^2 / \lambda_r^2)$, we can have $O(d \sqrt{\frac{n^3}{m}} / \lambda_r) = O(\sqrt{n})$. Hence

$$\|\mathbf{R}(t)\| \leq \|\mathbf{R}'(t)\| + O(\sqrt{n}) = O(\sqrt{n}). \quad (19)$$

Let $\|\mathbf{R}(t)\| \leq c_7 \sqrt{n}$.

Next we can use Assumption C.4 to prove that $\mathbf{D}(t)^\top \mathbf{v}_1$ increases geometrically when $\Lambda > 2/\eta$, which then causes the drop of the sharpness.

Lemma C.11. *Let $\rho^* = \frac{(c_2 + c_6)c_7}{2c_3}$, $\rho = \rho^* + \frac{\eta}{2} \frac{5c_2 + c_6}{m}$, $\epsilon_1 = 2\eta(\rho^* - \frac{(c_2 + c_6)c_7}{2c_3})$. If $\Lambda(t), \Lambda(t+1) > \frac{2}{\eta}(1 + \rho)$, we have $|\mathbf{D}(t+1)^\top \mathbf{v}_1| > (1 + \epsilon_1) \mathbf{D}(t)^\top \mathbf{v}_1$.*

Proof.

$$\begin{aligned} \mathbf{D}(t+1)^\top \mathbf{v}_1 &= \mathbf{D}(t)^\top (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{v}_1 \\ &= \mathbf{D}(t)^\top (\mathbf{I} - \eta \gamma^*(t) \mathbf{X}^\top \mathbf{X}) \mathbf{v}_1 - \eta \mathbf{D}(t)^\top \Gamma^*(t) \mathbf{v}_1 \\ &= (1 - \eta \gamma^*(t) \lambda_1) \mathbf{D}(t)^\top \mathbf{v}_1 - \eta \mathbf{D}(t)^\top \Gamma^*(t) \mathbf{v}_1 \end{aligned}$$

First by Lemma C.8 and Corollary C.2:

$$\begin{aligned} \gamma^*(t) \lambda_1 &= \mathbf{v}_1^\top (k_s (\mathbf{M}(t+1) - \Gamma(t+1)) + (1 - k_s) (\mathbf{M}(t) - \Gamma(t))) \mathbf{v}_1 \\ &\geq k_s \Lambda^*(t+1) + (1 - k_s) \Lambda^*(t) - \frac{2c_2}{m} \\ &\geq \frac{2(1 + \rho)}{\eta} - \frac{4c_2}{m} \end{aligned}$$

Also we have $|\mathbf{D}(t)^\top \mathbf{\Gamma}^*(t) \mathbf{v}_1| \leq \frac{c_2 + c_6}{m} \|\mathbf{D}(t)\| \leq \frac{c_2 + c_6}{m} (|\mathbf{D}(t)^\top \mathbf{v}_1| + \|\mathbf{R}(t)\|)$.

Hence

$$\begin{aligned} |\mathbf{D}(t+1)^\top \mathbf{v}_1| &\geq |1 - \eta\gamma^*(t)\lambda_1| |\mathbf{D}(t)^\top \mathbf{v}_1| - \eta |\mathbf{D}(t)^\top \mathbf{\Gamma}^*(t) \mathbf{v}_1| \\ &\geq | -1 - 2\rho + \frac{4c_2\eta}{m} + \frac{(c_2 + c_6)\eta}{m} | \\ &\geq (1 + 2\eta\rho^*) |\mathbf{D}(t)^\top \mathbf{v}_1| - \eta \frac{c_2 + c_6}{m} c_7 \sqrt{n} \\ &\geq (1 + \epsilon_1) |\mathbf{D}(t)^\top \mathbf{v}_1| \end{aligned}$$

Here the third inequality holds by the definition of ρ^* . \square

Now we state a lemma which proves that if $\mathbf{D}(t)^\top \mathbf{v}_1$ is large enough, then the sharpness will decrease in the next iteration.

Lemma C.12. *Assume $m = \Omega(\eta)$. There exists constant $c_8, c_9 > 0$ such that if $\mathbf{D}(t)^\top \mathbf{v}_1 > c_8 \sqrt{n}$, then $\mathbf{v}_1^\top \mathbf{M}(t+1) \mathbf{v}_1 - \mathbf{v}_1^\top \mathbf{M}(t) \mathbf{v}_1 < -c_9/m$.*

Proof. Recall Equation (17). Also we have $\|\mathbf{R}(t)\| \leq c_7 \sqrt{n}$ by (19). Hence

$$\begin{aligned} &\Lambda^*(t+1) - \Lambda^*(t) \\ &= -\frac{8\eta\lambda_1}{mn^2} \left(\mathbf{F}(t)^\top \mathbf{D}(t) + (\mathbf{F}(t)^\top \mathbf{v}_1) (\mathbf{D}(t)^\top \mathbf{v}_1) - \frac{\eta}{2} (\mathbf{D}(t)^\top \mathbf{v}_1)^2 \Lambda^*(t) \right. \\ &\quad \left. - \frac{\eta}{2} \mathbf{R}(t)^\top \mathbf{\Gamma}(t) \mathbf{R}(t) - \eta \mathbf{R}(t)^\top \mathbf{\Gamma}(t) (\mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t)) - \mathbf{R}(t)^\top \left(\frac{m}{d} \mathbf{X}^\top \mathbf{X} \right) \mathbf{R}(t) \frac{\eta}{mn} \right) \\ &\leq -\frac{8\eta\lambda_1}{mn^2} \left(2\beta (\mathbf{D}(t)^\top \mathbf{v}_1)^2 + \|\mathbf{R}(t)\|^2 - \frac{\eta c_2}{2m} \|\mathbf{R}(t)\|^2 \right. \\ &\quad \left. - \frac{c_2\eta}{m} \|\mathbf{R}(t)\| \cdot \|\mathbf{D}(t)\| - \frac{\eta\lambda_1}{dn} \|\mathbf{R}(t)\|^2 - 2\|\mathbf{Y}\| \cdot \|\mathbf{D}(t)\| \right) \\ &\leq -\frac{8\eta\lambda_1}{mn^2} \left(2\beta (\mathbf{D}(t)^\top \mathbf{v}_1)^2 + \|\mathbf{R}(t)\|^2 \left(1 - \frac{c_2\eta}{2m}\right) \right. \\ &\quad \left. - \frac{c_2\eta}{m} \|\mathbf{R}(t)\| \cdot (\|\mathbf{R}(t)\| + |\mathbf{D}(t)^\top \mathbf{v}_1|) - \frac{\eta\lambda_1}{dn} \|\mathbf{R}(t)\|^2 - 2\|\mathbf{Y}\| \cdot (\|\mathbf{R}(t)\| + |\mathbf{D}(t)^\top \mathbf{v}_1|) \right) \\ &\leq -\frac{8\eta\lambda_1}{mn^2} \left(2\beta (\mathbf{D}(t)^\top \mathbf{v}_1)^2 - \left(\frac{\eta c_7 c_2 \sqrt{n}}{m} + 2\sqrt{n} \right) \cdot |\mathbf{D}(t)^\top \mathbf{v}_1| \right. \\ &\quad \left. + c_7^2 n \left(1 - \frac{c_2\eta}{2m} - \frac{c_2\eta}{m} - \frac{\eta\lambda_1}{dn}\right) - 2c_7 n \right) \\ &= -\frac{8\eta\lambda_1}{mn^2} \left(2\beta (\mathbf{D}(t)^\top \mathbf{v}_1)^2 - c_{10} (\mathbf{D}(t)^\top \mathbf{v}_1) + c_{11} \right) \end{aligned}$$

Here $c_{10} := \eta c_7 \sqrt{n} \cdot \frac{c_2}{m} + 2\sqrt{n} = \Theta(\sqrt{n})$, $c_{11} := c_7^2 n \left(1 - \frac{c_2\eta}{2m} - \frac{c_2\eta}{m} - \frac{\eta\lambda_1}{dn}\right) - 2c_7 n = Nn$, and the first inequality holds because of Assumption C.5, Assumption C.3 and $\|\mathbf{D}(t)\|^2 = (\mathbf{D}(t)^\top \mathbf{v}_1)^2 + \|\mathbf{R}(t)\|^2$, the second inequality holds because of $\|\mathbf{D}(t)\| \leq |\mathbf{D}(t)^\top \mathbf{v}_1| + \|\mathbf{R}(t)\|$, and the third inequality follows from the upper bound (19) of $\|\mathbf{R}(t)\|$.

Note that it is a quadratic function of $\mathbf{D}(t)^\top \mathbf{v}_1$. Hence if $\mathbf{D}(t)^\top \mathbf{v}_1 > \frac{|c_{10}|}{2\beta} + \sqrt{\frac{|c_{11}|}{2\beta}}$, we have

$$\Lambda^*(t+1) - \Lambda^*(t) < -\frac{8\eta\lambda_1}{n^2 m} \cdot \frac{|c_{10}| \sqrt{|c_{11}|}}{2\beta} = \Theta\left(\frac{\eta}{m}\right)$$

Hence we finish the proof. \square

Now we can prove the final part (the first conclusion) of Theorem 4.

Lemma C.13. *Let ρ be the one defined in Theorem C.11. If $\Lambda(t_0) > \frac{2}{\eta}(1 + \rho)$ for some t_0 , there exists some $t_1 > t_0$ such that $\Lambda(t_1) < \frac{2}{\eta}(1 + \rho)$.*

Proof. Otherwise, for any $t > t_0$, $\Lambda(t) > \frac{2}{\eta}(1 + \rho)$. Then by Theorem C.11, $\mathbf{D}(t)^\top \mathbf{v}_1$ increases geometrically, hence there exists $t_2 > t_0$, such that $\mathbf{D}(t)^\top \mathbf{v}_1 > c_8 \sqrt{n}$. Now by Lemma C.12, each iteration $\Lambda^*(t)$ will decrease by at least a fixed amount. Hence there must exist a time $t_3 > t_0$ such that $\Lambda^*(t_3) < \frac{2}{\eta} - 2c_2/m$. Then by Corollary C.2, we get a contradiction. \square

C.4 Our Results and the NTK Regime

In this subsection, we explain why our results (Theorem 3, Theorem 4) are sufficiently different from the quadratic setting (e.g., linear regression) or the recent convergence analysis in NTK setting.

A key requirement in the convergence analysis in the NTK regime is that the learning rate is very small and the GD trajectory almost tracks the gradient flow, hence converges to the global minimum. However, we consider typical learning rate used in practice, which can be much larger. In particular, $\eta > 2/\Lambda$ can happen in our setting, which causes instability (i.e., such as the growth of loss in Lemma C.11) along the training trajectory. Such instability cannot be captured by any existing convergence analysis in NTK regime at all. Hence, all existing NTK convergence results do not directly apply here.

Equally importantly, we find that even when $\mathbf{W}(t)$ changes slightly (several orders of magnitude smaller than its initialization), PS and EOS still happen with a not so small learning rate η . To support our claim, we include the experimental results in Appendix D.2.3. In Figure 22, we can see that the initialization $\mathbf{W}(0)$ is much larger than the change of $\mathbf{W}(t)$ and the norm of $\mathbf{W}(t)$ grows larger when m becomes larger. However, we still observe that PS and EOS occur in this setting. Hence, the setting we study in this paper and our results are intrinsically different from the quadratic setting (in which case EOS cannot happen).

Last, in our proofs in Section C.2 and C.3, our current bound requires that $m = \Omega(n^2)$ and we also assume $\lambda_r = \Omega(1)$. This may create an impression that we need a very wide network which operates in the NTK regime. However, we remark that if our analysis can be tightened to $m = O(n)$, one can formally prove that $\|\mathbf{A}(t)\|$ can actually change significantly ($\|\mathbf{A}(t)\|^2 - \|\mathbf{A}(0)\|^2$) has the same scale as the initialization $\|\mathbf{A}(0)\|^2$, resulting a significant change of sharpness as well, hence beyond the NTK regime. For example, in the proof of EOS (Section C.3), we prove a loose $O(\sqrt{mn})$ upper bound of $\|\mathbf{D}(t)\|$ (Lemma C.9). However by Lemma C.12, when $\|\mathbf{D}(t)\|$ reaches $O(\sqrt{n})$, the sharpness starts dropping quickly. So if a better upper bound of $\|\mathbf{D}(t)\| = O(\sqrt{n})$ can be proven (this is true empirically for all of our experiments), the width m can be set to $\Theta(n)$, and this suffices to implies a significant change of $\|\mathbf{A}(t)\|^2$. We leave these improvements as future directions.

D Verification for Assumptions

In this section, we first justify the assumptions we made in Section 3 and Section 4 empirically. Then we present the experiment described in Section 5.

D.1 Assumptions in Section 3

In this subsection, we conduct experiments to verify the assumptions in Section 3. The detailed experiment settings can be found in Appendix A.

D.1.1 \mathbf{M}_A has small eigenvalues

In Section 3.2, we mentioned that the largest eigenvalue of $\mathbf{M}_A := (\frac{\partial \mathbf{F}}{\partial \mathbf{A}})(\frac{\partial \mathbf{F}}{\partial \mathbf{A}})^\top$ is much smaller than the sharpness. We verify this assumption under different settings in Figure 13, including a fully-connected linear network, a fully-connected network with tanh activation and a convolutional one. Observe that $\|\mathbf{M}_A\|$ (the blue curve) is very close to 0 and hardly increases during the training process along the whole trajectory.

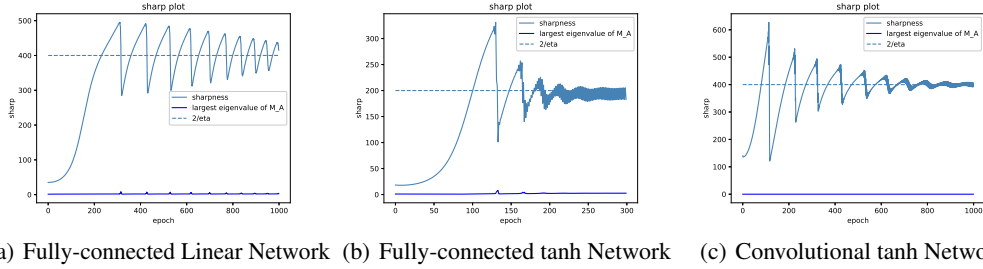


Figure 13: In this figure, we show that $\|M_A\|$ is much smaller than the sharpness. Note that $\|M_A\|$ (the blue curve) is very close to 0 and hardly increases during the training process along the whole trajectory. The sharpness (the steel-blue curve) strictly dominates $\|M_A\|$ for all time. The detailed experiment settings can be found in Appendix A.

D.1.2 Assumption 3.7

In this assumption, we assume that there is a large gap between the largest and the second largest eigenvalue, and thus the second largest eigenvalue is always below $1/\eta$. We verify the outlier assumption by calculating the largest and the second largest eigenvalue of $M(t)$. In Sagun et al. [28, 29], the sharpness is much larger than the largest eigenvalue in the bulk (the $(K + 1)$ -th largest eigenvalue of M where K is the number of classes). In our binary setting $K = 1$. In Figure 16, we show that the largest eigenvalue indeed dominates the second one, and the second one never reaches $1/\eta$, which verifies Assumption 3.7.




Figure 14: The second largest eigenvalue (the orange curve) of M is much smaller than the sharpness (the steel-blue curve). Also, the second largest eigenvalue never reaches $2/\eta$.

D.1.3 Assumption 3.2 (First Order Approximation of GD)

In Assumption 3.2, we assume the gradient descent trajectory is close to the first order approximation. To verify that the first order term is indeed dominant along the trajectory, for both the residual $D(t)$ and the output layer norm $\|A(t)\|^2$, we plot the norms of the actual GD update, the first order approximation order approximation and the higher order terms of the update rule in Figure 15. Observe that in the progressive sharpening phase, the first order approximation is almost the same as the actual gradient update; while in the EOS phase, the first order approximation is still close to the actual gradient most of the time. We can see that the norm of the higher order terms spikes occasionally, but when this happens the first order term spikes much higher.

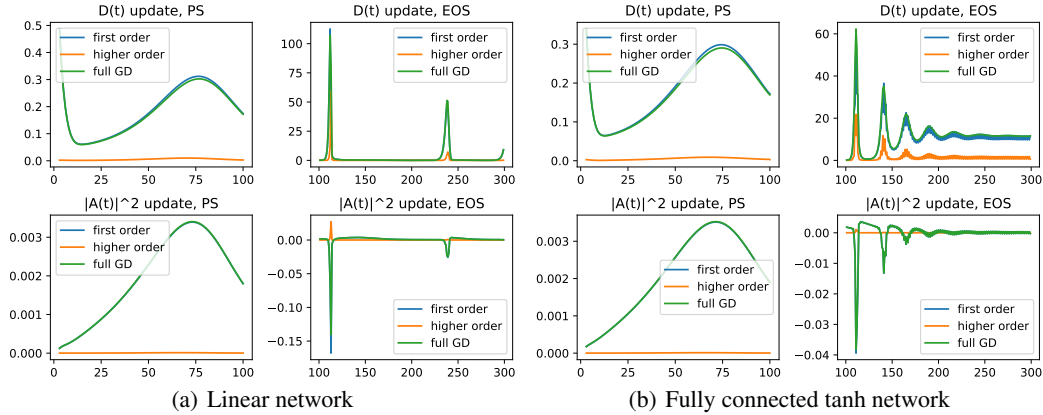


Figure 15: This figure shows the norm of the full gradient $\|D(t+1) - D(t)\|$, the first order approximation $\|\eta M(t)D(t)\|$ and the higher order terms $\|D(t+1) - D(t) + \eta M(t)D(t)\|$ in the training process. Observe that in both the $\|A(t)\|^2$ and $D(t)$ dynamics, the first order approximation (the blue curve) almost overlaps with the actual update (the green curve) in the PS phase. While in the EOS phase, though the norm of the higher order terms (the orange curve) spikes, the first order approximation is still close to the actual update.

D.1.4 Assumption 3.3 (Gradient flow for the PS phase)

In Assumption 3.3, we assume that in the progressive sharpening phase, the gradient descent trajectory is close to the gradient flow trajectory. We refer to Appendix J in Cohen et al. [6] for more experiments about this claim. Here for readers' convenience, we duplicate their Figure 29 in Appendix J as Figure 16.

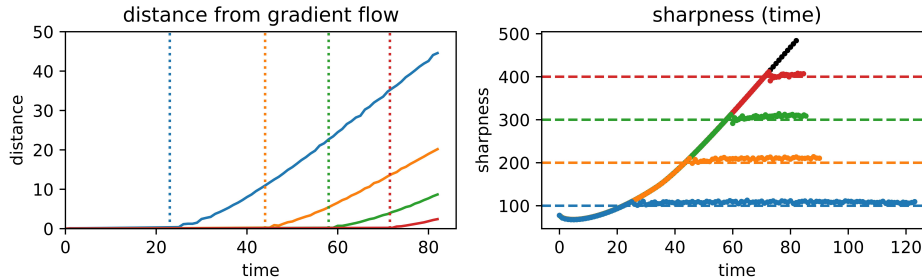


Figure 16: **Left:** the picture shows the ℓ_2 distance between the gradient flow trajectory at t iteration and the gradient descent trajectory at t/η iteration. The vertical dotted line is the time when the sharpness reaches $2/\eta$. **Right:** the plot of sharpness where (iteration \times learning rate) is on the x-axis.

D.1.5 Assumption 3.5 (iii) (principal directions moves slowly)

In Figure 17 we verify Assumption 3.5 (iii) and as well as the discussion after Assumption 3.5 (i). In general models, we find that the eigenvectors corresponding to small eigenvalues may change drastically. But for the largest eigenvector, it indeed changes slowly from the initialized direction. In Figure 17, we can see that over a long training time the similarity of $v_1(t)$ with its initialization is still larger than at least 0.98. Mulayoff et al. [24] also proved that near the minima, the top eigenvectors of the Hessian matrices tend to align. That is, the directions of these top eigenvectors are approximately parallel. This fact also corroborates our assumption near the minima.

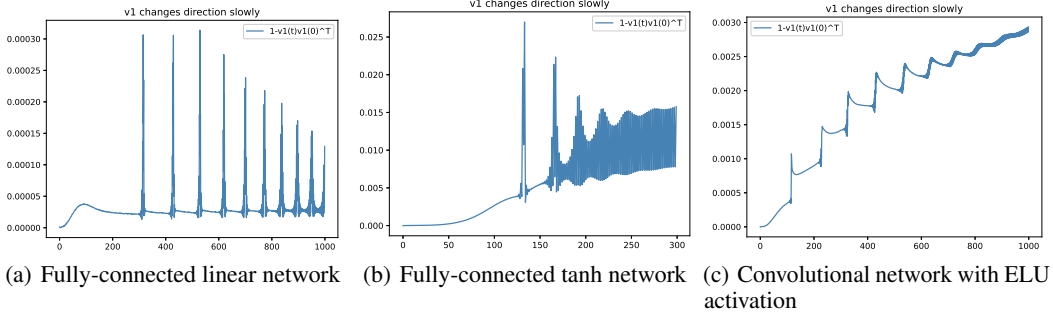


Figure 17: In this figure, we show that under various architectures, the first eigenvector of M changes slowly throughout the training process, even when Edge of Stability phenomenon occurs.

D.2 Assumptions in Section 4

In this subsection, we present the assumptions in Section 4, and conduct experiments to verify them. We add a scale coefficient $\frac{1}{\sqrt{m}}$ in the linear network to be consistent with the settings of theoretical analysis.

D.2.1 Assumption 4.1

In Assumption 4.1, we assume the ratio λ_i/λ_{i+1} between the two adjacent eigenvalues of $\mathbf{X}^\top \mathbf{X}$ is bounded by a small constant. In the 1000-example subset of CIFAR-10, we verify this assumption by experiments. We plot the eigenvalues and their ratio in Figure 18. It shows that Assumption 4.1 holds and the constant $\chi \approx 38$, since almost all the ratios are close to 1, and the largest ratio is $\lambda_1/\lambda_2 \approx 38$.

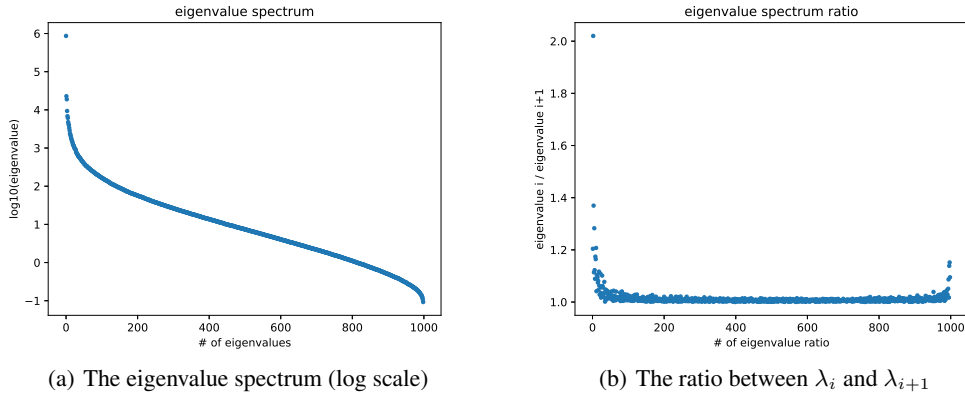


Figure 18: (a): The full eigenvalue spectrum of of $\mathbf{X}^\top \mathbf{X}$. Observe that the maximal ratio $\max\{\lambda_i/\lambda_{i+1}\} = \lambda_1/\lambda_2$. (b): The ratio between two adjacent eigenvalues λ_i/λ_{i+1} , $i \geq 2$. We exclude the largest eigenvalue to make the figure clearer.

If we further consider a mean-subtracted version of this subset of CIFAR-10 (See Figure 19), we can reduce the ratio to $\chi \approx 3$.

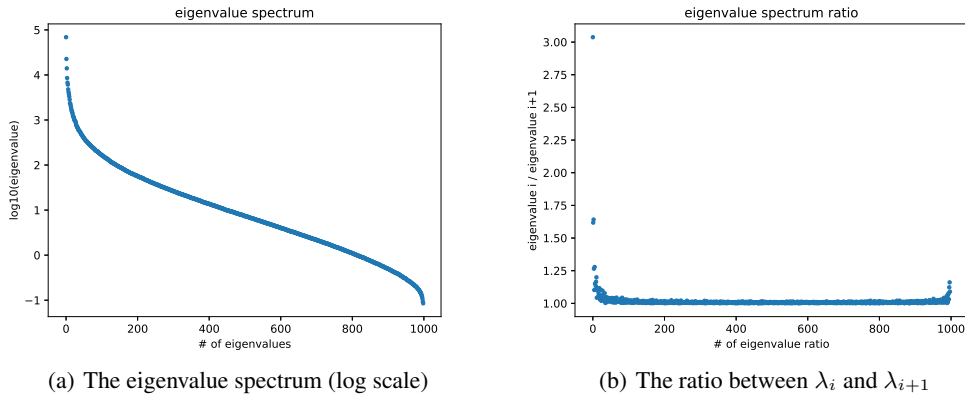


Figure 19: (a): The full eigenvalue spectrum of of $\mathbf{X}^\top \mathbf{X}$ when samples' mean is subtracted. (b): the ratio between two adjacent eigenvalues λ_i/λ_{i+1} . In Figure (b), observe that the maximal ratio $\max\{\lambda_i/\lambda_{i+1}\} = \lambda_1/\lambda_2 \approx 3$. **Remark.** Note we should exclude the minimal eigenvalue if the mean of examples is subtracted, because when we subtract the mean from the dataset, it cannot be full rank. The minimal eigenvalue of $\mathbf{X}^\top \mathbf{X}$ is 0, thus it is unnecessary to check it for our assumption.

D.2.2 Assumption 4.3

In Appendix C.2, by Theorem 3 we prove that $\|\Gamma(t)\|$ is bounded by $R_w = O(1/m)$ in the progressive sharpening phase. When gradient descent enters EOS, the proof does not hold and we make this bound into an assumption (Assumption 4.3). We empirically verify that the bound can only increase by a constant factor (despite some spikes in EOS). See Figure 20. In this figure $\|\Gamma(t)\| \leq \frac{24}{m}$ with $m = 40, 80, 160, 200$.

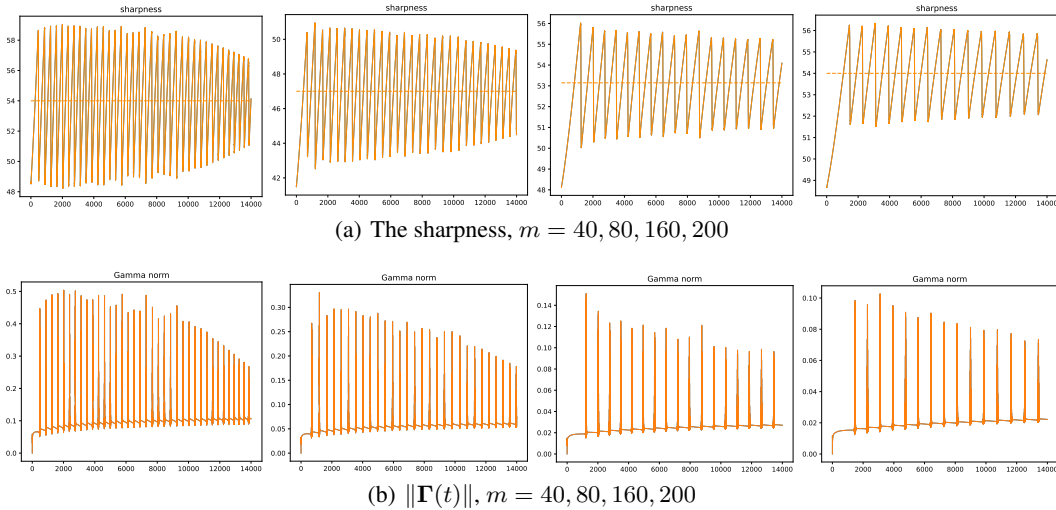


Figure 20: In this figure, we show that even when GD enters EOS, the bound of $\|\Gamma(t)\|$ is still $\Theta(1/m)$ along the whole trajectory. The corresponding constant $c_2 \approx 24$.

Actually there are two interesting empirical facts related to $\|\Gamma(t)\|$. One is the noticeable fact that $\|\Gamma(t)\|$ spikes, the other is that the values of $\|\Gamma(t)\|$ are overall quite small (despite the fact that \mathbf{W} changes non-trivially in our experiments) and decreases as m becomes larger. Our assumption tries to model the second fact (see Figure 20, $\|\Gamma(t)\|$) (despite the spikes) decreases as the width m grows, and the largest $\|\Gamma(t)\|$ is almost $24/m$ in all these experiments). However, we admit that our results do not reflect the first fact (the spikes of $\|\Gamma(t)\|$), and it is an interesting fact that is

worth investigating. We have strong intuition that $\|\Gamma(t)\|$ only grows by at most a constant factor, but currently do not have a formal proof yet. Nevertheless, the spiking behavior does not directly contradict our assumption.

Remark: Note that Assumption 4.3 is not equivalent to a small movement of $\mathbf{W}(t)$. Actually, in the experiments above ($m = 40, 80, 160, 200$) the movement of $\mathbf{W}(t)$ is quite significant compared to the norm of $\mathbf{W}(0)$ at initialization. See Figure 21.

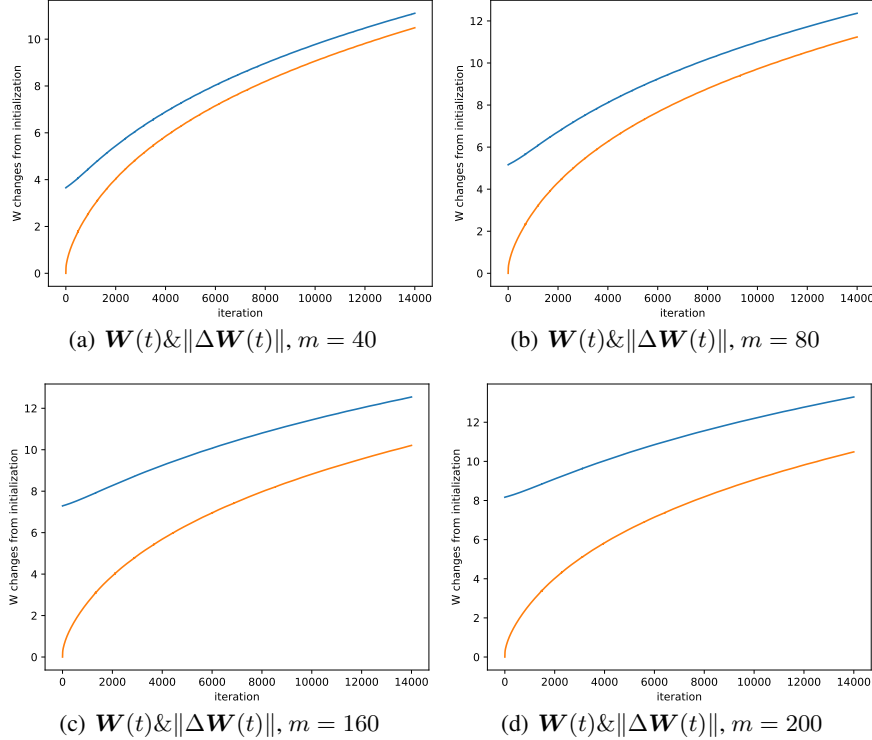


Figure 21: In this figure, we show that even the $\Gamma(t)$ is small and bounded, the movement $\|\Delta\mathbf{W}(t)\| := \|\mathbf{W}(t) - \mathbf{W}(0)\|$ (the dark orange curve) is close to $\|\mathbf{W}(t)\|$ (the blue curve), implying that $\mathbf{W}(t)$ moves considerably.

D.2.3 Comparison with the NTK regime: the non-quadratic property

Here we illustrate why our setting and results are sufficiently different from the quadratic setting (e.g., linear regression) or the recent convergence analysis in NTK setting. In particular, we show that even when the movement of $\mathbf{W}(t)$ is comparably negligible compared to the initialization $\mathbf{W}(0)$, EOS can still happen. Here we take a larger initialization of $\mathbf{W}(0)$, which is ten times of the standard initialization in order to dwarf the movement of $\mathbf{W}(t)$. The widths are $m = 1000, 2000, 4000, 8000$. We can see that the initialization $\mathbf{W}(0)$ is much larger than the change of $\mathbf{W}(t)$ and the norm of $\mathbf{W}(t)$ grows larger when m becomes larger. See Figure 22. Detailed comparison is in Appendix C.4.

D.3 Experiments in Section 5

In this subsection we show some empirical results that may reveal some interesting relation between the inner layers and the sharpness, which is not yet reflected in our theory. Our experimental results show that all layers seem to work together to influence the sharpness, and contribute to the progressive sharpening and edge of stability phenomena. In our experiment, while the sharpness is still calculated by the gradient of all parameters, we freeze some of the layers in the training process. The experimental results show that the more layers we freeze, the slower progressive sharpening happens and the weaker the oscillation of the sharpness is (See Figure 23, Figure 24). This indicates

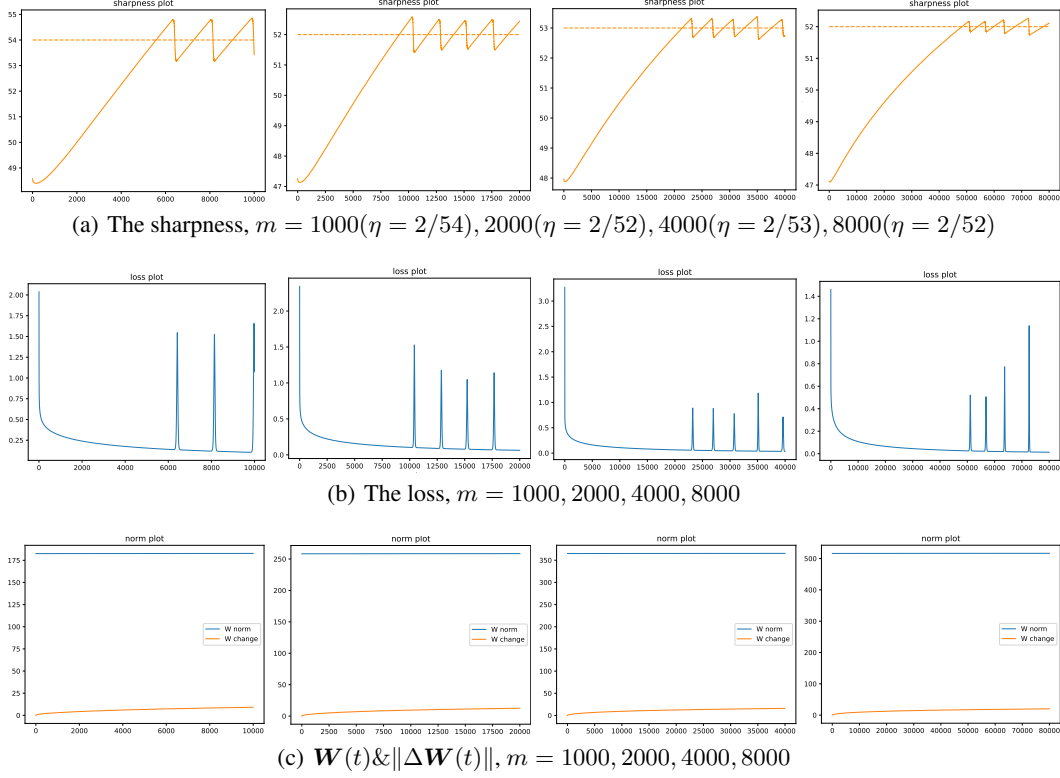


Figure 22: In this figure, we show that even when $\mathbf{W}(t)$ stays close to its initialization, EOS can still happen. In particular, Figure (c) shows the $\|\mathbf{W}(t)\|$ (the blue curve) and $\|\mathbf{W}(t) - \mathbf{W}(0)\|$ (the orange curve). Here for $m = 1000, 2000, 4000, 8000$, observe that the EOS still happens when $\|\mathbf{W}(t)\|/\|\mathbf{W}(t) - \mathbf{W}(0)\| > 10$. That indicates the intrinsic non-quadratic property of neural networks in the EOS regime.

that layers other than the output layer has nontrivial influence on the sharpness, but since different layers seem to work in the same direction, further justifying our assumption that $\|\mathbf{A}\|$ is positively related with the sharpness.

Fully-connected tanh Network:

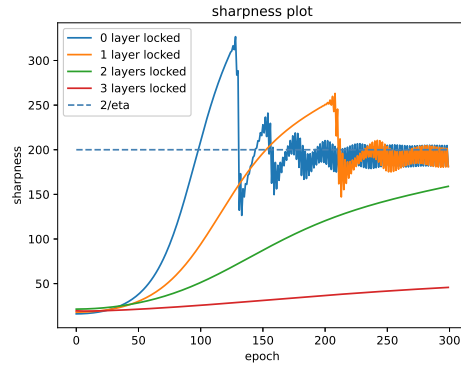
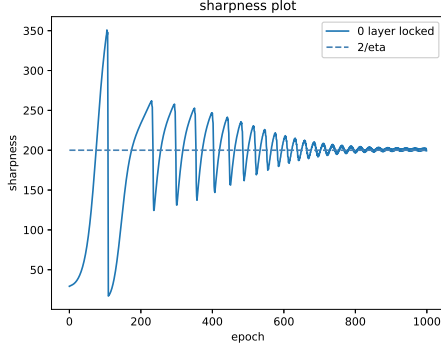



Figure 23: We conduct the experiment on a 5-layer tanh fully-connected network. Four lines in the plot show the sharpness for four independent training processes but with 0,1,2,3 outer layers locked. All other hyper-parameters are the same. The result shows that all layers have cooperative effect on the sharpness.


Fully-connected Linear Network:




(a) Fully-connected linear network 0 layer frozen



(b) Fully-connected linear network 1 layer frozen



(c) Fully-connected linear network 2 layers frozen



(d) Fully-connected linear network 3 layers frozen

Figure 24: The sharpness after 0,1,2,3 outer layers are locked. Details refer to Figure 23.

E Missing proofs

E.1 Proofs in Section 3

In this section, we provide the missing proofs in Section 3.

Lemma E.1. (Lemma 3.1) For all t in Phase I, under Assumption 3.5 (i) and 3.3, it holds that $\mathbf{D}(t)^\top \mathbf{F}(t) < 0$.

Proof. Consider the inner product $\mathbf{D}(t)^\top \mathbf{v}_i$. Recall \mathbf{v}_i is the i -th eigenvector of \mathbf{M} . By Assumption 3.5(i), \mathbf{v}_i does not change with t . By Assumption 3.3, we have

$$\frac{d\mathbf{D}(t)^\top \mathbf{v}_i}{dt} = -\mathbf{D}(t)^\top \mathbf{M}(t) \mathbf{v}_i = -\lambda_i(t) \mathbf{D}(t)^\top \mathbf{v}_i$$

where $\lambda_i(t) \geq 0$ is the corresponding eigenvalue of \mathbf{v}_i .

Since $\mathbf{F}(0) = 0$, $\mathbf{D}(0) = \mathbf{F}(0) - \mathbf{Y} = -\mathbf{Y}$. Then the differential equation has the solution

$$\mathbf{D}(t)^\top \mathbf{v}_i = e^{-\int_0^t \lambda_i(t) dt} \cdot \mathbf{D}(0)^\top \mathbf{v}_i = -e^{-\int_0^t \lambda_i(t) dt} \cdot \mathbf{Y}^\top \mathbf{v}_i$$

Plug it into the expression $\mathbf{D}(t)^\top \mathbf{F}(t) = \mathbf{D}(t)^\top (\mathbf{D}(t) + \mathbf{Y})$ and we have:

$$\mathbf{D}(t)^\top \mathbf{F}(t) = \sum_{i=1}^n \mathbf{D}(t)^\top \mathbf{v}_i (\mathbf{D}(t)^\top \mathbf{v}_i + \mathbf{Y}^\top \mathbf{v}_i) = -\sum_{i=1}^n (\mathbf{Y}^\top \mathbf{v}_i)^2 e^{-\int_0^t \lambda_i(t) dt} (1 - e^{-\int_0^t \lambda_i(t) dt})$$

since $(\mathbf{Y}^\top \mathbf{v}_i)^2 e^{-\int_0^t \lambda_i(t) dt} (1 - e^{-\int_0^t \lambda_i(t) dt}) > 0$ for all $i \in [n]$, $\mathbf{D}(t)^\top \mathbf{F}(t) < 0$. \square

Lemma E.2. (Lemma 3.2) Suppose Assumption 3.5 (iii) and 3.6 hold during this phase (with constants $\epsilon_2 > 0$ and $c > 1$). If $\Lambda(t) = (2 + \tau)/\eta$ and $\tau > \frac{1}{1-\epsilon_2-1/c} - 1$, then $\mathbf{D}(t)^\top \mathbf{v}_1(t)$ increases geometrically with factor $(1 + \tau)(1 - \epsilon_2 - 1/c) > 1$ for $t \geq t_0$ in this phase.

Proof. First by (3), we have $\mathbf{D}(t+1)^\top \mathbf{v}_1(t) = \mathbf{D}(t)^\top (\mathbf{I} - \eta \mathbf{M}(t)) \mathbf{v}_1(t) = (1 - \eta \Lambda(t)) \mathbf{D}(t)^\top \mathbf{v}_1(t)$. Then we have

$$\begin{aligned} |\mathbf{D}(t+1)^\top \mathbf{v}_1(t+1)| &\geq |\mathbf{D}(t+1)^\top \mathbf{v}_1(t)(1 - \epsilon_2)| - \epsilon_2 \|\mathbf{D}(t+1)\| \\ &\geq |\mathbf{D}(t)^\top \mathbf{v}_1(t)|(1 - \epsilon_2)(\eta \Lambda(t) - 1) - (\eta \Lambda(t) - 1)\epsilon_2 \|\mathbf{D}(t)\| \\ &\geq |\mathbf{D}(t)^\top \mathbf{v}_1(t)|(1 - \epsilon_2)(\eta \Lambda(t) - 1) - (\eta \Lambda(t) - 1)|\mathbf{D}(t)^\top \mathbf{v}_1(t)|/c \\ &= |\mathbf{D}(t)^\top \mathbf{v}_1(t)|(\eta \Lambda(t) - 1)(1 - \epsilon_2 - 1/c) \\ &= |\mathbf{D}(t)^\top \mathbf{v}_1(t)|(1 + \tau)(1 - \epsilon_2 - 1/c) \end{aligned}$$

where first inequality holds due to Assumption 3.5 (iii) and the third follows from Assumption 3.6. \square

Proposition E.1. (Proposition 3.1) If $\|\mathbf{D}(t)\| > \|\mathbf{Y}\|$, then $\mathbf{D}(t)^\top \mathbf{F}(t) > 0$.

Proof. This proposition can be proved by simply noticing that $\mathbf{D}(t)^\top \mathbf{F}(t) = \mathbf{D}(t)^\top (\mathbf{D}(t) - \mathbf{Y}(t)) \geq \|\mathbf{D}\|^2 - \|\mathbf{D}\|\|\mathbf{Y}\| > 0$ when $\|\mathbf{D}(t)\| \geq \|\mathbf{Y}\|$. \square

Proposition E.2. (Proposition 3.2) Under Assumption 3.2, if $\|\mathbf{D}(t)\| > \|\mathbf{Y}\|$, then $\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2 < -\frac{4\eta}{n}(\|\mathbf{D}(t)\| - \|\mathbf{Y}\|)^2$.

Proof. Use Assumption 3.2 and notice that $\mathbf{D}(t)^\top \mathbf{F}(t) \geq \|\mathbf{D}(t)\|^2 - \|\mathbf{D}(t)\|\|\mathbf{Y}\| > (\|\mathbf{D}(t)\| - \|\mathbf{Y}\|)^2$. \square

Lemma E.3. (Lemma 3.3) If $\Lambda(t) < 2/\eta$, then for any vector $\mathbf{u} \in \mathbb{R}^n$, $\|\mathbf{u}^\top (\mathbf{I} - \eta \mathbf{M}(t))\| \leq (1 - \eta\alpha)\|\mathbf{u}\|$, where $\alpha = \min\{2/\eta - \Lambda(t), \lambda_{\min}(\mathbf{M}(t))\}$. In particular, replacing \mathbf{u} with $\mathbf{D}(t)$, we can see $\|\mathbf{D}(t+1)\| \leq (1 - \eta\alpha)^2 \|\mathbf{D}(t)\|$.

Proof. For any vector $\mathbf{u} \in \mathbb{R}^n$,

$$\begin{aligned} \|\mathbf{u}^\top (\mathbf{I} - \eta \mathbf{M}(t))\|^2 &= \sum_{i=1}^n \|(\mathbf{u}^\top \mathbf{v}_i(t)) \mathbf{v}_i(t)^\top (\mathbf{I} - \eta \mathbf{M}(t))\|^2 \\ &= \sum_{i=1}^n (1 - \eta \lambda_i)^2 (\mathbf{u}^\top \mathbf{v}_i(t))^2 \leq (1 - \eta\alpha)^2 \sum_{i=1}^n (\mathbf{u}^\top \mathbf{v}_i(t))^2 = (1 - \eta\alpha)^2 \|\mathbf{u}\|^2. \end{aligned}$$

Then the second statement is due to Assumption 3.2. \square

Lemma E.4. (Lemma 3.4) Suppose Assumption 3.5 (iii) holds. $\mathbf{R}(t)$ satisfies the following:

$$\mathbf{R}(t+1) = (\mathbf{I} - \eta \mathbf{M}(t)) \mathbf{R}(t) + \mathbf{e}_1(t), \quad \text{where} \quad \|\mathbf{e}_1(t)\| \leq 6\sqrt{\epsilon_2} \|\mathbf{D}(t)\| (B_\Lambda - 1)$$

Proof. We consider the update rule for $\mathbf{R}(t)$:

$$\begin{aligned} \mathbf{R}(t+1) &= (\mathbf{I} - \mathbf{v}_1(t+1) \mathbf{v}_1(t+1)^\top) (\mathbf{I} - \eta \mathbf{M}(t)) \mathbf{D}(t) \\ &= (\mathbf{I} - \eta \mathbf{M}(t)) \mathbf{R}(t) - (\mathbf{v}_1(t+1) \mathbf{v}_1(t+1)^\top) (\mathbf{I} - \eta \mathbf{M}(t)) (\mathbf{I} - \mathbf{v}_1(t) \mathbf{v}_1(t)^\top) \mathbf{D}(t) \\ &\quad + (\mathbf{I} - \mathbf{v}_1(t+1) \mathbf{v}_1(t+1)^\top) (\mathbf{I} - \eta \mathbf{M}(t)) (\mathbf{v}_1(t) \mathbf{v}_1(t)^\top) \mathbf{D}(t) \\ &= (\mathbf{I} - \eta \mathbf{M}(t)) \mathbf{R}(t) - (1 - \eta \Lambda(t)) (\mathbf{v}_1(t+1) \mathbf{v}_1(t+1)^\top) (\mathbf{I} - \mathbf{v}_1(t) \mathbf{v}_1(t)^\top) \mathbf{D}(t) \\ &\quad + (1 - \eta \Lambda(t)) (\mathbf{I} - \mathbf{v}_1(t+1) \mathbf{v}_1(t+1)^\top) (\mathbf{v}_1(t) \mathbf{v}_1(t)^\top) \mathbf{D}(t) \end{aligned}$$

Now if we decompose $\mathbf{v}_1(t+1) = (1 - \delta(t)) \mathbf{v}_1(t) + \sqrt{1 - (1 - \delta)^2} \mathbf{v}_\perp(t)$, where $\langle \mathbf{v}_1(t), \mathbf{v}_\perp(t) \rangle = 0$. Then we have

$$\|(\mathbf{I} - \mathbf{v}_1(t+1) \mathbf{v}_1(t+1)^\top) \mathbf{v}_1(t) \mathbf{v}_1(t)^\top\|$$

$$\begin{aligned}
&= \|\mathbf{v}_1(t) - (\mathbf{v}_1(t)^\top \mathbf{v}_1(t+1))\mathbf{v}_1(t+1)\| \\
&= \|(1 - (1 - \delta(t))^2)\mathbf{v}_1(t) + \sqrt{1 - (1 - \delta(t))^2}\mathbf{v}_\perp(t)\| \\
&\leq \sqrt{(2\delta(t) - \delta(t)^2)(1 - \delta(t)) + (2\delta(t) - \delta(t)^2)^2} \\
&\leq \sqrt{2\delta(t) + 4\delta(t)^2} \\
&\leq 3\sqrt{\delta(t)} \\
&\leq 3\sqrt{\epsilon_2}.
\end{aligned}$$

Similar result can be obtained for the term $(\mathbf{v}_1(t+1)\mathbf{v}_1(t+1)^\top)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top)$.

Therefore, $\|(\eta\Lambda(t) - 1)(\mathbf{I} - \mathbf{v}_1(t+1)\mathbf{v}_1(t+1)^\top)(\mathbf{v}_1(t)\mathbf{v}_1(t)^\top)\mathbf{D}(t)\| \leq 6\sqrt{\epsilon_2}(B_\Lambda - 1)B_D$. \square

Lemma E.5. (Lemma 3.5) Define an auxiliary sequence $\mathbf{R}'(t)$ by $\mathbf{R}'(0) = \mathbf{R}(0)$, and $\mathbf{R}'(t+1) = (\mathbf{I} - \eta\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top))\mathbf{R}'(t)$. If Assumption 3.4, Assumption 3.5 (iii), Assumption 3.7 hold, and for any time t there exists a quantity $\lambda_r > 0$, such that the smallest eigenvalue of $\mathbf{M}(t)$, i.e. $\lambda_{\min}(\mathbf{M}(t)) > \lambda_r$, then there exists a constant $c_r > 0$ such that $\|\mathbf{R}(t) - \mathbf{R}'(t)\| \leq c_r \frac{B_D(B_\Lambda - 1)\sqrt{\epsilon_2}}{\eta\lambda_r}$.

Proof. First we can see that the eigenvalues of $\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top)$ are all the eigenvalues of \mathbf{M} except the largest one. Hence with Assumption 3.7, all eigenvalues of $\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top)$ are in $(\lambda_r, 1/\eta)$. Thus $\|\mathbf{I} - \eta\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top)\| \leq 1 - \eta\lambda_r$.

Hence by Assumption 3.4 we can get

$$\begin{aligned}
&\|\mathbf{R}(t+1) - \mathbf{R}'(t+1)\| \\
&= \|(\mathbf{I} - \eta\mathbf{M}(t))\mathbf{R}(t) - (\mathbf{I} - \eta\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top))\mathbf{R}'(t) + \mathbf{e}_1(t)\| \\
&= \|(\mathbf{I} - \eta\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top))(\mathbf{R}(t) - \mathbf{R}'(t)) + \mathbf{e}_1(t)\| \\
&\leq \|(\mathbf{I} - \eta\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top))(\mathbf{R}(t) - \mathbf{R}'(t))\| + \|\mathbf{e}_1(t)\| \\
&\leq \|\mathbf{I} - \eta\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1(t)\mathbf{v}_1(t)^\top)\| \|\mathbf{R}(t) - \mathbf{R}'(t)\| + \|\mathbf{e}_1(t)\| \\
&\leq \|\mathbf{R}(t) - \mathbf{R}'(t)\|(1 - \eta\lambda_r) + 3B_D(B_\Lambda - 1)\sqrt{\epsilon_2}.
\end{aligned}$$

Thus if we denote $\|\mathbf{R}(t) - \mathbf{R}'(t)\| - \frac{3B_D(B_\Lambda - 1)\sqrt{\epsilon_2}}{\eta\lambda_r}$ by $p(t)$, and replace $\|\mathbf{R}(t) - \mathbf{R}'(t)\|$ in the inequality above by $p(t) + \frac{3B_D(B_\Lambda - 1)\sqrt{\epsilon_2}}{\eta\lambda_r}$, we can get $|p(t+1)| \leq |p(t)|(1 - \eta\lambda_r)$. Hence, we can see that $|p(t)| < |p(0)| = \frac{3B_D(B_\Lambda - 1)\sqrt{\epsilon_2}}{\eta\lambda_r}$ for any time t . Therefore, we obtain that

$$\|\mathbf{R}(t) - \mathbf{R}'(t)\| < \frac{3B_D(B_\Lambda - 1)\sqrt{\epsilon_2}}{\eta\lambda_r} + |p(0)| < \frac{6B_D(B_\Lambda - 1)\sqrt{\epsilon_2}}{\eta\lambda_r}.$$

Taking $c_r = 6$, we finish the proof. \square

Lemma E.6. For all t in Phase I, under Assumption 3.5 (i) and Assumption 3.7, it holds that $\mathbf{R}'(t)^\top(\mathbf{R}'(t) + \mathbf{Y}) < 0$ where $\mathbf{R}'(t)$ is defined in Lemma 3.5.

Proof. Consider the inner product $\mathbf{R}'(t)^\top \mathbf{v}_i$. Recall \mathbf{v}_i is the i -th eigenvector of \mathbf{M} . By Assumption 3.5 (i), \mathbf{v}_i does not change with t .

By Assumption 3.7, $\lambda_i(t) < 1/\eta$ for $i > 1$, where $\lambda_i(t) \geq 0$ is the corresponding eigenvalue of \mathbf{v}_i .

By definition of $\mathbf{R}'(t)$ and Assumption 3.5 (i), $\mathbf{R}'(t+1) = (\mathbf{I} - \eta\mathbf{M}(t)(\mathbf{I} - \mathbf{v}_1\mathbf{v}_1^\top))\mathbf{R}'(t)$, hence $\mathbf{R}'(t+1)^\top \mathbf{v}_i = (1 - \eta\lambda_i(t))\mathbf{R}'(t)^\top \mathbf{v}_i$ for $i > 1$ and $\mathbf{R}'(t)^\top \mathbf{v}_1 = 0$. Hence for any $i > 1, t > 0$, $\frac{\mathbf{R}'(t)^\top \mathbf{v}_i}{\mathbf{R}'(0)^\top \mathbf{v}_i} = \prod_{j=0}^{t-1} (1 - \eta\lambda_i(j)) \in (0, 1)$.

Since $\mathbf{R}'(0) = \mathbf{R}(0) = (\mathbf{I} - \mathbf{v}_1\mathbf{v}_1^\top)\mathbf{Y}$, hence $\mathbf{R}'(0)^\top \mathbf{v}_i = -\mathbf{Y}^\top \mathbf{v}_i$ for any $i > 1$.

Plug it into the expression $\mathbf{R}'(t)^\top(\mathbf{R}'(t) + \mathbf{Y})$ and we have:

$$\mathbf{D}(t)^\top \mathbf{F}(t) = \sum_{i=2}^n \mathbf{R}'(t)^\top \mathbf{v}_i (\mathbf{R}'(t)^\top \mathbf{v}_i + \mathbf{Y}^\top \mathbf{v}_i) = \sum_{i=2}^n (\mathbf{R}'(t)^\top \mathbf{v}_i)^2 (1 - \frac{\mathbf{R}'(0)^\top \mathbf{v}_i}{\mathbf{R}'(t)^\top \mathbf{v}_i}) < 0$$

□

E.2 Progressive Sharpening under Weaker Assumptions

We prove Lemma 3.1 in Section 3, i.e. progressive sharpening happens under Assumption 3.5 (i) that the set of eigenvectors of the Gram matrix $M(t)$ is fixed throughout. However, empirically, the eigenvectors of $M(t)$ change non-negligibly (see Figure 26). In this section, we show that it is possible to relax Assumption 3.5 (i) to more realistic assumption on the eigenspace of $M(t)$. To this end, we first present a dynamical system view of the dynamics of $D(t)$, then prove Lemma E.7 which is analogue of Lemma 3.1 but under weaker assumptions.

E.2.1 A dynamical system view

Recall that we assume that $D(t)$ follows the gradient flow trajectory: $\frac{dD(t)}{dt} = -M(t)D(t)$. This is a linear dynamical system with changing coefficients $-M(t)$. In order to understand this dynamics, we first consider the linear dynamical system with constant coefficients $-M$, $\frac{dD(t)}{dt} = -MD(t)$, which is much better understood. The corresponding phase portrait is shown in Figure 25(a). Here, O is only fixed point. In this figure, we consider the ball $\mathbb{B} \subset \mathbb{R}^n$ such that the south pole of \mathbb{B} is the origin O and the north pole N of \mathbb{B} is fixed to be $-Y$, which is the initial point $-D(t)$. Recall that $F(t) = D(t) + Y$ and $D(0) = -Y$. We denote the tip of the vector $D(t)$ as D . Hence, $\overrightarrow{DO} = D(t)$ and $\overrightarrow{ND} = F(t)$. A useful geometric observation is the following:

Observation E.1. $D(t) \in \mathbb{B}$ if and only if $D(t)^\top F(t) = D(t)^\top (D(t) - (-Y)) < 0$ (equivalently, the angle $\angle(NDO) \geq \pi/2$).

Let $\{v_i(t)\}_i$ be the set of eigenvectors of $M(t)$. By symmetry, we can assume the direction of all $\{v_i(t)\}_i$ are pointing inside \mathbb{B} (or equivalently, $-v_i(t)^\top Y \geq 0$). We also define the hypercube $\mathbb{H}(t)$, which is defined by $\{v_i(t)\}_i$ as the edges and ON as the diagonal (see the dashed rectangle in Figure 25(a)). All vertices of $\mathbb{H}(t)$ are on the sphere $\partial\mathbb{B}$. In our proof of Lemma 3.1 (in which we assume the eigenvectors do not change), we have shown the entire trajectory of $D(t)$ is contained in $\mathbb{H}(t)$ which is always contained in \mathbb{B} . In particular, we prove that the projection of $D(t)$ onto each v_i decreases monotonically).

Now, we consider the more general case when the eigenvectors $\{v_i(t)\}_i$ may change. In Lemma E.7 presented below, we show that if the eigenvectors change directions relatively slowly (precisely the change rate satisfies (20)), $D(t)$ is still contained in $\mathbb{H}(t)$, hence also in \mathbb{B} . Then by Observation E.1 and (3), the norm $\|A\|$ increases (hence the sharpness increases).

We also note that it is impossible to prove $D(t) \in \mathbb{B}$ for all $t > 0$ without any assumption on the change of eigenvector directions. In particular, if the eigenvector changes fast enough so that $D(t)$ is out of the hypercube $\mathbb{H}(t)$, the trajectory may eventually go out of the ball \mathbb{B} (e.g., the region inside the dashed red box in Figure 25(b)). However, it seems that this region is fairly close to the origin O , and the trajectory does not go very far away from \mathbb{B} (implying $D(t)^\top F(t)$ is not a large positive number). Hence, the above geometric intuition suggests that it is possible to further relax the assumption (inequality (20)) made in E.7 and prove that A-norm increases (hence sharpness increases) until the sharpness reaches close to $2/\eta$, after which gradient flow is not a good approximation of gradient descent anymore. We leave it as an interesting future direction.

E.2.2 Progressive sharpening when eigendirections change slowly

Suppose at initialization $F(0) = 0$. By $D(t) = F(t) - Y$, we know $D(0) = -Y$. With loss of generality, we assume $D(0)^\top v_i(0) > 0$ at initialization for all i (Otherwise, we define $v_i(0) = -v_i(0)$ and this condition holds).

Lemma E.7. *Suppose Assumption 3.3 holds, and the set of (unit) eigenvectors of $M(t)$ is $\{v_i(t)\}$. Assume at all time t and direction $v_i(t)$, the following condition holds:*

$$F(t)^\top \frac{dv_i(t)}{dt} < \lambda_i(t) D(t)^\top v_i(t). \quad (20)$$

Then we have $D(t)^\top F(t) < 0$ for all $t < t_1$, where t_1 is the first time when there exists some i , $D(t)^\top v_i(t) \leq 0$.

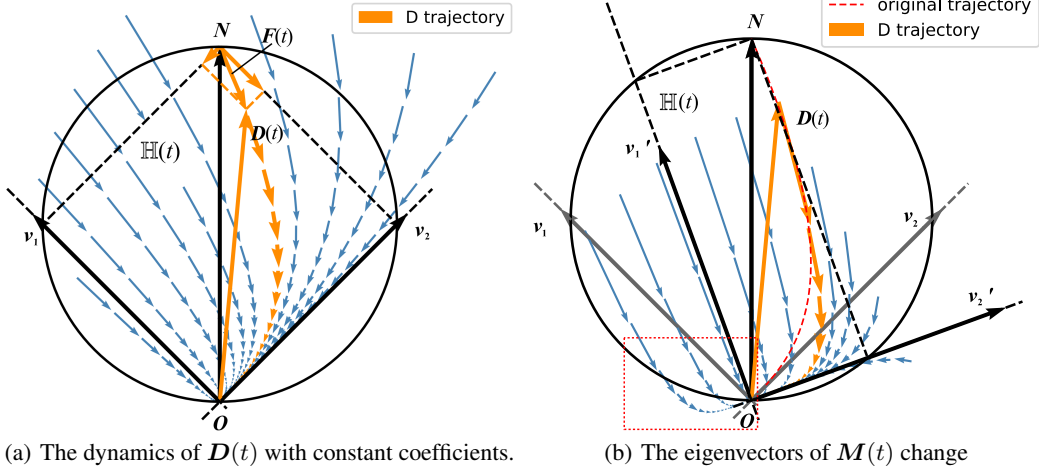


Figure 25: The phase portraits of the dynamics of $D(t)$ (projected to the first two principle directions). We rotate the coordinate system so that south pole of \mathbb{B} is the origin O and the north pole is $-Y$ which is the initial point $D(t)$. v_1 and v_2 are the first two eigenvectors of M . In (b), v_1 and v_2 change directions to v_1' and v_2' . The trajectory may go out of \mathbb{B} from certain regions (the dashed red box).

In the above lemma, we can see that inequality (20) holds if $v_i(t)$ changes relatively slowly. In particular, if $\frac{dv_i(t)}{dt} = 0$ for all i , the lemma reduces to Lemma 3.1. Note that in contrast to the fixed eigenvector assumption, we cannot show that $D(t)^\top v_i(t)$ decreases in all directions all the time (since $M(t)$ may change). Instead, we show that $D(t)$ is always in the hypercube $\mathbb{H}(t)$ under the condition of the lemma.

Proof. We first notice the decomposition

$$D(t)^\top F(t) = \sum_{i=1}^r D(t)^\top v_i(t) F(t)^\top v_i(t).$$

We show for each $v_i(t)$, $D(t)^\top v_i(t) F(t)^\top v_i(t) < 0$ under the conditions of the lemma.

We first consider the dynamics of inner product $D(t)^\top v_i(t)$. According to Assumption 3.3, when $t < t_1$, $D(t)^\top v_i(t) > 0$ due to the continuity. Moreover, by the dynamics of $D(t)$ in Assumption 3.3, we can see

$$\begin{aligned} \frac{dD(t)^\top v_i(t)}{dt} &= -D(t)^\top M(t) v_i(t) + D(t)^\top \frac{dv_i(t)}{dt} \\ &= -\lambda_i(t) D(t)^\top v_i(t) + D(t)^\top \frac{dv_i(t)}{dt} \end{aligned}$$

With this dynamics above and $F(t) = D(t) + Y$, we can get F 's dynamics:

$$\begin{aligned} \frac{dF(t)^\top v_i(t)}{dt} &= -D(t)^\top M(t) v_i(t) + F(t)^\top \frac{dv_i(t)}{dt} \\ &= -\lambda_i(t) D(t)^\top v_i(t) + F(t)^\top \frac{dv_i(t)}{dt} \end{aligned} \quad (D1)$$

With (D1) and $F(t)^\top \frac{dv_i(t)}{dt} \leq \lambda_i(t) D(t)^\top v_i(t)$, we know

$$\frac{dF(t)^\top v_i(t)}{dt} = -\lambda_i(t) D(t)^\top v_i(t) + F(t)^\top \frac{dv_i(t)}{dt} < 0.$$

In this way, we know $F(t)^\top v_i(t)$ always decreases for $t < t_1$. Recall that $F(0) = D(0) + Y = 0$, which makes $F(0)^\top v_i(0) = 0$. So $F(t)^\top v_i(t) < 0$ according to the dynamics. Geometrically, this implies $D(t)$ does not cross the facets of $\mathbb{H}(t)$ incident on the north pole N ⁶ (otherwise $F(t)^\top v_i(t)$

⁶The hypercube $\mathbb{H}(t)$ has $2n$ facets, n of them incident on the south pole O and the others the north pole N .

needs to change sign). See Figure 25). On the other hand, $\mathbf{D}(t)^\top \mathbf{v}_i(t) > 0$ for $t < t_1$, by the definition of t_1 , which is equivalent to say that $\mathbf{D}(t)^\top \mathbf{v}_i(t)$ does not change sign or $\mathbf{D}(t)$ does not cross the facets of $\mathbb{H}(t)$ incident on the south pole O .

Therefore we have $\mathbf{D}(t)^\top \mathbf{v}_i(t) \mathbf{F}(t)^\top \mathbf{v}_i(t) < 0$ and $\mathbf{D}(t) \in \mathbb{H}(t) \subset \mathbb{B}$ for all $t < t_1$, where t_1 is the first time when $\mathbf{D}(t_1)^\top \mathbf{v}_i(t_1) \leq 0$. Summing over all the components $\mathbf{D}(t)^\top \mathbf{v}_i(t) \mathbf{F}(t)^\top \mathbf{v}_i(t)$ for all i , we prove the lemma. \square

E.2.3 Empirical verification of inequality (20)




Figure 26: Verification for inequality (20) when $i = 1, 5, 50$.

We verify the condition (20) in Lemma E.7 via experiments. See Figure 26. We can see that (20) holds for $i = 1, 5, 50$ during the first hundred iterations of the training process. However, for very large i , the eigenvectors corresponding to very small eigenvalues may change quickly and (20) does not hold. The contributions (i.e., $\mathbf{D}(t)^\top \mathbf{v}_i(t) \mathbf{F}(t)^\top \mathbf{v}_i(t)$) from those corresponding to very small eigenvalues tend to cancel each other empirically, hence does not affect the sign of the total sum $\sum_{i=1}^r \mathbf{D}(t)^\top \mathbf{v}_i(t) \mathbf{F}(t)^\top \mathbf{v}_i(t)$. Hence, it is possible to further relax the condition (20) and we leave it as a future direction.

E.3 Proofs of equations in Appendix C

The proof of some equations in Appendix C is omitted, and we list them in this subsection.

Lemma E.8. (Lemma C.1) *The update rule of the residual vector of $\mathbf{D}(t)$:*

$$\mathbf{D}(t+1)^\top - \mathbf{D}(t)^\top = -\eta \mathbf{D}(t)^\top \mathbf{M}(t) + \frac{4\eta^2}{n^2 m} \mathbf{D}(t)^\top (\mathbf{F}(t)^\top \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X} \quad (21)$$

Proof.

$$\begin{aligned} \mathbf{D}(t+1)^\top - \mathbf{D}(t)^\top &= \frac{1}{\sqrt{m}} (\mathbf{A}(t+1)^\top \mathbf{W}(t+1) - \mathbf{A}(t)^\top \mathbf{W}(t)) \mathbf{X} \\ &= \frac{1}{\sqrt{m}} ((\mathbf{A}(t+1)^\top - \mathbf{A}(t)^\top) \mathbf{W}(t+1) + \mathbf{A}(t)^\top (\mathbf{W}(t+1) - \mathbf{W}(t))) \mathbf{X} \\ &= \frac{1}{\sqrt{m}} ((\mathbf{A}(t+1)^\top - \mathbf{A}(t)^\top) \mathbf{W}(t) + \mathbf{A}(t)^\top (\mathbf{W}(t+1) - \mathbf{W}(t))) \mathbf{X} \\ &\quad + \frac{1}{\sqrt{m}} (\mathbf{A}(t+1)^\top - \mathbf{A}(t)^\top) (\mathbf{W}(t+1) - \mathbf{W}(t)) \mathbf{X} \\ &= \frac{2\eta}{nm} ((-\mathbf{W}(t) \mathbf{X} \mathbf{D}(t))^\top \mathbf{W}(t) - \mathbf{A}(t)^\top \mathbf{A}(t) \mathbf{D}(t)^\top \mathbf{X}^\top) \mathbf{X} \\ &\quad + \frac{1}{\sqrt{m}} \left(-\frac{2\eta}{n\sqrt{m}} \mathbf{W}(t) \mathbf{X} \mathbf{D}(t) \right)^\top \frac{2\eta}{n\sqrt{m}} \mathbf{A}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} \end{aligned}$$

$$\begin{aligned}
&= -\mathbf{D}(t)^\top \frac{2\eta}{nm} (\mathbf{X}^\top \mathbf{W}(t)^\top \mathbf{W}(t) \mathbf{X} + \|\mathbf{A}(t)\|^2 \mathbf{X}^\top \mathbf{X}) \\
&\quad + \frac{4\eta^2}{n^2 m} \mathbf{D}(t)^\top (\mathbf{F}(t)^\top \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X} \\
&= -\eta \mathbf{D}(t)^\top \mathbf{M}(t) + \frac{4\eta^2}{n^2 m} \mathbf{D}(t)^\top (\mathbf{F}(t)^\top \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X}
\end{aligned}$$

□

Lemma E.9. (Lemma C.2) *The update rule of the Gram matrix $\mathbf{M}(t)$ is,*

$$\begin{aligned}
\mathbf{M}(t+1) - \mathbf{M}(t) &= -\frac{4\eta}{n^2 m} (2(\mathbf{F}(t)^\top \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X} + \mathbf{F}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{F}(t)^\top) \\
&\quad + \frac{8\eta^2}{n^3 m^2} (\mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}(t)^\top \mathbf{W}(t) \mathbf{X} \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X} \\
&\quad + \frac{8\eta^2}{n^3 m^2} \|\mathbf{A}(t)\|^2 \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X}
\end{aligned}$$

Proof. By the update rule of $\mathbf{A}(t)$ and $\mathbf{W}(t)$ in equation (6) and (7), we have

$$\begin{aligned}
&\mathbf{M}(t+1) - \mathbf{M}(t) \\
&= \frac{2}{nm} ((\|\mathbf{A}(t+1)\|^2 - \|\mathbf{A}(t)\|^2) \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top (\mathbf{W}(t+1)^\top \mathbf{W}(t+1) - \mathbf{W}(t)^\top \mathbf{W}(t)) \mathbf{X}) \\
&\hspace{15em} \text{(Use equation (7))} \\
&= \frac{2}{nm} \left(-\frac{4\eta}{n} \mathbf{D}(t)^\top \mathbf{F}(t) + \eta^2 \left\| \frac{\partial \mathcal{L}}{\partial \mathbf{A}} \right\|^2 \right) \mathbf{X}^\top \mathbf{X} \\
&\quad + \frac{2}{nm} (\mathbf{X}^\top ((\mathbf{W}(t+1)^\top - \mathbf{W}(t)^\top) \mathbf{W}(t) + \mathbf{W}(t)^\top (\mathbf{W}(t+1) - \mathbf{W}(t))) \mathbf{X}) \\
&\quad + \mathbf{X}^\top (\mathbf{W}(t+1)^\top - \mathbf{W}(t)^\top) (\mathbf{W}(t+1) - \mathbf{W}(t)) \mathbf{X}) \\
&\hspace{15em} \text{(Use equation (6))} \\
&= -\frac{8\eta}{n^2 m} \mathbf{D}(t)^\top \mathbf{F}(t) \mathbf{X}^\top \mathbf{X} + \frac{8}{n^3 m^2} \eta^2 (\mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}(t)^\top \mathbf{W}(t) \mathbf{X} \mathbf{D}(t)) \mathbf{X}^\top \mathbf{X} \\
&\quad - \frac{4\eta}{n^2 m} (\mathbf{F}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{F}(t)^\top) \\
&\quad + \frac{8\eta^2}{n^3 m^2} \|\mathbf{A}(t)\|^2 \mathbf{X}^\top \mathbf{X} \mathbf{D}(t) \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{X}
\end{aligned}$$

Rearrange the terms and we complete the proof. □

Lemma E.10. *For all iteration t , if any vector $\mathbf{u} \in \mathbb{R}^n$ satisfies $\mathbf{X}^\top \mathbf{X} \mathbf{u} = 0$, then*

$$\mathbf{Y}^\top \mathbf{u} = \mathbf{D}(t)^\top \mathbf{u} = 0$$

Proof. If $\mathbf{X}^\top \mathbf{X} \mathbf{u} = 0$, then

$$\mathbf{u}^\top \mathbf{X}^\top \mathbf{X} \mathbf{u} = 0 \Rightarrow \mathbf{X} \mathbf{u} = 0.$$

So we have

$$\mathbf{Y}^\top \mathbf{u} = -\mathbf{D}(0)^\top \mathbf{u} = \mathbf{A}^{*\top} \mathbf{W}^* \mathbf{X} \mathbf{u} = 0$$

With the dynamics of $\mathbf{D}(t)$, we have

$$\mathbf{D}(t+1)^\top \mathbf{u} = \mathbf{D}(t)^\top (\mathbf{I}_n - \eta \mathbf{M}^*(t)) \mathbf{u}$$

While

$$\begin{aligned}
\mathbf{M}^*(t) \mathbf{u} &= \mathbf{M}(t) \mathbf{u} - \frac{4\eta}{n^2 m} (\mathbf{D}(t)^\top \mathbf{F}(t)) \mathbf{X}^\top \mathbf{X} \mathbf{u} = \mathbf{M}(t) \mathbf{u} \\
&= \frac{2}{mn} (\|\mathbf{A}(t)\|^2 \mathbf{X}^\top \mathbf{X} + \mathbf{X}^\top \mathbf{W}^\top(t) \mathbf{W}(t) \mathbf{X}) \mathbf{u} = 0
\end{aligned}$$

Thus we have

$$\mathbf{D}(t+1)^\top \mathbf{u} = \mathbf{D}(t)^\top (\mathbf{I}_n - \eta \mathbf{M}^*(t)) \mathbf{u} = \mathbf{D}(t)^\top \mathbf{u} = \dots = \mathbf{D}(0)^\top \mathbf{u} = 0$$

□

Lemma E.11. (Lemma C.7) Here we give the proof of the equation (17):

$$\begin{aligned} \Lambda^*(t+1) - \Lambda^*(t) &= -\frac{8\eta\lambda_1}{mn^2} \left(\mathbf{F}(t)^\top \mathbf{D}(t) + (\mathbf{F}(t)^\top \mathbf{v}_1)(\mathbf{D}(t)^\top \mathbf{v}_1) - \frac{\eta}{2} (\mathbf{D}(t)^\top \mathbf{v}_1)^2 \Lambda^*(t) \right. \\ &\quad \left. - \frac{\eta}{2} \mathbf{R}(t)^\top \mathbf{\Gamma}(t) \mathbf{R}(t) - \eta \mathbf{R}(t)^\top \mathbf{\Gamma}(t) (\mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t)) - \frac{\eta}{mn} \mathbf{R}(t)^\top \left(\frac{m}{d} \mathbf{X}^\top \mathbf{X} \right) \mathbf{R}(t) \right). \end{aligned}$$

Proof. First, since $\mathbf{W}^\top \mathbf{W}$ is initialized as $\frac{m}{d} \mathbf{X}^\top \mathbf{X}$, by definition, we have

$$\mathbf{X}^\top \mathbf{W}^\top \mathbf{W} \mathbf{X} = \frac{mn}{2} \mathbf{\Gamma}(t) + \frac{m}{d} \mathbf{X}^\top \mathbf{X}.$$

Then

$$\begin{aligned} &\Lambda^*(t+1) - \Lambda^*(t) \\ &= \mathbf{v}_1^\top \mathbf{M}(t+1) \mathbf{v}_1 - \mathbf{v}_1^\top \mathbf{M}(t) \mathbf{v}_1 \\ &= \mathbf{v}_1^\top (\mathbf{M}(t+1) - \mathbf{M}(t)) \mathbf{v}_1 \\ &= -\frac{8\eta}{n^2 m} (\lambda_1 \mathbf{F}(t)^\top \mathbf{D}(t) + \lambda_1 (\mathbf{F}(t)^\top \mathbf{v}_1)(\mathbf{D}(t)^\top \mathbf{v}_1)) \\ &\quad + \frac{8\eta^2}{n^3 m^2} (\lambda_1 \mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}^\top \mathbf{W} \mathbf{X} \mathbf{D}(t) + \|\mathbf{A}\|^2 \lambda_1^2 (\mathbf{D}(t)^\top \mathbf{v}_1)^2) \end{aligned}$$

Because

$$\begin{aligned} &\mathbf{D}(t)^\top \mathbf{X}^\top \mathbf{W}^\top \mathbf{W} \mathbf{X} \mathbf{D}(t) \\ &= \mathbf{D}(t)^\top \left(\frac{m}{d} \mathbf{X}^\top \mathbf{X} + \frac{mn}{2} \mathbf{\Gamma}(t) \right) \mathbf{D}(t) \\ &= \frac{m}{d} \mathbf{D}(t)^\top (\mathbf{v}_1 \mathbf{v}_1^\top + \mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{X}^\top \mathbf{X} (\mathbf{D}(t)^\top (\mathbf{v}_1 \mathbf{v}_1^\top + \mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top))^\top \\ &\quad + \frac{mn}{2} \mathbf{D}(t)^\top \mathbf{\Gamma}(t) \mathbf{D}(t) \\ &= \frac{m}{d} \lambda_1 (\mathbf{D}(t)^\top \mathbf{v}_1)^2 + \frac{m}{d} \mathbf{R}(t)^\top \mathbf{X}^\top \mathbf{X} \mathbf{R}(t) + \frac{mn}{2} \mathbf{D}(t)^\top \mathbf{\Gamma}(t) \mathbf{D}(t) \end{aligned}$$

and

$$\lambda_1 \left(\frac{m}{d} + \|\mathbf{A}\|^2 \right) = \frac{mn}{2} \mathbf{v}_1^\top (\mathbf{M}(t) - \mathbf{\Gamma}(t)) \mathbf{v}_1 = \frac{mn}{2} \Lambda^*(t) - \frac{mn}{2} \mathbf{v}_1^\top \mathbf{\Gamma}(t) \mathbf{v}_1$$

we can see that

$$\begin{aligned} &\Lambda^*(t+1) - \Lambda^*(t) \\ &= -\frac{8\eta\lambda_1}{n^2 m} \left(\mathbf{F}(t)^\top \mathbf{D}(t) + (\mathbf{F}(t)^\top \mathbf{v}_1)(\mathbf{D}(t)^\top \mathbf{v}_1) - \frac{\eta}{mn} (\mathbf{D}(t)^\top \mathbf{v}_1)^2 \left(\frac{m}{d} + \|\mathbf{A}\|^2 \right) \lambda_1 \right. \\ &\quad \left. - \frac{\eta}{mn} \cdot \frac{m}{d} \mathbf{R}(t)^\top \mathbf{X}^\top \mathbf{X} \mathbf{R}(t) - \frac{\eta}{2} \mathbf{D}(t)^\top \mathbf{\Gamma}(t) \mathbf{D}(t) \right) \\ &= -\frac{8\eta\lambda_1}{mn^2} \left(\mathbf{F}(t)^\top \mathbf{D}(t) + (\mathbf{F}(t)^\top \mathbf{v}_1)(\mathbf{D}(t)^\top \mathbf{v}_1) - \frac{\eta}{2} (\mathbf{D}(t)^\top \mathbf{v}_1)^2 \Lambda^*(t) \right. \\ &\quad \left. + \frac{\eta}{2} (\mathbf{v}_1^\top \mathbf{\Gamma}(t) \mathbf{v}_1 \cdot (\mathbf{D}(t)^\top \mathbf{v}_1)^2 - \mathbf{D}(t)^\top \mathbf{\Gamma}(t) \mathbf{D}(t)) - \mathbf{R}(t)^\top \left(\frac{m}{d} \mathbf{X}^\top \mathbf{X} \right) \mathbf{R}(t) \right) \end{aligned}$$

Note that

$$\begin{aligned} &\mathbf{v}_1 \mathbf{\Gamma}(t) \mathbf{v}_1^\top (\mathbf{D}(t)^\top \mathbf{v}_1)^2 - \mathbf{D}(t)^\top \mathbf{\Gamma}(t) \mathbf{D}(t) \\ &= \mathbf{D}(t)^\top \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{\Gamma}(t) \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t) - (\mathbf{R}(t)^\top + \mathbf{D}(t)^\top \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{\Gamma}(t) (\mathbf{R}(t) + \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t)) \\ &= -\mathbf{R}(t)^\top \mathbf{\Gamma}(t) \mathbf{R}(t) - 2\mathbf{R}(t)^\top \mathbf{\Gamma}(t) (\mathbf{D}(t)^\top \mathbf{v}_1 \mathbf{v}_1^\top)^\top \end{aligned}$$

Now we finish the proof. □

Lemma E.12. (Equation 18) $\mathbf{R}(t+1) = (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{R}(t) + \eta (\mathbf{v}_1 \mathbf{v}_1^\top \Gamma(t) - \Gamma(t) \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t)$

Proof.

$$\begin{aligned}
\mathbf{R}(t+1) &= (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t+1) \\
&= (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{D}(t) \\
&= (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t) \\
&\quad + (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) (\mathbf{I} - \eta \mathbf{M}^*(t)) (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t) \\
&= (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{R}(t) - \mathbf{v}_1 \mathbf{v}_1^\top (\mathbf{I} - \eta \mathbf{M}^*(t)) (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t) \\
&\quad + (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t) \\
&= (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{R}(t) - \mathbf{v}_1 \mathbf{v}_1^\top (-\eta \Gamma(t)) (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t) \\
&\quad + (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top) (-\eta \Gamma(t)) \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{D}(t) \\
&= (\mathbf{I} - \eta \mathbf{M}^*(t)) \mathbf{R}(t) + \eta (\mathbf{v}_1 \mathbf{v}_1^\top \Gamma(t) - \Gamma(t) \mathbf{v}_1 \mathbf{v}_1^\top) \mathbf{D}(t)
\end{aligned}$$

□

F Missing Preliminaries

Hessian, Fisher information matrix and NTK: The Hessian of the objective $\mathcal{L}(f(\boldsymbol{\theta}))$ is:

$$\nabla_{\boldsymbol{\theta}}^2 \mathcal{L}(f(\boldsymbol{\theta})) = \frac{1}{n} \sum_{i=1}^n \nabla_{\boldsymbol{\theta}}^2 \ell(f(\boldsymbol{\theta}, \mathbf{x}_i))$$

We need the relationship between the Hessian matrix, the Fisher information matrix (FIM) and the neural tangent kernel (NTK). As shown in Papan [26], Martens [23], Bottou et al. [4], the Hessian can be decomposed into two components:

$$\nabla_{\boldsymbol{\theta}}^2 \mathcal{L}(f(\boldsymbol{\theta})) = \frac{1}{n} \sum_{i=1}^n \left(\left. \frac{\partial f(\mathbf{x}_i, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right|_{f(\mathbf{x}_i)}^\top \left. \frac{\partial^2 \ell(z, y_i)}{\partial z^2} \right|_{f(\mathbf{x}_i)} \frac{\partial f(\mathbf{x}_i, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} + \left. \frac{\partial \ell(z, y_i)}{\partial z} \right|_{f(\mathbf{x}_i)} \nabla_{\boldsymbol{\theta}}^2 f(\mathbf{x}_i, \boldsymbol{\theta}) \right)$$

Here $\left. \frac{\partial f(\mathbf{x}_i, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right|_{f(\mathbf{x}_i)} \in \mathbb{R}^{1 \times p}$ is the gradient of f . Papan [26] also demonstrate empirically that the first term, known as ‘‘Gauss-Newton matrix’’, G-term or Fisher information matrix (FIM), dominates the second term in terms of the largest eigenvalue. Thus when considering the sharpness, we can analyze the largest eigenvalue of FIM which is a close proxy of the largest eigenvalue of Hessian. In the binary MSE loss case, $\left. \frac{\partial^2 \ell(z, y_i)}{\partial z^2} \right|_{f(\mathbf{x}_i)} = 2$, which implies that

$$\|\nabla_{\boldsymbol{\theta}}^2 \mathcal{L}(f(\boldsymbol{\theta}))\|_2 \approx \|G\|_2 := \frac{2}{n} \left\| \sum_{i=1}^n \left. \frac{\partial f(\mathbf{x}_i, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right|_{f(\mathbf{x}_i)}^\top \left. \frac{\partial f(\mathbf{x}_i, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right|_{f(\mathbf{x}_i)} \right\|_2 = \frac{2}{n} \left\| \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right|_{f(\mathbf{x}_i)}^\top \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right\|_2$$

where $\frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} := \left(\left. \frac{\partial f(\mathbf{x}_1, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right|_{f(\mathbf{x}_1)}^\top, \dots, \left. \frac{\partial f(\mathbf{x}_n, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right|_{f(\mathbf{x}_n)}^\top \right)^\top \in \mathbb{R}^{n \times p}$. Meanwhile, Karakida et al. [16] pointed out the duality between the FIM and a Gram matrix \mathbf{M} , defined as

$$\mathbf{M} = \frac{2}{n} \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}^\top \quad (22)$$

It also known as the neural tangent kernel NTK (Karakida et al. [16, 15]), which has been studied extensively in recent years (see e.g., [13],[8],[2],[5]). Note that in this paper, we do not assume the training is in NTK regime, in which the Hessian does not change much during training. It is not hard to see that \mathbf{M} and FIM share the same non-zero eigenvalues: if $\mathbf{G}\mathbf{u} = \lambda\mathbf{u}$ for some eigenvector $\mathbf{u} \in \mathbb{R}^p$,

$$\mathbf{M} \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{u} = \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{G}\mathbf{u} = \lambda \frac{\partial \mathbf{F}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \mathbf{u},$$

in other words, λ is also an eigenvalue of \mathbf{M} .

Sharpness: There are various definitions of sharpness in the literature ([20, 30, 9]). In particular, a popular definition of the sharpness is the largest eigenvalue of the Hessian ([30]). Based on the above discussion, in this paper, we adopt the largest eigenvalue of \mathbf{M} , $\Lambda(\boldsymbol{\theta}) = \lambda_{\max}(\mathbf{M})$ as the definition of the sharpness, which is a close approximation of the largest eigenvalue of the Hessian empirically.

Gradient Descent: In this paper, we study the trajectory of gradient descent and gradient flow⁷. We use $\boldsymbol{\theta}(t)$ to denote the parameter at iteration t (or time t) and the sharpness at time t as $\Lambda(t) = \Lambda(\boldsymbol{\theta}(t))$. We similarly define $\mathbf{M}(t)$, $\mathbf{F}(t)$, $\mathbf{D}(t)$, $\mathcal{L}(t)$.

Along GD trajectory, the weight vector $\boldsymbol{\theta}(t)$ is updated in the following way:

$$\boldsymbol{\theta}(t+1) = \boldsymbol{\theta}(t) - \eta \frac{\partial \mathcal{L}(\boldsymbol{\theta}(t))}{\partial \boldsymbol{\theta}(t)}$$

When the learning rate η is infinitesimal, the GD trajectory above is equivalent to gradient flow trajectory. Here we show the gradient flow dynamics of the residual vector $\mathbf{D}(t)$:

$$\frac{d\mathbf{D}(t)}{dt} = \frac{\partial \mathbf{D}(t)}{\partial \boldsymbol{\theta}} \frac{d\boldsymbol{\theta}(t)}{dt} = -\frac{\partial \mathbf{F}(t)}{\partial \boldsymbol{\theta}} \frac{\partial \mathcal{L}(t)}{\partial \boldsymbol{\theta}} = -\frac{2}{n} \frac{\partial \mathbf{F}(t)}{\partial \boldsymbol{\theta}} \frac{\partial \mathbf{F}(t)}{\partial \boldsymbol{\theta}}^\top \mathbf{D}(t) = -\mathbf{M}(t)\mathbf{D}(t) \quad (23)$$

In light of (23), we have the following approximate update rule of $\mathbf{D}(t)$ under gradient descent:

$$\mathbf{D}(t+1) - \mathbf{D}(t) \approx -\eta \mathbf{M}(t)\mathbf{D}(t)$$

⁷Gradient flow is a good approximation of gradient descent in the beginning of the training. See Cohen et al. [6] for more experiments and Ahn et al. [1] for theoretical justification. We assume this fact during the progressive sharpening phase. See Assumption 3.3.