

GRADIENT DESCENT CONVERGES LINEARLY TO FLATTER MINIMA THAN GRADIENT FLOW IN SHALLOW LINEAR NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

We study the gradient descent (GD) dynamics of a depth-2 linear neural network with a single input and output. We show that GD converges at an explicit linear rate to a global minimum of the training loss, even with a large stepsize—about $2/\text{sharpness}$. It still converges for even larger stepsizes, but may do so very slowly. We also characterize the solution to which GD converges, which has lower norm and sharpness than the gradient flow solution. Our analysis reveals a trade off between the speed of convergence and the magnitude of implicit regularization. This sheds light on the benefits of training at the “Edge of Stability”, which induces additional regularization by delaying convergence and may have implications for training more complex models.

1 INTRODUCTION

Training modern machine learning (ML) models like deep neural networks via empirical risk minimization (ERM) requires solving difficult high-dimensional, non-convex, under-determined optimization problems. Although such optimization problems are generally intractable to solve in theory, we train models effectively in practice using algorithms like stochastic gradient descent (SGD). This highlights a disconnect between the worst-case convergence rate of SGD and its convergence on specific ERM problems that arise from training, e.g., neural networks. Even we can solve the ERM problem, typical minimizers of the under-determined objective will overfit the training data and generalize poorly. That said, the specific solutions found by SGD and its variants usually *do* successfully generalize. Understanding how and why we are able to successfully optimize and generalize with these models is of great interest to the ML community and could help fuel continued progress in applied ML.

A key feature of popular ML models, including neural networks, is that the model output is related to the product of model parameters in successive layers. For instance, the output of a 2 layer feed-forward network with ReLU activations has output $\mathbf{W}_2 \text{ReLU}(\mathbf{W}_1 x + \mathbf{b}_1)$, which is closely related to the product of the weight matrices $\mathbf{W}_2 \mathbf{W}_1$. Ultimately, this “self-multiplication” of different model parameters gives rise to the non-convex and under-determined ERM problems that cause such (theoretical) difficulties.

In this work, we distill this parameter self-multiplication property down to its simplest form and comprehensively explain how it affects the training optimization dynamics, the “implicit regularization” of the model parameters, and the “edge-of-stability” dynamics that arise in certain regimes. In particular, we consider the extremely simple problem of learning a univariate linear model $\hat{y} = mx$ to minimize the squared error, except we parameterize the slope as $m = m(\mathbf{a}, \mathbf{b}) = \mathbf{a}^\top \mathbf{b}$ in terms of self-multiplying parameters $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$. This can also be thought of as a depth-2 linear neural network with d hidden units. For training data $\{(x_i, y_i) \in \mathbb{R} \times \mathbb{R}\}_{i=1}^n$, this results in the loss

$$\min_{\mathbf{a}, \mathbf{b} \in \mathbb{R}^d} \bar{L}(\mathbf{a}, \mathbf{b}) := \frac{1}{2n} \sum_{i=1}^n (\mathbf{a}^\top \mathbf{b} x_i - y_i)^2. \quad (1)$$

This objective is equivalent—by rescaling and subtracting a constant—to the even simpler loss¹

$$\min_{\mathbf{a}, \mathbf{b} \in \mathbb{R}^d} L(\mathbf{a}, \mathbf{b}) \quad := \quad \frac{1}{2} (\mathbf{a}^\top \mathbf{b} - \Phi)^2. \quad (2)$$

In what follows, we focus on this formulation and assume $\Phi \geq 0$ w.l.o.g. for simplicity and clarity.

Despite its simplicity, the objective (2), which has also been studied by prior work (Lewkowycz et al., 2020; Wang et al., 2022; Chen & Bruna, 2023; Ahn et al., 2024), is a useful object of study because it has a number of qualitative similarities to more complex and realistic problems like deep learning training objectives. First, it has similar high-level properties—the problem (2) is non-convex and highly under-determined because the set of minimizers constitutes the $(2d - 1)$ -dimensional hyperboloid in \mathbb{R}^{2d} that solves $\mathbf{a}^\top \mathbf{b} = \Phi$. It also exhibits some of the same symmetries as realistic neural networks; for example, $\mathbf{a}^\top \mathbf{b}$ is invariant to swapping “neurons” $(\mathbf{a}_i, \mathbf{b}_i) \leftrightarrow (\mathbf{a}_j, \mathbf{b}_j)$ or to rescaling $(\mathbf{a}_i, \mathbf{b}_i) \rightarrow (c\mathbf{a}_i, c^{-1}\mathbf{b}_i)$. More importantly, the dynamics when optimizing (2) with gradient descent are qualitatively similar to the dynamics of training more complex models. Simultaneously, the problem (2) is simple enough that we can provide a detailed and nearly comprehensive characterization of several different aspects of training. Specifically, we will present the following results:

Convergence of gradient descent Despite the non-convexity of the optimization problem (2), prior work has shown that GD converges to a global minimum from a.e. initialization (Wang et al., 2022). We show that, in fact, it converges at a linear rate. In addition, we identify several phases that depend on the relationship between the stepsize, η ; the scale of the parameters, $\lambda := \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2$; and the residuals, $\epsilon := \mathbf{a}^\top \mathbf{b} - \Phi$. Several of these phases are closely related to the so-called Edge of Stability (EoS) phenomenon (Cohen et al., 2021), where gradient descent decreases the objective (although non-monotonically) despite the largest eigenvalue of the objective’s Hessian matrix being larger than the critical threshold $2/\eta$.

Location of convergence In addition to showing how fast gradient descent converges to *some* global minimizer, we can also describe *which* of the many possible solutions, $\mathbf{a}^\top \mathbf{b} = \Phi$, gradient descent will converge to. To do so, we show that gradient descent implicitly regularizes the “imbalance” of the parameter vectors, quantified by $Q := \sum_{i=1}^d |\mathbf{a}_i^2 - \mathbf{b}_i^2|$, with a larger stepsize generally leading to stronger regularization. This is notably different from the behavior of Gradient Flow (GF) (the $\eta \rightarrow 0$ limit of gradient descent), which *conserves* Q . Since GF is often employed in the literature as an easier-to-analyze approximation of gradient descent (e.g. Du et al., 2018; Tarmoun et al., 2021), our results highlight a potential danger of over-reliance on this approximation. The balance of the parameters is also closely related to the “sharpness” of the solution, i.e. the maximum eigenvalue of the Hessian, which is equal to $\lambda := \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2$ at solutions $\mathbf{a}^\top \mathbf{b} = \Phi$. For the problem (2) specifically, the actual prediction function defined by any solution $\mathbf{a}^\top \mathbf{b} = \Phi$ is the same—after all at any minimizer, $\hat{y} = \Phi x$ regardless of the parameters—so the sharpness is not relevant to generalization. Nevertheless, there is a large body of work in other contexts showing that less sharp minima of the loss tend to generalize better (Hochreiter & Schmidhuber, 1997; Keskar et al., 2016; Smith & Le, 2017; Park et al., 2019), and our analysis shows how the self-multiplying structure of (2) tends to regularize the sharpness.

The key to our analysis is the following pair of observations. On the one hand, gradient descent iterations change the imbalance like $Q(t+1) = (1 - \eta^2 \epsilon(t)^2)Q(t)$, so the imbalance decreases throughout optimization for $0 < \eta < \sqrt{2/|\epsilon(t)|}$. At the same time, the objective L does *not* globally satisfy the Polyak-Łojasiewicz (PL) condition (Polyak, 1963) because the origin is a saddle point, but it *does* satisfy a version of the PL condition along the GD trajectory (see Definition 2), which is sufficient to prove linear convergence of GD to a global minimizer. Interestingly, the PL constant along the GD trajectory, which controls the speed of convergence, is equal to the smallest value of $\lambda(t)$ encountered along the way, which is itself approximately equal to the value of $Q(t)$ at the first time that $\mathbf{a}(t)^\top \mathbf{b}(t) > 0$. Thus, the *stronger* the implicit regularization of Q , the *slower* the convergence of GD, and vice versa, which puts these goals directly at odds with each other.

¹See Lemma 4 in Appendix A for a simple proof.

RELATED WORK

A large body of research has shown empirically that training neural networks with larger learning rates tends to lead to better generalization (LeCun et al., 2002; Bjorck et al., 2018; Li et al., 2019; Jastrzebski et al., 2020). However, in classical settings, convergence can only be guaranteed when the stepsize is small enough that $\lambda_{\max}(\nabla^2 L) < 2/\eta$ throughout optimization (Bottou et al., 2018). Nevertheless, a recent line of work starting with Cohen et al. (2021) observed that when training neural networks, the maximum eigenvalue of the Hessian, or “sharpness”, tends to grow throughout training until it reaches, or even surpasses the critical $2/\eta$ threshold. But rather than diverging, the loss continues to decrease (non-monotonically) while the sharpness continues to hover around $2/\eta$, which is referred to as the Edge of Stability (EoS) phenomenon. Understanding more deeply the training of neural networks with large stepsizes is of great interest.

Problems closely resembling (2) have been studied previously. Ahn et al. (2024) study losses of the form $\ell(ab)$ with $a, b \in \mathbb{R}$ and ℓ any convex, Lipschitz, and even function. The assumption that ℓ is even means it is minimized at zero (this is analogous to $\Phi = 0$ in our case), and they prove convergence to zero from any initialization with any stepsize, but without a rate. However, this result relies crucially on both the loss being Lipschitz and minimized at zero—and this is not surprising, we know that GD diverges on realistic objectives when the stepsize is too large. They also show that the limit point of gradient descent satisfies $|a_\infty^2 - b_\infty^2| \approx \min\{2/\eta, |a_0^2 - b_0^2|\}$, i.e. the imbalance between the weights is implicitly regularized down to the level of $2/\eta$. Chen & Bruna (2023) study (2) with scalar $a, b \in \mathbb{R}$ and prove that the limit point of GD satisfies $a - b \rightarrow 0$ when the stepsize is chosen slightly too large for convergence to any minimizer $ab = \Phi$ to be possible. This is qualitatively similar to our work, but they intentionally choose a too-large stepsize in order to highlight the implicit regularization of the imbalance, while we provide conditions on η under which convergence to a minimizer and some amount of regularization happen simultaneously.

In the most closely related work, Wang et al. (2022) study the exact objective (2) and show that gradient descent using any stepsize up to $\eta \lesssim 4/\text{sharpness}$ —approximately twice as large as the classical threshold of $2/\text{sharpness}$ —eventually converges to a minimizer, but without a rate. They also show some level of implicit regularization of λ , e.g. at convergence $\lambda \leq \frac{2}{\eta}$. In comparison, we provide an explicit convergence rate for GD and give a more detailed connection between this rate and the implicit regularization.

Finally, many papers have studied other models such as matrix factorization or linear neural networks (Saxe et al., 2013; Arora et al., 2018; Gidel et al., 2019; Tarmoun et al., 2021; Xu et al., 2023; Nguegnang et al., 2024), which are more faithful representations of realistic neural networks, but they are also much more difficult to analyze. Due to this difficulty, these results often only apply to gradient flow, or to GD with a very small learning rate, or to GD under additional, hard to interpret assumptions. In this work, we focus on the problem (2) in order to obtain a simpler, easier to interpret set of results.

2 NOTATIONS AND SETTING

The gradient descent dynamics of the parameters are described by

$$\begin{bmatrix} \mathbf{a}(t+1) \\ \mathbf{b}(t+1) \end{bmatrix} = \begin{bmatrix} \mathbf{a}(t) \\ \mathbf{b}(t) \end{bmatrix} - \eta \nabla L(\mathbf{a}(t), \mathbf{b}(t)) = \begin{bmatrix} \mathbf{a}(t) \\ \mathbf{b}(t) \end{bmatrix} - \eta (\mathbf{a}(t)^\top \mathbf{b}(t) - \Phi) \begin{bmatrix} \mathbf{b}(t) \\ \mathbf{a}(t) \end{bmatrix} \quad (3)$$

However, tracking the dynamics of \mathbf{a} and \mathbf{b} is somewhat unwieldy due to the overparametrization of the model $\mathbf{a}^\top \mathbf{b}$. Therefore, it will be convenient to reparametrize the dynamics in terms of the following three quantities, rather than the parameters themselves.

Residuals. We define the residuals, $\varepsilon := \mathbf{a}^\top \mathbf{b} - \Phi = \pm \sqrt{2L(\mathbf{a}, \mathbf{b})}$, which measure the distance to the manifold of minima. We will control the magnitude of ε in order to prove convergence.

Norm of the parameters. We denote by $\lambda := \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2$ the squared Euclidean norm of the parameter vectors. This is interesting because the Hessian of the loss at a solution $\mathbf{a}^\top \mathbf{b} = \Phi$ is

$$\nabla^2 L(\mathbf{a}, \mathbf{b}) = \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix} \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix}^\top + \varepsilon \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix} \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix}^\top \quad (4)$$

so it is rank-1 and has top eigenvalue, or “sharpness”, equal to λ . Also, by the Cauchy-Schwarz and Young inequalities, at any solution $\Phi = \mathbf{a}^\top \mathbf{b} \leq \|\mathbf{a}\| \|\mathbf{b}\| \leq \frac{1}{2} (\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2) = \frac{1}{2} \lambda$, with equality occurring when $\mathbf{a} = \mathbf{b}$. Therefore, the minimum sharpness of any solution is 2Φ .

The scale λ is also useful in our analysis because the evolution of the residuals due to (3) is closely related to λ :

$$\varepsilon(t+1) = \varepsilon(t)(1 - \eta\lambda(t) + \eta^2\varepsilon(t)(\varepsilon(t) + \Phi)). \quad (5)$$

Thus, the term $-\eta\lambda$ is the main cause of the decrease in the magnitude of ε . Finally, the GD dynamics of ε and λ are completely determined by each other

$$\lambda(t+1) = (1 + \eta^2\varepsilon(t)^2)\lambda(t) - 4\eta\varepsilon(t)(\varepsilon(t) + \Phi). \quad (6)$$

The imbalance. To complement the above, we also define $Q_i := \mathbf{a}_i^2 - \mathbf{b}_i^2$ and $Q := \sum_{i=1}^d |Q_i|$. The gradient descent dynamics on the Q_i ’s due to (3) is described by

$$Q_i(t+1) = (1 - \eta^2\varepsilon(t)^2) \cdot Q_i(t) \implies Q(t+1) = |1 - \eta^2\varepsilon(t)^2| \cdot Q(t). \quad (7)$$

From the lack of a term linear in η , we can see that the Q_i ’s are conserved by gradient flow, but *not* by gradient descent, which decreases their magnitude for any $\eta < \sqrt{2}/|\varepsilon|$. This is the essential cause of GD’s implicit regularization effect. In our analysis, we use Q in two ways: (i) the lower bound $\lambda \geq Q$ is key our the proof of GD’s convergence speed, and (ii) we use Q to help characterize the location of convergence.

3 LOCATION OF CONVERGENCE

Our first result concerns which solution is reached by gradient descent:

Theorem 1. *For $\eta < \min \left\{ \frac{1}{2|\varepsilon(0)|}, \frac{2}{\sqrt{\lambda(0)^2 + 4\Phi^2}} \right\}$, at the limit point of gradient descent*

$$0 < \exp\left(-\frac{\sqrt{\eta}\varepsilon(0)^2}{\Phi}\right) < \left| \frac{Q_i(\infty)}{Q_i(0)} \right| < \exp\left(-\eta^2 \sum_{t=0}^{\infty} \varepsilon(t)^2\right) < \exp(-\eta^2\varepsilon(0)^2)$$

A full proof is located in Appendix H. The main idea of Theorem 1 follows from the Q updates (7). In particular, (7) implies that at the limit point of GD:

$$Q(\infty) = Q(0) \prod_{t=0}^{\infty} |1 - \eta^2\varepsilon(t)^2| \stackrel{\eta \leq \sqrt{3}/2|\varepsilon(0)|^{-1}}{\approx} \exp\left(-\eta^2 \sum_{t=0}^{\infty} \varepsilon(t)^2\right) \quad (8)$$

Hence, the amount that Q is regularized is essentially determined by how quickly the loss converges to zero. If the loss decreases quickly, then $\sum_{t=0}^{\infty} \varepsilon(t)^2$ will be small and the amount of regularization will be small. Conversely, if the loss decreases slowly, then $\sum_{t=0}^{\infty} \varepsilon(t)^2$ would be large and $Q(\infty)$ would be made very small. In the next section, we will show that for the stepsize described in Theorem 1, the loss actually converges to zero at a linear rate, so $\sum_{t=0}^{\infty} \varepsilon(t)^2 \approx \varepsilon(0)^2/\eta\mu$ for an certain value of μ . In conclusion, we have the following takeaway:

Takeaway 1: *Gradient descent converges to a solution with lower imbalance than gradient flow, but the imbalance always remains non-zero.*

For a geometric perspective on why reduces the imbalance compared with gradient flow, see Figure 1. Gradient flow conserves Q :

$$\dot{Q}_i = 2(\mathbf{a}_i \dot{\mathbf{a}}_i - \mathbf{b}_i \dot{\mathbf{b}}_i) = 2(\mathbf{a}_i(-\varepsilon \mathbf{b}_i) - \mathbf{b}_i(-\varepsilon \mathbf{a}_i)) = 0 \implies \dot{Q} = 0. \quad (9)$$

On the other hand, because gradient flow curves away from the origin in such a way that Q is constant throughout the trajectory, the discretization error introduced by a gradient descent step results in GD moving “inward” towards the line $\mathbf{a} = \mathbf{b}$, resulting in less imbalance.




Figure 1: The GF trajectory curves away from the origin, so the discretization error of each GD step brings it closer to the line $a = b$.




Figure 2: The qualitative behavior of GD steps differs in each of the three regions.

4 SPEED OF CONVERGENCE

Our second result is a characterization of the convergence speed of gradient descent.

Theorem 2. Let $\eta < \min \left\{ \frac{1}{2|\varepsilon(0)|}, \frac{2}{\sqrt{\lambda(0)^2 + 4\Phi^2}} \right\}$, denote $\bar{\eta} := \min \left\{ \eta, \frac{2}{\sqrt{\lambda(0)^2 + 4\Phi^2}} - \eta \right\}$. Then if $Q(0) \neq 0$, for any $\delta > 0$, gradient descent reaches a point $L(\mathbf{a}(T), \mathbf{b}(T)) \leq \delta$ with

$$T \leq \mathcal{O} \left(\frac{\max\{\log |\varepsilon(0)|, 0\}}{\bar{\eta}Q(0) \exp(\min\{-\mathbf{a}(0)^\top \mathbf{b}(0), 0\})} + \frac{\log \frac{1}{\delta}}{\bar{\eta}Q(0) \exp(\min\{-\mathbf{a}(0)^\top \mathbf{b}(0), 0\}) + \bar{\eta}\Phi} \right)$$

If $\min \left\{ \frac{\sqrt{2}}{\varepsilon}, \frac{2}{\sqrt{\lambda(0)^2 + 4\Phi^2}} \right\} < \eta < \min \left\{ \frac{2}{\varepsilon}, \frac{2}{\lambda} + \frac{2\varepsilon(\varepsilon + \Phi)}{\lambda^3} \right\}$ we have convergence but it could be logarithmically slow.

A full proof can be found in Appendix F. The key idea of the proof is to show from the update equation of the residuals, (5), that roughly speaking

$$\varepsilon(t+1) \approx (1 - \eta\lambda(t))\varepsilon(t) \quad (10)$$

Since we also show that $0 < \eta\lambda(t) < 2$ for all t , this means that GD will converge at a linear rate $(1 - \eta \min_t \lambda(t))$. However, care must be taken to lower bound the smallest parameter norm $\lambda(t)$ visited by GD because if the iterates stray too close to the origin, where $\lambda = 0$, then this linear convergence could be arbitrarily slow. Therefore, the main technical challenge is to bound the GD iterates away from the origin. The key observation is that $\lambda \geq Q$, and we can easily track the evolution of Q using (7) (the updates on λ , (6), are more difficult to control).

To lower bound λ along the trajectory, we break the parameter space into three regions, as depicted in Figure 2. In region A, by the Cauchy-Schwarz and Young inequalities, $2\Phi < 2\mathbf{a}^\top \mathbf{b} \leq 2\|\mathbf{a}\|\|\mathbf{b}\| \leq \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 = \lambda$. So, within region A we have linear convergence with rate $(1 - 2\eta\Phi)$. Region B contains points where λ is arbitrarily small, but since $\varepsilon < 0$ but $\mathbf{a}^\top \mathbf{b} > 0$ here, by (6)

$$\lambda(t+1) = (1 + \eta^2\varepsilon(t)^2)\lambda(t) - 4\underbrace{\eta\varepsilon(t)}_{<0} \underbrace{\mathbf{a}(t)^\top \mathbf{b}(t)}_{>0} \geq (1 + \eta^2\varepsilon(t)^2)\lambda(t) \quad (11)$$

Thus, within region B, GD *increases* the value of λ , so we can lower bound $\lambda(t) \geq \lambda(\tau)$, where τ is the time that the GD iterates entered region B. Furthermore, we always have $\lambda(\tau) \geq Q(\tau)$, so all that is needed is to control $Q(\tau)$ at the time that GD enters region B. Finally, to address region C, the Q updates (7) show $Q(t+1) = (1 - \eta^2\varepsilon(t)^2)Q(t)$ while we prove that $\varepsilon(t+1) \approx (1 - \eta Q(t))\varepsilon(t)$. We use this to argue that the time it takes to substantially decrease Q scales with η^{-2} while the time to exit region C by making $\varepsilon(t+1) > -\Phi$ scales with only η^{-1} . So at the time that GD exits region C, $Q(\tau)$ remains only slightly smaller than $Q(0)$.

Putting this all together, if GD is initialized in region C, it takes $O(\eta^{-1})$ steps to leave, at which point it enters region B with $Q(\tau) \approx Q(0)$, which serves as a lower bound on $\lambda(t)$ until convergence, assuming that the remaining iterates stay in region B. If GD is initialized and remains in region A, then $\lambda(t) \geq 2\Phi$ throughout optimization. To complete the argument, we account for the case where the GD trajectory leaves and/or re-enters a region more than once.

The structure of the proof—which relies on lower bounding the imbalance, $Q(\tau)$, of the GD iterate closest to the origin—is essentially a mirror image of the proof of Theorem 1, and leads to another key takeaway:

Takeaway 2: *The stronger the implicit regularization of the imbalance, Q , the slower the convergence and vice versa.*

The trade-off between the convergence speed and implicit regularization of Q is illustrated by Figure 3. This experiment also indicates that in the EoS regime where the stepsize is larger than what is allowed by Theorem 2 but smaller than approximately $4/\lambda(0)$, GD still converges to a solution, but both the rate of convergence and the amount of regularization of Q have a intricate and chaotic dependence on the initialization and stepsize. Nonetheless, also in this EoS regime we see that the convergence rate and amount of regularization have a generally negative relationship.




Figure 3: We used GD to minimize the problem (2) with $\Phi = 1$ using different stepsizes from various initializations with $\varepsilon(0) = -2$ held constant. The x-axis corresponds to the initial scale $\lambda(0)$ of the initialization, while the y-axis corresponds to the $\eta\lambda(0)$. For stepsizes $\eta \lesssim 2/\lambda(0)$, the amount of regularization $Q(T)/Q(0)$ is limited, but convergence is quick, and quickest around $\eta \approx 1/\lambda(0)$. On the other hand, for very large stepsizes $\eta > 2/\lambda(0)$, convergence is more chaotic; the convergence rate and regularization have a negative relationship but depend sensitively on η and $\lambda(0)$.

5 ON THE STEP SIZE AND THE EDGE OF STABILITY

We characterize here the regime of the dynamics of ε given the size of the learning rate and we sketch a proof for that.

5.1 THE STABLE REGIME.

We know from Theorem 2 that $\eta < 2/\lambda$ along the all trajectory implies linear convergence and we manage to bound λ along the whole trajectory thanks to the following lemma and a characterization of the effect of discretization.

Lemma 1. *The maximum of λ along the trajectory taken by gradient flow on L starting at $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ is*

$$\max_{\text{trajectory}} (\lambda) \leq \sqrt{\lambda(0)^2 + 4\Phi^2}.$$

Instrumental to prove this theorem is the observation that:

Lemma 2. *The quantity $\lambda^2 - 8\epsilon(\epsilon + \Phi) + 4\epsilon^2$ is conserved by the gradient flow on L .*

6 THE CASE OF EDGE OF STABILITY

Recent findings by Cohen et al. (2021) demonstrate that during the training of neural networks with full batch gradient descent at a step size of η , the largest eigenvalue of the Hessian stabilizes right above $2/\eta$. This is somewhat surprising as when the gradients are a linear function of the parameters, if $\eta > \frac{2}{\lambda}$, gradient descent diverges. A very good exemplification of this fact is the case of one dimensional parabola, see Cohen et al. (2021). However, neural networks, surprisingly, converge even if $\eta \geq \frac{2}{\lambda}$. Our results here show that the reason may be the product structure and its interaction with discrete dynamics. Indeed, our results show how convergence happens also for $\eta > 2/\lambda$ but slower. This implies that training happens even though at the edge of stability. Moreover, we prove that the slower the training and the bigger the learning rate, the lower the parameter norm of the solution gradient descent will eventually converge to. This suggests that training at the edge of stability may induce increased implicit regularization.

In the case in which the learning rate is slightly bigger than threshold above $\eta > \frac{2}{\sqrt{\lambda(0)^2 + 4\Phi^2}}$ we show that convergence still happens but at a lower speed. This can be noticed in the case of Fig 3. Precisely, define $\tilde{\eta}$ as

Definition 1. Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, denote by $\tilde{\eta}(\mathbf{a}, \mathbf{b})$ the value

$$\tilde{\eta}(\mathbf{a}, \mathbf{b}) := \frac{2}{\lambda} (1 + \alpha + 4\alpha^2 + 20\alpha^3 + 112\alpha^4 + O(\alpha^5)) \quad \text{with} \quad \alpha := \frac{\epsilon(\epsilon + \Phi)}{\lambda^2}.$$

Then we have that

Proposition 1. *Let $\frac{2}{\sqrt{\lambda(0)^2 + 4\Phi^2}} \leq \eta < \tilde{\eta}$ gradient descent on L converges. However, convergence may happen at only logarithmic speed.*

In the case in which η is even bigger and approaches $4/\lambda$ (Wang et al., 2022, Theorem 1) suggests that there exists a time in which $\eta \leq \frac{2}{\sqrt{\lambda(0)^2 + 4\Phi^2}}$. Just applying our result then we ensure linear convergence, although from that time on. As we showed above, this initial oscillatory phase may, however, last for arbitrarily long time.

7 CONVERGENCE ANALYSIS

In this section we connect the speed of convergence with a lower bound on λ along the trajectory.

7.1 PL CONDITION ALONG THE TRAJECTORIES

To assess convergence rates, we introduce a convergence criterion we call the Polyak-Łojasiewicz Condition Along the Trajectories (PLAT Condition). This criterion serves as a generalization of the traditional PL condition. Specifically adapted to non-convex settings where empirical data indicate quick convergence although the problem itself is not necessarily PL.

Definition 2 (PL Condition Along the Trajectories (PLAT)). Consider the optimization problem $\min_x f(x)$ for $f: S \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$, paired with an optimization algorithm \mathcal{A} . This problem-algorithm pair satisfies the PLAT at $x_0 \in \mathbb{R}^d$ if there exists a constant $\mu(x_0) > 0$ and a stationary point f^* such that for all points x visited by \mathcal{A} starting from x_0

$$\frac{1}{2} \|\nabla_x f(x)\| \geq \mu(x_0) \cdot (f(x) - f^*).$$

While a function f satisfying the traditional PL condition implies it meets the PLAT criteria when equipped with gradient descent (GD) and gradient flow (GF) for every initialization. Notably, however, if a function is PLAT with GD and GF almost everywhere, it may admit saddle points. Almost everywhere linear convergence to global minima is anyway ensured.

7.2 DYNAMICS AND IMPLICATIONS

The introduction of the PLAT criterion facilitates a nuanced understanding of optimization trajectories. In the context of gradient descent with step size η , satisfying PLAT at x_0 means that the function, when restricted to the trajectory of the algorithm,

$$\{x \text{ s.t. } \exists n \in \mathbb{N} \text{ s.t. } x = X_n \text{ where } X_0 = x_0 \text{ and } X_{k+1} := X_k - \eta \nabla f(X_k)\}$$

adheres to the PL condition:

$$\frac{1}{2} \|\nabla_x f(x)\|^2 \geq \mu(x_0) \cdot (f(x) - f^*(x_0)).$$

This, thus, primarily concerns the speed of linear convergence which we lower bound by μ , not whether convergence happen.

Note also that the proof of the fact that PL-condition implies exponential convergence for gradient flow and linear convergence for gradient descent is actually subtly using the weaker assumption of PL-condition along trajectories.

7.3 WARM UP: GRADIENT FLOW ON OUR PROBLEM

We illustrate that the loss function $L(\mathbf{a}, \mathbf{b})$ we take into consideration when equipped with gradient flow (GF), satisfies PLAT almost everywhere. As a first step, note that for any initialization $\alpha_0, \beta_0 \in \mathbb{R}^m$ we have

Lemma 3 (L is locally PL). *L admits local PL constant $\mu(\mathbf{a}, \mathbf{b}) = \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 = \lambda$.*

$$\frac{1}{2} \|\nabla L(\mathbf{a}, \mathbf{b})\|^2 = (\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2) \cdot \frac{1}{2} (\mathbf{a}^\top \mathbf{b} - \Phi)^2 = \lambda \cdot L(\mathbf{a}, \mathbf{b}).$$

However, L does not meet the PL condition globally as $\lambda = 0$ at the saddle point $\alpha, \beta = 0$.

Next note that gradient flow conserves the quantity $Q(\mathbf{a}, \mathbf{b}) := \sum_{i=1}^n |a_i^2 - b_i^2|$ which is always smaller than λ , and the equality is reached only if for all $i \in \{1, 2, \dots, n\}$ at least one between \mathbf{a}_i and \mathbf{b}_i is equal to zero, more formally $Q = \lambda$ if and only if $\sum_i \min\{|\mathbf{a}_i|, |\mathbf{b}_i|\} = 0$. This implies that if at initialization \mathbf{a}, \mathbf{b} satisfy

- $Q(\mathbf{a}, \mathbf{b}) > 0$, then L and gradient flow are PLAT with constant Q .
- $Q(\mathbf{a}, \mathbf{b}) = 0$, then we are on the one dimensional manifold $\mathbf{a} = \pm \mathbf{b}$.
 - If $\mathbf{a} = -\mathbf{b}$ then the problem becomes $L = (\|\mathbf{a}\|^2 + \Phi)^2$ and it converges to the minimum $\mathbf{a} = \mathbf{b} = 0$ with $\mu(\mathbf{a}, \mathbf{a}) = (\|\mathbf{a}\|^2 + \Phi)^2 > \Phi^2 > 0$.
 - If $\mathbf{a} \neq -\mathbf{b}$ and there exists a component i such that $\mathbf{a}_i = \mathbf{b}_i$ instead the components $n_1 < n$ components satisfying $\mathbf{a}_i = -\mathbf{b}_i$ will converge to $\mathbf{a}_i = \mathbf{b}_i = 0$, the the $n - n_1 \geq 1$ other components will converge to the global minimum of L with PL constant given by their norm at initialization $2 \sum_{i \text{ s.t. } \mathbf{a}_i \neq -\mathbf{b}_i} \mathbf{a}_i^2$.

This proves that

Proposition 2. *The loss $L(\mathbf{a}, \mathbf{b})$ equipped with gradient flow is PLAT almost everywhere.*

This proposition demonstrates that L equipped with GF, reliably conforms to the PLAT, showcasing linear convergence, although with a clear differentiation of behavior based on initial conditions and trajectory dynamics.

7.4 GRADIENT DESCENT CONVERGES LINEARLY

We are now ready to demonstrate the formal and clean version of Theorem 2 which is about the robust PL structure of $L(\mathbf{a}, \mathbf{b})$ equipped with gradient descent.

Proposition 3. *$L(\mathbf{a}, \mathbf{b})$ satisfies the PL-condition along the trajectories of gradient descent with almost every step size in $0 < \eta < \min \left\{ \frac{\sqrt{2}}{\varepsilon}, \frac{2}{\sqrt{\lambda(0)^2 + 4\Phi^2}} \right\}$ for almost every initialization $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$.*

Unexpectedly however, we can prove that no matter the step size and the initialization we have linear convergence on our loss. However, the dynamics exhibit particular sensitivity to this hyper parameters. Note, indeed, that for a fixed initialization where $Q \neq 0$, if η is such that there exists a step k along the trajectory where $\eta \cdot (\mathbf{a}^\top \mathbf{b} - \Phi) = 1$ exactly, convergence happen to $\mathbf{a} = \mathbf{b} = 0$ instead of the global minimum. Indeed, in this case, on the next step we have $\mathbf{a}(k+1) = -\mathbf{b}(k+1) = \mathbf{a}(k) - \mathbf{b}(k)$. This implies that when $Q \neq 0$, for almost every η in the allowed range we have linear convergence to a global minimum, in the remaining set of measure zero, linear convergence to the saddle $\mathbf{a} = \mathbf{b} = 0$.

8 CONCLUSION

In this paper, we analyzed the gradient descent dynamics of a depth-2 linear neural network, offering a simplified model to explore training behaviors observed in more complex neural networks. Our key contributions are:

1. **Linear convergence with large step sizes:** We demonstrated that gradient descent converges at a linear rate to a global minimum, even with larger-than-expected step sizes—up to approximately $2/\text{sharpness}$. For even larger step sizes, convergence still occurs but slows down significantly.
2. **Implicit Regularization:** We characterized the solution reached by gradient descent, showing that it implicitly regularizes the parameter imbalance and sharpness, leading to a lower norm solution compared to gradient flow. Notably, as the step size increases, the implicit regularization effect strengthens, flattening the solution.
3. **Trade-off Between Speed and Regularization:** Our analysis uncovered a trade-off between the convergence rate and the degree of implicit regularization. Training at the edge of stability, while slower, induces additional regularization, which may be beneficial for generalization.

These findings provide insight into how larger step sizes affect neural network training dynamics, shedding light on the phenomenon of training near the "Edge of Stability" and its potential benefits for regularization in more complex models.

REFERENCES

- Kwangjun Ahn, Sébastien Bubeck, Sinho Chewi, Yin Tat Lee, Felipe Suarez, and Yi Zhang. Learning threshold neurons via edge of stability. *Advances in Neural Information Processing Systems*, 36, 2024.
- Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. *arXiv preprint arXiv:1810.02281*, 2018.
- Nils Bjorck, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. Understanding batch normalization. *Advances in neural information processing systems*, 31, 2018.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM review*, 60(2):223–311, 2018.
- Lei Chen and Joan Bruna. Beyond the edge of stability via two-step gradient updates. In *International Conference on Machine Learning*, pp. 4330–4391. PMLR, 2023.
- Jeremy Cohen, Simran Kaur, Yanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *International Conference on Learning Representations*, 2021.
- Simon S Du, Wei Hu, and Jason D Lee. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. *Advances in neural information processing systems*, 31, 2018.

- Gauthier Gidel, Francis Bach, and Simon Lacoste-Julien. Implicit regularization of discrete gradient dynamics in linear neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural computation*, 9(1):1–42, 1997.
- 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.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
- Yann LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pp. 9–50. Springer, 2002.
- 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.
- Yuanzhi Li, Colin Wei, and Tengyu Ma. Towards explaining the regularization effect of initial large learning rate in training neural networks. *Advances in neural information processing systems*, 32, 2019.
- Gabin Maxime Nguegnang, Holger Rauhut, and Ulrich Terstiege. Convergence of gradient descent for learning linear neural networks. *Advances in Continuous and Discrete Models*, 2024(1):1–28, 2024.
- Daniel Park, Jascha Sohl-Dickstein, Quoc Le, and Samuel Smith. The effect of network width on stochastic gradient descent and generalization: an empirical study. In *International Conference on Machine Learning*, pp. 5042–5051. PMLR, 2019.
- Boris T Polyak. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963.
- Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- Samuel L Smith and Quoc V Le. A bayesian perspective on generalization and stochastic gradient descent. *arXiv preprint arXiv:1710.06451*, 2017.
- Salma Tarmoun, Guilherme Franca, Benjamin D Haeffele, and Rene Vidal. Understanding the dynamics of gradient flow in overparameterized linear models. In *International Conference on Machine Learning*, pp. 10153–10161. PMLR, 2021.
- Yuqing Wang, Minshuo Chen, Tuo Zhao, and Molei Tao. Large learning rate tames homogeneity: Convergence and balancing effect. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=3tbDrs77LJ5>.
- Ziqing Xu, Hancheng Min, Salma Tarmoun, Enrique Mallada, and René Vidal. Linear convergence of gradient descent for finite width over-parametrized linear networks with general initialization. In *International Conference on Artificial Intelligence and Statistics*, pp. 2262–2284. PMLR, 2023.

A ON THE OBJECTIVE

Lemma 4. For any $\mathbf{a}, \mathbf{b}, \{(x_i, y_i)\}_{i=1}^n$,

$$\bar{L}(\mathbf{a}, \mathbf{b}) = \frac{\sum_{i=1}^n x_i^2}{2n} (\mathbf{a}^\top \mathbf{b} - c)^2 + \text{Const} = \left[\frac{1}{n} \sum_i x_i^2 \right] L(\mathbf{a}, \mathbf{b}) + \text{Const}.$$

where $c = \frac{\sum_{i=1}^n x_i y_i}{\sum_{i=1}^n x_i^2}$ and $\text{Const} = \frac{1}{2n} \left(\sum_{i=1}^n y_i^2 - \frac{(\sum_{i=1}^n x_i y_i)^2}{\sum_{i=1}^n x_i^2} \right)$ does not depend on \mathbf{a}, \mathbf{b} .

Proof. Let \mathbf{x} denote the vector whose i th entry is x_i , and let \mathbf{y} denote the vector whose i th entry is y_i . Then we can write

$$\begin{aligned}\bar{L}(\mathbf{a}, \mathbf{b}) &= \frac{1}{2n} \|\mathbf{a}^\top \mathbf{b} \mathbf{x} - \mathbf{y}\|^2 = \frac{1}{2n} ((\mathbf{a}^\top \mathbf{b})^2 \|\mathbf{x}\|^2 - 2\mathbf{a}^\top \mathbf{b} \langle \mathbf{x}, \mathbf{y} \rangle + \|\mathbf{y}\|^2) \\ &= \frac{\|\mathbf{x}\|^2}{2n} \left((\mathbf{a}^\top \mathbf{b})^2 - 2\mathbf{a}^\top \mathbf{b} \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{x}\|^2} + \frac{\|\mathbf{y}\|^2}{\|\mathbf{x}\|^2} \right) \\ &= \frac{\|\mathbf{x}\|^2}{2n} \left(\left(\mathbf{a}^\top \mathbf{b} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{x}\|^2} \right)^2 + \frac{\|\mathbf{y}\|^2}{\|\mathbf{x}\|^2} - \frac{\langle \mathbf{x}, \mathbf{y} \rangle^2}{\|\mathbf{x}\|^4} \right) \quad (12)\end{aligned}$$

Rewriting this in terms of the x_i 's and y_i 's completes the proof. \square

Note, thus, that all our proofs work on \bar{L} , we thus have to rescale $\varepsilon, \lambda, Q, \eta$ accordingly. Precisely,

$$\lambda \curvearrowright \left[\frac{1}{n} \sum_i x_i^2 \right] \lambda, \quad Q \curvearrowright \left[\frac{1}{n} \sum_i x_i^2 \right] Q, \quad \varepsilon \curvearrowright \left[\frac{1}{n} \sum_i x_i^2 \right] \varepsilon, \quad \text{and} \quad \eta \curvearrowright \left[\frac{1}{n} \sum_i x_i^2 \right] \eta. \quad (13)$$

Analogously, note that if $\Phi < 0$ nothing changes in the analysis of the dynamics. When $\mathbf{a} \neq -\mathbf{b}$ just change \mathbf{a} to $-\mathbf{a}$ and apply the same analysis as before.

B FROM THE RESIDUALS TO THE LOSS

First note that if ε converges exponentially to zero, then loss L converges exponentially to its minimum.

Lemma 5. *Assume $|\varepsilon(k)|$ converges linearly fast with rate $(1 - \eta\mu) < 1$. Then L converges linearly fast with rate $(1 - \eta\mu)^2$. In particular, let $\delta > 0$, the loss L is smaller than δ in a number of steps t that satisfies*

$$t \leq \frac{\log L_0 - \log(\delta)}{\eta\mu}.$$

Indeed note that for how we defined ε we have that $L = \varepsilon^2$, thus $L(k+1) = |\varepsilon(k+1)| \leq (1 - \eta\mu)|\varepsilon(k)|^2$. Note that this lemma allows us to deal with the convergence of ε instead of the convergence of L and infer the convergence of L . Indeed, if the residuals ε converge linearly with rate $(1 - \eta\mu) < 1$, then the time it takes to converge is such that $\sqrt{\delta} \geq (1 - \eta\mu)^t \sqrt{L_0}$ which is

$$t \leq \frac{\log L_0 - \log(\delta)}{-\log(1 - \eta\mu)} \leq \frac{\log L_0 - \log(\delta)}{\eta\mu}. \quad (14)$$

From now on we will deal with convergence of residuals only.

C BOUNDING THE FINAL SHARPNESS

C.1 SIZE OF λ FOR GRADIENT FLOW

Note that we can characterize the norm λ_∞ found by gradient flow by noticing that

Lemma 6. *Along the gradient flow trajectory, the following quantity is conserved*

$$\lambda^2 - 8\varepsilon(\varepsilon + \Phi) + 4\varepsilon^2.$$

Proof. The gradient flow dynamics are described by

$$\begin{bmatrix} \dot{\mathbf{a}} \\ \dot{\mathbf{b}} \end{bmatrix} = -\nabla L(\mathbf{a}, \mathbf{b}) = -\varepsilon \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix} \quad (15)$$

First, we compute

$$\dot{\lambda} = \frac{d}{dt} [\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2] \quad (16)$$

$$= 2 \langle \mathbf{a}, \dot{\mathbf{a}} \rangle + 2 \langle \mathbf{b}, \dot{\mathbf{b}} \rangle \quad (17)$$

$$= -2\varepsilon \langle \mathbf{a}, \mathbf{b} \rangle - 2\varepsilon \langle \mathbf{b}, \mathbf{a} \rangle \quad (18)$$

$$= -4\varepsilon(\varepsilon + \Phi) \quad (19)$$

and

$$\dot{\varepsilon} = \frac{d}{dt} [\langle \mathbf{a}, \mathbf{b} \rangle - \Phi] \quad (20)$$

$$= \langle \mathbf{a}, \dot{\mathbf{b}} \rangle + \langle \dot{\mathbf{a}}, \mathbf{b} \rangle \quad (21)$$

$$= -\varepsilon \|\mathbf{a}\|^2 - \varepsilon \|\mathbf{b}\|^2 \quad (22)$$

$$= -\lambda \varepsilon \quad (23)$$

Finally, straightforward calculation confirms:

$$\frac{d}{dt} [\lambda^2 - 8\varepsilon(\varepsilon + \Phi) + 4\varepsilon^2] = 2\lambda\dot{\lambda} - 8\varepsilon\dot{\varepsilon} - 8\dot{\varepsilon}(\varepsilon + \Phi) + 8\varepsilon\dot{\varepsilon} \quad (24)$$

$$= 2\lambda\dot{\lambda} - 8\dot{\varepsilon}(\varepsilon + \Phi) \quad (25)$$

$$= 2\lambda(-4\varepsilon(\varepsilon + \Phi)) - 8(-\lambda\varepsilon)(\varepsilon + \Phi) \quad (26)$$

$$= 0 \quad (27)$$

which completes the proof. \square

These two lemmas imply that the norm of the solution found by gradient descent will always be smaller than λ_∞ and since $\varepsilon_\infty = 0$ we can compute it using the formula above, precisely.

Lemma 7. *Let $\eta \leq \frac{\sqrt{2}}{\varepsilon}$ along the whole trajectory. Then the norm λ every point along the trajectory of gradient descent is*

$$\lambda \leq \sqrt{\lambda(0)^2 - 8\varepsilon(0)(\varepsilon(0) + \Phi) + 4\varepsilon(0)^2} \leq \sqrt{\lambda(0)^2 + 4\Phi^2}.$$

Note indeed that

$$\begin{aligned} \sqrt{\lambda(0)^2 - 8\varepsilon(0)(\varepsilon(0) + \Phi) + 4\varepsilon(0)^2} &= \sqrt{(\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2)^2 + 8\Phi\mathbf{a}^\top\mathbf{b} - 4(\mathbf{a}^\top\mathbf{b})^2} \\ &\leq \max_{\mathbf{a}^\top\mathbf{b}} \sqrt{(\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2)^2 + 8\Phi\mathbf{a}^\top\mathbf{b} - 4(\mathbf{a}^\top\mathbf{b})^2} \\ &= \sqrt{(\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2)^2 + 4\Phi^2} \end{aligned} \quad (28)$$

attained when $\mathbf{a}^\top\mathbf{b} = \Phi$.

Definition 3 (Maximal Sharpness $\bar{\lambda}$). We denote by $\bar{\lambda}$ and we call maximal sharpness the value

$$\sqrt{(\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2)^2 + 4\Phi^2}.$$

C.2 SIZE OF λ FOR GRADIENT DESCENT

Equation 6 tells us that λ decreases if $\eta \leq \frac{4\mathbf{a}^\top\mathbf{b}}{\lambda\varepsilon}$. Note that if $\varepsilon, \mathbf{a}^\top\mathbf{b} > 0$ this is implied by $\eta \leq 2/\varepsilon$. If $\mathbf{a}^\top\mathbf{b} > 0$ but $\varepsilon \leq 0$, and $\eta \leq \frac{\sqrt{2}}{\varepsilon}$, λ grows but it has a maximum for λ is $\bar{\lambda}$. And if $\varepsilon, \mathbf{a}^\top\mathbf{b} < 0$ this is again implied by $\eta \leq 2/\varepsilon$.

This implies that if we have $\eta \leq \frac{\sqrt{2}}{\varepsilon}$ for all ε along the trajectory, at every step, either λ decreases or is smaller than 2Φ .

Note that for every \mathbf{a}, \mathbf{b} , if $\eta \leq \frac{2}{\varepsilon}$ and $\eta \leq \frac{2}{\bar{\lambda}}$, then we have that

$$\varepsilon(1) = \left(1 - \frac{2}{\bar{\lambda}}\right) \varepsilon + \frac{4}{\bar{\lambda}^2} \varepsilon^2 (\varepsilon + \Phi) \quad (29)$$

note that the min of $\varepsilon^2(\varepsilon + \Phi)$ is for $\varepsilon = \Phi/2$.

$$\left\| \left(1 - \frac{2}{\lambda} \lambda\right) \varepsilon + \frac{1}{\lambda^2} \varepsilon \Phi^2 \right\| < |\varepsilon|. \quad (30)$$

This shows that

Lemma 8. *If $\eta < \min \left\{ \frac{\sqrt{2}}{\varepsilon}, \frac{2}{\lambda} \right\}$ along the whole gradient descent trajectory, then the norm λ is smaller than $\bar{\lambda}$.*

D SPEED OF CONVERGENCE

D.1 CONTINUOUS DYNAMICS

In the case of gradient flow the pairs (a, b) along the trajectory satisfy a PL condition with $\mu(a, b) = \|a\|^2 + \|b\|^2$, indeed note that $L(a, b)$ satisfies

$$(\mathbf{a}^\top \mathbf{b} - \Phi)^2 (\|a\|^2 + \|b\|^2) = \|\nabla L(a, b)\|^2 = \mu(t) \cdot L(a, b).$$

Note that for all i the quantity $Q_i = a_i^2 - b_i^2$ is conserved along the trajectory, indeed

$$\frac{d}{dt} (a_i(t)^2 - b_i(t)^2) = 2\varepsilon(a_i b_i - a_i b_i) = 0.$$

Thus we have that $Q(0) \neq 0$ is a lower bound to μ along the whole trajectory.

Proposition 4 (Convergence of gradient flow). *Let $a, b \in \mathbb{R}$ such that $Q \neq 0$. The gradient flow starting from a, b converges exponentially fast with rate at least Q to the point \tilde{a}, \tilde{b} which satisfies that (i) $\tilde{a}^\top \tilde{b} = \Phi$ and (ii) for all i that $a_i(0)^2 - b_i(0)^2 = \tilde{a}_i^2 - \tilde{b}_i^2$ and $\text{sign}(\tilde{a}_i - \tilde{b}_i) = \text{sign}(a_i(0) - b_i(0))$.*

D.2 LOWER BOUND TO $\mu(t)$ IN THE DISCRETE CASE.

Note that the derivative in time of $\mu(\mathbf{a}(t), \mathbf{b}(t))$ is

$$\dot{\mu} = -4 \left[\frac{1}{n} \sum x_i^2 \right]^2 (\mathbf{a}^\top \mathbf{b} - \Phi) \mathbf{a}^\top \mathbf{b}. \quad (31)$$

It thus decreases when $\mathbf{a}^\top \mathbf{b} > \Phi$ and when $\mathbf{a}^\top \mathbf{b} < 0$ and $\mathbf{a}^\top \mathbf{b} < \Phi$, it grows when $\mathbf{a}^\top \mathbf{b} > 0$ and $\mathbf{a}^\top \mathbf{b} < \Phi$. This means that

- When $\mathbf{a}(0)^\top \mathbf{b}(0) > \Phi$, in Area A of Figure 2, we can bound

$$\mu(t) \geq \inf_{\mathbf{a}^\top \mathbf{b} > \Phi} \left[\frac{1}{n} \sum x_i^2 \right] (\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2) = 2\Phi.$$

Thus in this area we have that $2\Phi \leq \lambda \leq \bar{\lambda}$.

- When $\mathbf{a}(0)^\top \mathbf{b}(0) > 0$ and $\mathbf{a}(0)^\top \mathbf{b}(0) < \Phi$, in Area B of Figure 2, we can bound

$$\mu(t) \geq \mu(0) = \left[\frac{1}{n} \sum x_i^2 \right] (\|\mathbf{a}(0)\|^2 + \|\mathbf{b}(0)\|^2).$$

Thus in this area we have that $\lambda(0) \leq \lambda \leq 2\Phi \leq \bar{\lambda}$, where $\lambda(0) \geq Q_0 > 0$ is the norm of the first step in this area, when $Q \neq 0$.

- When $\mathbf{a}(0)^\top \mathbf{b}(0) < 0$, in Area C of Figure 2, the residuals decreases until $\mathbf{a}^\top \mathbf{b} = 0$. Thus the lowest point for Q will be at the step τ that is the first step in which $\mathbf{a}^\top \mathbf{b} \geq 0$. This implies that the quantity $a^2 + b^2$ will be at its minimum either at time τ or $\tau - 1$

$$\mu(t) \geq \min\{\mu(\tau - 1), \mu(\tau)\} \quad \text{where } \tau = \min_{t \in \mathbb{N}} \{a(t)^\top b(t) > 0\}.$$

In particular $\mu(t) \geq \min\{\mu(\tau - 1), \mu(\tau)\} \geq Q(\tau_1)$, we need to show that when $Q \neq 0$ then $Q(\tau_1) \neq 0$. Thus in this area we will prove in the next section that we have that $Q(\tau_1) \leq \lambda \leq 2\Phi \leq \bar{\lambda}$.

This concludes the argument for all the cases except for $\mathbf{a}(0)^\top \mathbf{b}(0) < 0$. We will now bound $|Q_i(\tau_1)|$ in terms of $|Q(0)|$, the learning rate $\eta > 0$, and $\mathbf{a}(0)^\top \mathbf{b}(0)$.

E LOWER BOUND ON μ IN AREA C

E.1 OPTIMAL STEP SIZE

Note that if $\eta \geq \frac{\lambda + \sqrt{\lambda^2 + 4(\epsilon + \Phi)^2}|\epsilon|}{2\epsilon(\epsilon + \Phi)}$ then in the next step we are landing directly in Area B or A. Note in Area C, whenever $\eta\epsilon(0) > 1$ we are always in this setting above. In these cases gradient descent makes one step in Area C and leaves, so the convergence analysis continues with the ones in Areas B and A.

E.2 OTHER STEP SIZES

The difficult phase to deal with analytically is the one where the dynamics stays in Area C for long.

We compute here a lower bound on $|Q_i(\tau_1)|$. The idea here is that the residuals $\mathbf{a}^\top \mathbf{b} - \Phi$ will converge as $\exp(-\eta t)$ and the quantity $Q_i(t)$ at most as $\exp(-\eta^2 t)$, thus $a(t)^\top b(t)$ crosses 0 before $|Q_i(t)|$ gets too small.

Note that at every step of gradient descent we have the following updates on the following quantities

$$\begin{aligned} a(t+1)^\top b(t+1) - \Phi &= (1 - \eta(\|a(t)\|^2 + \|b(t)\|^2)) (a(t)^\top b(t) - \Phi) \\ &\quad + \underbrace{\eta^2 (a(t)^\top b(t) - \Phi)^2 a(t)^\top b(t)}_{\text{positive}} \end{aligned} \quad (32)$$

and

$$a_i(t+1)^2 - b_i(t+1)^2 = (1 - \eta^2 (a(t)^\top b(t) - \Phi)^2) (a_i(t)^2 - b_i(t)^2). \quad (33)$$

Thus we have that

$$a(t+1)b(t+1) - \Phi > (1 - \eta(a^2(t) + b^2(t))) (a(t)^\top b(t) - \Phi) \quad (34)$$

when $a(t)^\top b(t) < 0$.

Assumption 1. Assume that $\eta < \min \left\{ \frac{\sqrt{2}}{\epsilon}, \frac{2}{\lambda} \right\}$.

Definition 4. Define $\alpha = \eta < \min \left\{ \frac{\sqrt{2}}{\epsilon}, \frac{2}{\lambda} \right\}$. Define the sequence $\{z_k, w_k\}_{k=0}^\infty$ such that $z_0 = \epsilon(0) < -\Phi$, $w_0 = Q_0 > 0$, and for all $k \in \mathbb{N}$ we have

$$\begin{aligned} z_{k+1} &= \left(1 - \alpha \max \left\{ w_k, -2z_k - 2\Phi \right\}\right) z_k \\ w_{k+1} &= (1 - \alpha^2 z_k^2) w_k. \end{aligned} \quad (35)$$

Define $\tau_1 := \min_{k \in \mathbb{N}} \{z_k > -\Phi\}$.

Note that we have

Lemma 9 (Bounding with the sequences). *Under Assumption 1, for all $k < \tau_1$ such that $z_k < 0$ we have*

$$z_k \leq \epsilon_k \quad \text{and} \quad w_k < Q_k.$$

Moreover, $w_k, -z_k \geq 0$ are strongly monotone decreasing for $k < \tau_1$ and $a(\tau_1)^\top b(\tau_1) > 0$, thus for all $k \leq \tau_1$ we have $\alpha < 2/\max\{-2z_k, w_k\}$.

Proof. Note that this is the case for $k = 0$. As for the inductive step, Eq. 34, Cauchy-Schwartz inequality, and Eq. 35 establish the first point. Note that $z_{\tau_1-1} + \Phi < a(\tau_1 - 1)b(\tau_1 - 1) < 0$, then the first point and the definition of τ imply that $0 < z_{\tau_1} + \Phi < a(\tau_1)^\top b(\tau_1)$. Note that after the first step $w_1 < w_0$ and $z_1 > z_0$ since Cauchy-Schwartz implies that $\alpha < 2/\max\{-2z_0, w_0\}$. Inductively, for all i we have $\alpha < 2/\max\{-2z_i, w_i\}$, thus fact that $z_i < 0$ for all $k < \tau_1$ implies that $w_{i+1} < w_i$ is strongly monotonically decreasing, that $z_{i+1} > z_i$ is strongly monotonically increasing, and that $\alpha < 2/\max\{-2z_{i+1}, w_{i+1}\}$. \square

As explained before, for all t we have $\mu(t) \geq \max\{\mu(\tau_1), \mu(\tau_1 - 1)\}$ and $\mu(t) \geq \sum_i |a_i^2(t) - b_i^2(t)| \geq w_t \geq w_{t+1}$. Thus for all t we have $\mu(t) \geq w_{\tau_1}$. This and the lemma above show that

Lemma 10. *Under Assumption 1, we have that $\mathbf{a}(\tau_1)^\top \mathbf{b}(\tau_1) > -\Phi$ and for all $t \in \mathbb{N}$ we have $\mu(t) \geq w_{\tau_1}$.*

Lemma 11 (Rate of convergence 1 - Sequence.). *Under Assumption 1, if $w_0 \geq -2z_0 - 2\Phi$, define*

$$c_1 := \frac{w_0}{2} - \frac{\sqrt{w_0(w_0 - 4\alpha z_0^2)}}{2} > 0, \text{ then}$$

$$c_1 < w_{\tau_1} < w_0 - \alpha^2 (\Phi)^2 \frac{(\mathbf{a}(0)^\top \mathbf{b}(0))^2}{w_0} \quad (36)$$

and

$$\frac{1}{\alpha(w_0)^{3/2}} < \tau_1 < \frac{\mathbf{a}(0)^\top \mathbf{b}(0)}{\alpha c_1^{3/2}} (\Phi)^{-1} + 1. \quad (37)$$

Proof. Note that for all $k < \tau_1$ we have

$$\frac{(z_{k+1} - z_k)^2}{w_{k+1} - w_k} = \frac{\alpha^2 w_k^2 z_k^2}{-\alpha^2 z_k^2 w_k} = -w_k. \quad (38)$$

Note that $z_{\tau_1-1} - z_0 < |\mathbf{a}(0)^\top \mathbf{b}(0)| \leq z_{\tau_1} - z_0$. We thus obtain that

$$\mathbf{a}(0)^\top \mathbf{b}(0) \sim z_0 - z_{\tau_1} = \sum_{k=0}^{\tau_1-1} z_k - z_{k+1} = \sum_{k=0}^{\tau_1-1} \sqrt{w_k - w_{k+1}} \cdot w_k = \alpha \sum_{k=0}^{\tau_1-1} z_k (w_k)^{3/2} \quad (39)$$

This implies that

$$\alpha(\tau_1 - 1)\Phi(w_{\tau-2})^{3/2} < \mathbf{a}(0)^\top \mathbf{b}(0) \leq \alpha\tau_1 z_0 (w_0)^{3/2}. \quad (40)$$

Next we proceed bounding w_{τ_1} so that we can bound τ_1 . Note that the fact that $z_k < -\Phi < 0$ for all $k < \tau_1$ and Sedrakyan's lemma imply that

$$\begin{aligned} w_0 - w_{\tau_1} &= \sum_{i=0}^{\tau_1-1} w_k - w_{k+1} = \sum_{k=0}^{\tau_1-1} \frac{(z_{k+1} - z_k)^2}{w_k} > \frac{(z_{\tau_1} - z_0)^2}{\sum_{i=0}^{\tau_1-1} w_k} \\ &> \frac{(\mathbf{a}(0)^\top \mathbf{b}(0))^2}{\sum_{i=0}^{\tau_1-1} (1 - \alpha^2 z_{\tau_1-1}^2) w_0} = \alpha^2 z_{\tau_1}^2 \frac{(\mathbf{a}(0)^\top \mathbf{b}(0))^2}{w_0} \\ &> \frac{\mathbf{a}(0)^\top \mathbf{b}(0)}{\tau_1 w_0}. \end{aligned} \quad (41)$$

And this implies that

$$w_{\tau_1} = w_0 + (w_{\tau_1} - w_0) < w_0 - \alpha^2 (\Phi)^2 \frac{(\mathbf{a}(0)^\top \mathbf{b}(0))^2}{w_0} \quad (42)$$

Moreover, we have

$$\begin{aligned} w_0 - w_{\tau_1} &= \sum_{k=0}^{\tau_1-1} w_k - w_{k+1} = \alpha^2 \sum_{i=0}^{\tau_1-1} z_k^2 w_k < \alpha^2 \sum_{i=0}^{\tau_1-1} z_k^2 (1 - \alpha^2 z_{\tau_1}^2)^k w_0 \\ &< \alpha^2 z_0^2 w_0 \sum_{i=0}^{\tau_1-1} (1 - \alpha w_{\tau_1})^k (1 - \alpha^2 z_{\tau_1}^2)^k < \alpha z_0^2 w_0 \frac{1}{w_{\tau_1} + \alpha z_{\tau_1}^2 - \alpha^2 w_{\tau_1} z_{\tau_1}^2} \\ &< \alpha \frac{z_0^2 w_0}{w_{\tau_1}}. \end{aligned} \quad (43)$$

This implies with $0 < w_{\tau_1} < w_0$ that $w_{\tau_1}(w_0 - w_{\tau_1}) < \alpha z_0^2 w_0$, thus

$$w_{\tau_1}^2 - w_0 w_{\tau_1} + \alpha z_0^2 w_0 > 0. \quad (44)$$

Note that $w_0 > -2z_0$ and $\alpha < -2/z_0$ implies that $w_0 \geq \alpha z_0^2$ then, solving, we obtain

$$0 < c_1 := \frac{w_0}{2} - \frac{\sqrt{w_0(w_0 - 4\alpha z_0^2)}}{2} < w_{\tau_1} < w_0 - \alpha^2 (\Phi)^2 \frac{(\mathbf{a}(0)^\top \mathbf{b}(0))^2}{w_0} \quad (45)$$

$w_{\tau_1} > \frac{w_0}{2} - \frac{\sqrt{w_0(w_0 - 4\alpha z_0^2)}}{2}$ Thus, opportunely bounding $w_{\tau-2}$ we obtain

$$\alpha(\tau_1 - 1)\Phi c_1^{3/2} < \mathbf{a}(0)^\top \mathbf{b}(0). \quad (46)$$

That we can reorganize as

$$\tau_1 < \frac{\mathbf{a}(0)^\top \mathbf{b}(0)}{\alpha c_1^{3/2}} (\Phi)^{-1} + 1. \quad (47)$$

Thus

$$\frac{1}{\alpha(w_0)^{3/2}} \leq \tau_1 \leq \frac{\mathbf{a}(0)^\top \mathbf{b}(0)}{\alpha c_1^{3/2}} (\Phi)^{-1} + 1. \quad (48)$$

□

Lemma 12 (Rate of convergence 1.). *Under Assumption 1, if $\sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)| \geq -2\mathbf{a}(0)^\top \mathbf{b}(0) > 0$, define $c_1 := \frac{w_0}{2} - \frac{\sqrt{w_0(w_0 - 4\alpha z_0^2)}}{2} > 2\sqrt{\alpha}\Phi$ as above, then*

$$c_1 < \sum_i |a_i^2(\tau_1) - b_i^2(\tau_1)| < \sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)| - \alpha^2 (\Phi)^2 \frac{(\mathbf{a}(0)^\top \mathbf{b}(0))^2}{\sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)|}, \quad (49)$$

$$\mathbf{a}(0)^\top \mathbf{b}(0) > 0, \quad (50)$$

and

$$\frac{1}{\alpha(\sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)|)^{3/2}} < \tau_1 < \frac{\mathbf{a}(0)^\top \mathbf{b}(0)}{\alpha c_1^{3/2}} (\Phi)^{-1} + 1. \quad (51)$$

Proof. One bound comes from Lemma 11 and Lemma 9. The other one comes by just following the proof of Lemma 11. □

Lemma 13 (Rate of convergence 2 - Sequence.). *Let $c_2 := |w_0/\mathbf{a}(0)^\top \mathbf{b}(0)|$. Under Assumption 1, if $w_0 < -2z_0 - 2\Phi$ then by time τ_2 as below we are under the assumptions of Lemma 11, so we have $w_{\tau_2} > -2z_{\tau_2} - 2\Phi$ and we have*

$$\frac{|c_2|}{2\alpha|z_0|} < \tau_2 < \frac{|c_2|}{\log(1 + 2\alpha\Phi)} + 1, \quad (52)$$

$$w_0(1 - \alpha^2 z_0^2) \exp\left(-\frac{\alpha|c_2|}{2} \frac{z_0^2 \sum x_i^2}{\sum x_i y_i}\right) < w_{\tau_2} < w_0 \exp\left(-\frac{\alpha|c_2|}{2} \frac{[\sum x_i y_i]^2}{z_0 [\sum x_i^2]^2}\right), \quad (53)$$

$$(z_0 + \Phi)(1 - 2\alpha|z_0|) \exp\left(-|c_2| \frac{|z_0| \sum x_i^2}{\sum x_i y_i} (1 - \alpha\Phi)^{-1}\right) - \Phi < z_{\tau_2}, \quad (54)$$

and

$$z_{\tau_2} < (z_0 + \Phi) \exp\left(-|c_2| \frac{\sum x_i y_i}{|z_0| \sum x_i^2}\right) - \Phi. \quad (55)$$

Proof. Note that the quantity $z_k + \Phi$ is shrinking exponentially as $\dot{X}_t = -2\alpha|z_k|X_t$ and w_k is shrinking exponentially as $\dot{X}_t = -\alpha^2 z_k^2 X_t$. We have that $\alpha < \frac{2}{|z_0|} < \frac{2}{|z_k|}$, thus the rate at which $z_k + \Phi$ is decreasing is faster than w_k , this implies that at a certain point we will have $w_k > -2z_k - 2\Phi$. Let us study this time τ_2 . Note that the fact that $|z_0| > |z_k| > \Phi$ implies that $2\alpha|z_k| > 2\alpha\Phi$ and $\alpha^2|z_k| > \alpha^2(\Phi)^2$

$$2\alpha|z_0| > 2\alpha|z_k| > 2 \frac{\sum x_i y_i}{|z_0| \sum x_i^2} \quad (56)$$

and

$$\alpha^2 |z_0|^2 > \alpha^2 |z_k|^2 > \frac{(\Phi)^2}{z_0^2}. \quad (57)$$

By the time $\tau_2 = \min_k \{w_0 \geq -2z_0 - 2\Phi, z_k > -\Phi\}$ we have the thesis. Note that definition of τ_2 and the iterative formulas

$$\begin{aligned} w_{\tau_1} &= w_0 \prod_{k=1}^{\tau_2} (1 - \alpha^2 z_k^2), \quad \text{and} \\ z_{\tau_1} + \Phi &= (z_0 + \Phi) \prod_{k=1}^{\tau_2} (1 - 2\alpha z_k) \end{aligned} \quad (58)$$

imply that

$$\prod_{k=0}^{\tau_2-1} (1 - \alpha^2 z_k^2) \left| \frac{w_0}{z_0 + \Phi} \right| = \left| \frac{w_{\tau_2}}{z_0 + \Phi} \right| > 2 \frac{z_{\tau_2} + \Phi}{z_0 + \Phi} = 2 \prod_{k=0}^{\tau_2-1} (1 - 2\alpha z_k). \quad (59)$$

This implies

$$\left| \frac{w_0}{z_0 + \Phi} \right| > 2 \prod_{k=0}^{\tau_2-1} \frac{1 - 2\alpha z_k}{1 - \alpha^2 z_k^2} > 2 \left(\frac{1 - 2\alpha z_0}{1 - \alpha^2 (\Phi)^2} \right)^{\tau_2} > 2 (1 - 2\alpha z_0)^{\tau_2}. \quad (60)$$

Analogously we have

$$\left| \frac{w_0}{z_0 + \Phi} \right| < 2 \prod_{k=0}^{\tau_2-2} \frac{1 - 2\alpha z_k}{1 - \alpha^2 z_k^2} = 2 \prod_{k=0}^{\tau_2-2} \frac{1 - 4\alpha^2 z_k^2}{1 - \alpha^2 z_k^2} \frac{1}{1 + 2\alpha |z_k|} < 2 \prod_{k=0}^{\tau_2-2} (1 + 2\alpha \Phi)^{-1}. \quad (61)$$

Thus defining $c_2 := \log \left(\frac{1}{2} \left| \frac{w_0}{z_0 + \Phi} \right| \right) < 0$ we obtain

$$\begin{aligned} -c_2 &< 2\alpha |z_0| \tau_2 \\ -c_2 &> \log(1 + 2\alpha \Phi) (\tau_2 - 1). \end{aligned} \quad (62)$$

Thus we conclude that

$$\begin{aligned} \frac{|c_2|}{2\alpha |z_0|} &< \tau_2 \\ \frac{|c_2|}{2\alpha \Phi} (1 - \alpha \Phi)^{-1} + 1 &> \frac{|c_2|}{\log(1 + 2\alpha \Phi)} + 1 > \tau_2 \end{aligned} \quad (63)$$

This implies that

$$(z_0 + \Phi) \exp \left(-|c_2| \frac{\sum x_i y_i}{|z_0| \sum x_i^2} \right) - \Phi > z_{\tau_2} \quad (64)$$

and

$$(z_0 + \Phi) (1 - 2\alpha |z_0|) \exp \left(-|c_2| \frac{|z_0| \sum x_i^2}{\sum x_i y_i} (1 - \alpha \Phi)^{-1} \right) - \Phi < z_{\tau_2}. \quad (65)$$

Analogously

$$w_0 (1 - \alpha^2 z_0^2)^{\frac{|c_2|}{\log(1 + 2\alpha \Phi)} + 1} < w_{\tau_2} < w_0 \left(1 - \alpha^2 (\Phi)^2 \right)^{\frac{|c_2|}{2\alpha |z_0|}}. \quad (66)$$

This means approximately that

$$w_0 (1 - \alpha^2 z_0^2) \exp \left(-\frac{\alpha |c_2|}{2} \frac{z_0^2 \sum x_i^2}{\sum x_i y_i} \right) < w_{\tau_2} < w_0 \exp \left(-\frac{\alpha |c_2|}{2} \frac{[\sum x_i y_i]^2}{z_0 [\sum x_i^2]^2} \right). \quad (67)$$

□

Lemma 14 (Rate of convergence 2.). *Let $c_2 := \left| \sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)| / \mathbf{a}(0)^\top \mathbf{b}(0) \right|$. Under Assumption 1, if $0 < \sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)| \leq -2\mathbf{a}(0)^\top \mathbf{b}(0)$, then by time τ_2 as below we are under the assumptions of Lemma 11, so we have $\sum_i |\mathbf{a}_i^2(\tau_2) - \mathbf{b}_i^2(\tau_2)| \geq -2\mathbf{a}(\tau_2)^\top \mathbf{b}(\tau_2) > 0$ and we have*

$$\begin{aligned}
& \bullet \frac{|c_2|}{2\alpha|\mathbf{a}(0)^\top \mathbf{b}(0) - \Phi|} < \tau_2, \\
& \bullet \tau_2 < \frac{|c_2|}{\log(1+2\alpha\Phi)} + 1, \\
& \bullet \sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)| (1 - \alpha^2 z_0^2) \exp\left(-\frac{\alpha|c_2|}{2} \frac{z_0^2 \sum x_i^2}{\sum x_i y_i}\right) < \sum_i |a_i^2(\tau_2) - b_i^2(\tau_2)|, \\
& \bullet \sum_i |a_i^2(\tau_2) - b_i^2(\tau_2)| < \sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)| \exp\left(-\frac{\alpha|c_2|}{2} \frac{[\sum x_i y_i]^2}{z_0 [\sum x_i^2]^2}\right), \\
& \bullet (z_0 + \Phi)(1 - 2\alpha|z_0|) \exp\left(-|c_2| \frac{|z_0| \sum x_i^2}{\sum x_i y_i} (1 - \alpha\Phi)^{-1}\right) < a(\tau_2)^\top b(\tau_2), \text{ and} \\
& \bullet a(\tau_2)^\top b(\tau_2) < (z_0 + \Phi) \exp\left(-|c_2| \frac{\sum x_i y_i}{|z_0| \sum x_i^2}\right).
\end{aligned}$$

Proof. One bound comes from Lemma 13 and Lemma 9. The other one comes by just following the proof of Lemma 13. \square

This shows that

Proposition 5. *Under Assumption 1, if $\mathbf{a}(0)^\top \mathbf{b}(0) < 0$ and $\Phi > 0$, there exists τ such that $a(\tau)^\top b(\tau) > 0$ and*

$$2\sqrt{\alpha}\Phi < Q(\tau) = O(\exp(-\alpha)) \cdot Q(0).$$

- If $\sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)| \geq 2|\mathbf{a}(0)^\top \mathbf{b}(0)|$ then $\tau \leq O(\alpha^{-1})$ and
- If $\sum_i |\mathbf{a}_i^2(0) - \mathbf{b}_i^2(0)| < 2|\mathbf{a}(0)^\top \mathbf{b}(0)|$ then $\tau \leq O\left(\frac{1}{\log(1+2\alpha)}\right)$.

Analogously if $\mathbf{a}(0)^\top \mathbf{b}(0) > 0$ and $\Phi < 0$, after the same τ , we have $\mathbf{a}(\tau)^\top \mathbf{b}(\tau) < 0$ and $\sum_i |\mathbf{a}_i^2(\tau) - \mathbf{b}_i^2(\tau)|$ is the same as above.

F CONVERGENCE SPEED CASE BY CASE

This section serves as merger for all the theory made before. Precisely, here we use the analysis developed to prove Theorem 2 and Proposition 1.

We prove below and in Appendix E.2 that in the three different regions of the landscape we have different PL constants μ for ε -and then for L . Precisely, if $\varepsilon \geq 0$ then $\mu > 2\Phi$, if $\varepsilon < -\Phi/2$ then $\mu = Q(\tau)$, and if $-\Phi/2 \leq \varepsilon < 0$ then $\mu = \Phi$. This implies that we have convergence with the minimum of $Q(\tau)$ and 2Φ as PL constant until $|\varepsilon| > \Phi/2$, then we have convergence with Φ as PL constant from then on.

F.1 POSITIVE RESIDUALS

First note that $\eta\lambda\varepsilon > \eta^2\varepsilon^2(\varepsilon + \Phi)$, indeed $\lambda > 2(\varepsilon + \Phi)$ by Cauchy Schwartz and $\eta \leq \frac{\sqrt{2}}{\varepsilon} < \frac{2}{\varepsilon}$. This implies that when η is infinitesimal, the gain is at least

$$\begin{aligned}
\varepsilon(k+1) &= (1 - \eta\lambda)\varepsilon(k) + \eta^2\varepsilon(k)(\varepsilon + \Phi) \\
&\leq \varepsilon(k) - \eta\lambda\varepsilon(k) \left(1 - \frac{\eta}{2}\varepsilon(k)\right) \\
&\leq \left(1 - \frac{2 - \sqrt{2}}{2}\eta\lambda\right)\varepsilon(k).
\end{aligned} \tag{68}$$

Next note that $\frac{x}{\sqrt{x^2+y^2}} = \sqrt{1 - \frac{y^2}{x^2+y^2}} \leq 1 - \frac{1}{2} \frac{y^2}{x^2+y^2}$. When $\eta \sim \frac{2}{\lambda}(1 - \delta)$, $\delta > 0$ we have that

$$\begin{aligned}
 |\varepsilon(k+1)| &= |(1 - \eta\lambda)\varepsilon(k) + \eta^2\varepsilon^2(k)(\varepsilon(k) + \Phi)| \\
 &\leq \left| \varepsilon(k) - \frac{2\lambda(1 - \delta)}{\sqrt{\lambda^2 + 4\Phi^2}}\varepsilon(k) + \frac{4(1 - \delta)^2}{\lambda^2 + 4\Phi^2}\varepsilon(k)^2(\varepsilon(k) + \Phi) \right| \\
 &\leq \left| (-1 + 2\delta)\varepsilon(k) + \frac{4\Phi^2(1 - \delta)}{\lambda^2 + 4\Phi^2}\varepsilon(k) + \frac{4(1 - \delta)}{\lambda^2 + 4\Phi^2}\varepsilon(k)^2(\varepsilon(k) + \Phi) \right| \\
 &\leq (1 - 2\delta)\varepsilon(k).
 \end{aligned} \tag{69}$$

This implies that within our learning rate boundaries we have exponential convergence with rate either controlled by η or δ at power 1.

In case $\frac{2}{\lambda} \leq \eta \leq \frac{2}{\lambda}(1 - \delta)$ then convergence happens exponentially but in time $O(\eta^{-2})$. For instance $\eta \sim \frac{2}{\lambda}$ we have that

$$\begin{aligned}
 |\varepsilon(k+1)| &= |(1 - \eta\lambda)\varepsilon(k) + \eta^2\varepsilon^2(k)(\varepsilon(k) + \Phi)| \\
 &\leq \left| \varepsilon(k) - \frac{2\lambda}{\sqrt{\lambda^2 + 4\Phi^2}}\varepsilon(k) + \frac{4}{\lambda^2 + 4\Phi^2}\varepsilon(k)^2(\varepsilon(k) + \Phi) \right| \\
 &\leq \left| -\varepsilon(k) + \frac{4\Phi^2}{\lambda^2 + 4\Phi^2}\varepsilon(k) + \frac{4}{\lambda^2 + 4\Phi^2}\varepsilon(k)^2(\varepsilon(k) + \Phi) \right| \\
 &\leq (1 - \eta^2\Phi^2)\varepsilon(k).
 \end{aligned} \tag{70}$$

In case $\eta \geq \frac{2}{\lambda}$ but convergence happen, then convergence happens only at most logarithmically fast at least for a first phase, precisely

$$\begin{aligned}
 |\varepsilon(k+1)| &= |(1 - \eta\lambda)\varepsilon(k) + \eta^2\varepsilon^2(k)(\varepsilon(k) + \Phi)| \\
 &\leq \left| -\varepsilon(k) + \frac{4}{\lambda^2}\varepsilon(k)^2(\varepsilon(k) + \Phi) \right| \\
 &\leq \left(1 - \frac{4}{\lambda^2}\Phi\varepsilon(k) \right) \varepsilon(k).
 \end{aligned} \tag{71}$$

F.2 NEGATIVE RESIDUALS $-\Phi/2 < \varepsilon < 0$

When the residuals are small negative we have exponential convergence, precisely, for very small $\eta \ll 1$ we have rate at least $(1 - \eta\Phi)$:

$$\begin{aligned}
 |\varepsilon(k+1)| &= |(1 - \eta\lambda)\varepsilon(k) + \eta^2\varepsilon^2(k)(\varepsilon(k) + \Phi)| \\
 &\leq \left| (1 - \eta\lambda) + \frac{\eta^2}{4}\Phi^2 \right| |\varepsilon(k)| \\
 &\leq (1 - \eta\Phi) |\varepsilon(k)|.
 \end{aligned} \tag{72}$$

For bigger $\eta = \frac{2}{\lambda}$, we have convergence with rate about ~ 2 . The maximum over $\lambda(0), \varepsilon(k)$ in the region in which $\varepsilon(k) = -c\frac{\Phi}{2}$ with $c \in (0, 1]$

$$\begin{aligned}
 \max |\varepsilon(k+1)| &= \max |(1 - \eta\lambda)\varepsilon(k) + \eta^2\varepsilon^2(k)(\varepsilon(k) + \Phi)| \\
 &\leq \max \left| \varepsilon(k) - \frac{2\lambda}{\sqrt{\lambda(0)^2 + 4\Phi^2}}\varepsilon(k) + \frac{4}{\lambda(0)^2 + 4\Phi^2}\varepsilon(k)^2(\varepsilon(k) + \Phi) \right|.
 \end{aligned} \tag{73}$$

Note that the minimum in $\lambda(0)$ of this last equation is for $\sqrt{\lambda(0)^2 + 4\Phi^2} = \delta + 2\Phi$ for some $\delta > 0$ which satisfies $\delta \ll 1$. This is independent of the size of ε . Along this trajectory, $\mathbf{a}^\top \mathbf{b} = (1 - c/2)\Phi$

and $\lambda \leq (2 - c)\Phi + \delta$ This implies

$$\begin{aligned}
\max_{\lambda(0), \varepsilon(k)} |\varepsilon(k+1)| &\leq \max_{\varepsilon(k)} \left| \varepsilon(k) - \frac{2(2-c)\Phi}{\delta+2\Phi} \varepsilon(k) + \frac{4}{(\delta+2\Phi)^2} \varepsilon(k)^2 (\varepsilon(k) + \Phi) \right| \\
&\leq \max_{\varepsilon(k)} \left| -\frac{(2-2c)\Phi + \delta}{\delta+2\Phi} + \frac{c(2-c)}{(\delta+2\Phi)^2} \Phi^2 \right| |\varepsilon(k)| \\
&\leq \left| -1 + c \frac{2\Phi}{\delta+2\Phi} + c(2-c) \frac{\Phi^2}{(\delta+2\Phi)^2} \right| |\varepsilon(k)| \leq \left| c - 1 - \frac{c^2}{4} + \frac{c}{2} \right| |\varepsilon(k)| \\
&\stackrel{c=1}{\leq} \frac{1}{4} |\varepsilon(k)| = \frac{1}{8} \Phi.
\end{aligned} \tag{74}$$

The maximum of $|c^2/4 - 3c/2 + 1|$ over $c \in (0, 1]$ is $c = 1$.

In the case of $c = 1$, on the next step, in this case, we are in the positive residuals setting with λ as follows $2 \cdot \mathbf{a}^\top \mathbf{b} = \frac{9}{4} \Phi + \delta$. Here, then

$$\begin{aligned}
|\varepsilon(k+2)| &\leq \left| -\frac{5}{4} + \frac{1}{4\Phi^2} \varepsilon(k+1)(\varepsilon(k+1) + \Phi) \right| |\varepsilon(k+1)| \\
&\leq \frac{1}{16} \left(5 - \frac{1}{16} \right) |\varepsilon(k)|.
\end{aligned} \tag{75}$$

So after 2 steps, we had a linear shrink of $5/16$ and the linear convergence with constant $\mu = \Phi$ restarts, this is the plus 2 of the theorem.

F.3 NEGATIVE RESIDUALS $\varepsilon \leq \Phi/2$

This case is taken care of in Appendix E.2 until $\varepsilon = 0$. With the same $\mu > 0$ we have exponential convergence until $\Phi/2$. As we said in Appendix E.2 as ε crosses Φ , the norm λ restarts increasing. This implies that a good lower bound remains Q of the time of crossing. The evolution of ε

$$\varepsilon(k+1) = (1 - \eta\lambda)\varepsilon(k) + \eta^2 \varepsilon(k)^2 (\varepsilon(k) + \Phi) \geq (1 - \eta Q)\varepsilon(k). \tag{76}$$

The time t taken to ε to go from Φ to $\Phi/2$ is thus

$$\Phi/2 \geq (1 - \eta Q)^t \Phi \tag{77}$$

so we have

$$t \leq \frac{\log(\Phi) - \log(\Phi/2)}{-\log(1 - \eta Q)} \leq \frac{\log(\Phi) - \log(\Phi/2)}{\eta Q_\tau}. \tag{78}$$

F.4 CLOSING UP: TIGHT RATE

The previous sections and Lemma 5 allow us to conclude that we have loss convergence, i.e., $L \leq \delta$, in a number of steps which is

$$t \leq \tau + \frac{\log(\Phi) - \log(\Phi/2)}{\eta \min\{Q_\tau, 2\Phi\}} + 2 + \frac{\log(\Phi/2) - \log(\delta)}{\eta \min\{Q_\tau, \Phi\}}, \tag{79}$$

where τ is the τ_1 defined in Definition 4 and evaluated in Proposition 5. This establishes Theorem 2.

G CURIOSITY: JUMPS BETWEEN REGIONS

Note that if the dynamics does not jump from one side to the other of the landscape, then we have a clean exponential convergence and we can control the implicit regularization. We will see under which hypothesis on the learning rate this happens.

Note that Equation 5 tells us that after one step ε does not change sign (thus you remain in the same region in which you started) if and only if we have the following bound on the learning rate.

Definition 5. For all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, let $\alpha = \frac{\varepsilon(\varepsilon + \Phi)}{\lambda^3}$, define

$$\eta_1 := \frac{1}{\lambda} (1 + \alpha + 2\alpha^2 + 5\alpha^3 + 14\alpha^8 + \dots), \quad (80)$$

$$\eta_2 := \frac{2}{\lambda} (1 + 2\alpha + 8\alpha^2 + 40\alpha^3 + 224\alpha^4 + \dots). \quad (81)$$

The way we obtain η_1 is by seeing for what η we have that $\varepsilon(t+1) = 0$. Precisely,

Lemma 15. *If $\eta = \eta_1$, we have that the residuals at the next steps are 0. If $\eta = \eta_2$, then the residuals at the next steps are the same but changed of sign. Moreover,*

- If $\eta \in (0, \eta_1)$ then $\text{sign}(\varepsilon(1)) = \text{sign}(\varepsilon)$ and $|\varepsilon(1)| < |\varepsilon|$.
- If $\eta \in (\eta_1, \eta_2)$ then $\text{sign}(\varepsilon(1)) \neq \text{sign}(\varepsilon)$ and $|\varepsilon(1)| < |\varepsilon|$.

Proof of Lemma 15. Note that the residuals after one step are the same sign as the previous residuals if and only if

$$1 - \eta\lambda + \eta^2\varepsilon(\varepsilon + \Phi) \geq 0. \quad (82)$$

Solving this one as a second degree equation gives

$$\eta \leq \frac{\lambda - \sqrt{\lambda^2 - 4\varepsilon(\varepsilon + \Phi)}}{2\varepsilon(\varepsilon + \Phi)} \quad \text{or} \quad \eta \geq \frac{\lambda + \sqrt{\lambda^2 - 4\varepsilon(\varepsilon + \Phi)}}{2\varepsilon(\varepsilon + \Phi)} \quad (83)$$

Now expanding in Taylor the square root, we obtain that

$$\eta_1 \leq \frac{1}{2\varepsilon(\varepsilon + \Phi)} \left(\frac{4\varepsilon(\varepsilon + \Phi)}{2\lambda} + \frac{16\varepsilon^2(\varepsilon + \Phi)^2}{8\lambda^3} + \dots \right) = \frac{1}{\lambda} + \frac{\varepsilon(\varepsilon + \Phi)}{\lambda^3} + \dots \quad (84)$$

This implies that the residuals are the same sign as the starting ones if

$$\eta \leq \eta_1 \quad \text{or} \quad \eta \geq \frac{\lambda}{\varepsilon(\varepsilon + \Phi)} - \eta_1. \quad (85)$$

Analogously, for η_2 we have that the absolute value of the residuals is smaller than the absolute value of the residuals one step before, if and only if

$$2 - \eta\lambda + \eta^2\varepsilon(\varepsilon + \Phi) \geq 0. \quad (86)$$

This implies that

$$\eta \leq \frac{\lambda - \sqrt{\lambda^2 - 8\varepsilon(\varepsilon + \Phi)}}{2\varepsilon(\varepsilon + \Phi)} \quad \text{or} \quad \eta \geq \frac{\lambda + \sqrt{\lambda^2 - 8\varepsilon(\varepsilon + \Phi)}}{2\varepsilon(\varepsilon + \Phi)} \quad (87)$$

and analogously to before

$$\eta \leq \eta_2 \quad \text{or} \quad \eta \geq \frac{\lambda}{\varepsilon(\varepsilon + \Phi)} - \eta_2. \quad (88)$$

Also note that for $\varepsilon > 0$ we have that $\varepsilon(1) < \varepsilon$ or for $\varepsilon < 0$ we have that $\varepsilon(1) > \varepsilon$ if and only if

$$1 - \eta\lambda + \eta^2\varepsilon(\varepsilon + \Phi) \leq 1. \quad (89)$$

This solves when

$$\eta \leq \frac{\lambda}{\varepsilon(\varepsilon + \Phi)} \quad (90)$$

□

Note that what we did here implies that if $\eta \leq \eta_2$ and $\eta \leq \frac{\sqrt{2}}{\varepsilon}$ for all the λ s along the trajectory we thus always have exponential convergence if such PL condition holds. We know from the previous section that in this setting λ is always smaller than $\bar{\lambda}$. So if such a μ exists and $\eta \leq \eta_2$ with $\bar{\lambda}$ and we converge and we can properly bound the implicit regularization.

H LOCATION OF CONVERGENCE

We will bound here the final sharpness for two reasons:

- Understanding the location of convergence.
- Picking the right learning rate.

Note that assuming $\eta \leq \frac{\sqrt{2}}{\varepsilon}$ along the whole trajectory we have that $|Q|$ strictly monotonically shrinks along the trajectory. This means that the dynamics may seem to oscillate around in an uncontrollable way, but every time it oscillates is landing on a trajectory that takes to a global minimum with lower sharpness.

Note that this is true almost everywhere, indeed if the trajectory is such that at a certain point in time t satisfies $\eta = \varepsilon(t)^{-1}$ exactly, then the trajectory would land on the trajectory taking to the saddle, indeed

$$\mathbf{a}(t+1) = -\mathbf{b}(t+1) = \mathbf{a}(t) - \mathbf{b}(t). \quad (91)$$

Luckily, fixing a learning rate size, the set of starting points for which this is the case has measure zero. Observe also that $\eta = -\varepsilon^{-1}$ is instead optimal and results in $\mathbf{a}(t+1) = \mathbf{b}(t+1)$, implying convergence to the lowest norm solution.

This means that assuming $\eta \leq \frac{\sqrt{2}}{\varepsilon}$ implies that the dynamics may diverge or converge, but for sure at every step is getting closer and closer to the line $\mathbf{a} = \mathbf{b}$.

Note that picking any learning rate η such that $|\varepsilon(t+1)| \leq |\varepsilon(t)|$ and $\eta \leq \frac{\sqrt{2}}{\varepsilon}$ is thus sufficient for the convergence. It is not necessary though, it may be that the absolute value of the residuals grows and then descends.

Regarding the proof of the upperbound of Theorem 1 note that for all k

$$Q_\infty = Q_0 \cdot \prod_{k=1}^{\infty} (1 - \eta^2 \varepsilon_k^2) = Q_0 \exp \left(\sum_{k=1}^{\infty} \log(1 - \eta^2 \varepsilon_k^2) \right) < Q_0 \exp \left(-\eta^2 \sum_{k=1}^{\infty} \varepsilon_k^2 \right). \quad (92)$$

Regarding the proof of the lower bound of Theorem 1, notice that we have from Appendix D.2 that the rate of convergence of ε is at least $Q(\tau_1)$ in Area B and at least 2Φ in Area A. This implies that if the initialization is in Area B or C, then

$$\begin{aligned} Q_\infty &= Q(0) \cdot \prod_{k=0}^{\infty} (1 - \eta^2 \varepsilon_k^2) \geq Q(0) \cdot \exp \left(-\sum_1^{\infty} \eta^2 \varepsilon_k^2 - \eta^4 \varepsilon_k^4 \right) \\ &\geq Q(0) \cdot \exp \left(-\eta^2 \varepsilon(0)^2 \sum_1^{\infty} (1 - \eta Q(\tau))^{2k} - \eta^4 \varepsilon_0^4 \sum_1^{\infty} (1 - \eta Q(\tau))^{4k} \right) \\ &\geq Q(0) \cdot \exp \left(-\frac{\eta \varepsilon(0)^2}{Q(\tau)(2 - \eta Q(\tau))} - \frac{\eta^3 \varepsilon(0)^4}{Q(\tau)(8 - \eta Q(\tau))} \right) \\ &\geq Q(0) \cdot \exp \left(-\frac{\sqrt{\eta} \varepsilon(0)^2}{2\Phi} \left(1 + \frac{\eta^2 \varepsilon(0)^2}{8} \right) \right) \\ &\geq Q(0) \cdot \exp \left(-\frac{\sqrt{\eta} \varepsilon(0)^2}{\Phi} \right). \end{aligned} \quad (93)$$

By plugging in the lower bound in Lemma 12.