

# Flavors of Margin: Implicit Bias of Steepest Descent in Homogeneous Neural Networks

**Nikolaos Tsilivis**

*Center for Data Science  
New York University  
NYC, NY 10011, USA*

NT2231@NYU.EDU

**Eitan Gronich**

*Department of Computer Science and Applied Mathematics  
Weizmann Institute of Science  
Rehovot 7610001, Israel*

EITAN.GRONICH@WEIZMANN.AC.IL

**Julia Kempe**

*Courant Institute and Center for Data Science  
New York University  
NYC, NY 10011, USA  
Meta FAIR Paris, France*

KEMPE@META.COM

**Gal Vardi**

*Department of Computer Science and Applied Mathematics  
Weizmann Institute of Science  
Rehovot 7610001, Israel*

GAL.VARDI@WEIZMANN.AC.IL

**Editor:** Benjamin Guedj

## Abstract

We study the implicit bias of the general family of steepest descent algorithms with infinitesimal learning rate in deep homogeneous neural networks. We show that: (a) an algorithm-dependent geometric margin starts increasing once the networks reach perfect training accuracy, and (b) any limit point of the training trajectory corresponds to a KKT point of the corresponding margin-maximization problem. We experimentally zoom into the trajectories of neural networks optimized with various steepest descent algorithms, highlighting connections to the implicit bias of popular adaptive methods (Adam and Shampoo).<sup>1</sup>

**Keywords:** Implicit bias of optimization, Homogeneous neural networks, Steepest descent

## 1. Introduction

Overparameterized neural networks excel in many natural supervised learning applications. A theory that aims to explain their strong generalization performance places optimization at the forefront: in problems where many candidate models are available, the optimization algorithm implicitly selects well-generalizing ones (Neyshabur et al., 2015b). The term “implicitly” indicates that the loss/objective function does not explicitly favor simple, well-generalizing solutions, rather this occurs due to the properties of the optimization algorithm.

---

1. The present paper is an extended version of an earlier conference paper (Tsilivis et al., 2025).

Most existing theoretical results on this *implicit bias of optimization* demonstrate, to some extent, that gradient descent in overparameterized problems biases the solution to be the *simplest*, in terms of the lowest possible  $\ell_2$  norm of the weights (Soudry et al., 2018; Ji and Telgarsky, 2019, 2020; Lyu and Li, 2020; Nacson et al., 2019).

Simplicity, however, is a term that depends on the setting. For instance, in logistic regression with many irrelevant features, an  $\ell_1$ -regularized solution is simpler than an  $\ell_2$ -regularized one (Ng, 2004). Moreover, in contemporary deep learning, Adam (Kingma and Ba, 2015), AdamW (Loshchilov and Hutter, 2019), and related optimization algorithms are preferred for language modeling (Zhang et al., 2020), and their implicit bias might be better suited for such applications than gradient descent. It is therefore important to understand the types of solutions favored by optimization algorithms beyond (stochastic) gradient descent, in order to address the current (and future) applications of deep learning.

In this work, we contribute to this line of research by studying the large family of *steepest descent* algorithms (Equation 1) with respect to an arbitrary norm  $\|\cdot\|$  in deep, non-linear, homogeneous neural networks. This class of methods extends gradient descent to optimization geometries other than the Euclidean, allowing the update rule to operate under a different norm. It includes coordinate descent (which has strong ties to boosting (Mason et al., 1999)), sign gradient descent (which is closely related to Adam (Kunstner et al., 2023)) and spectral steepest descent (which is similar to Shampoo (Gupta et al., 2018)) as special cases.

**Our contributions.** We provide a unifying, rigorous analysis of any steepest descent algorithm in classification settings with locally-Lipschitz, homogeneous neural networks trained using an exponentially-tailed loss. Specifically, we focus on the late stage of training (after the network has achieved perfect training accuracy) in the limit of an infinitesimal learning rate.

Our first result characterizes the algorithm’s tendency to increase an algorithm-dependent margin (Theorem 1): similar to prior work on gradient descent (Lyu and Li, 2020), we show that a *soft* version of the geometric margin starts increasing immediately after fitting the training data. We then turn our attention to the asymptotic properties of the algorithm. As we show, the limit points of training are along the direction of a Karush-Kuhn-Tucker (*KKT*) point of the algorithm-dependent margin maximization problem (Theorem 4).

In total, these results establish (geometric) margin-maximization in *any* steepest descent algorithm and significantly generalize prior results that concerned gradient descent only (Lyu and Li, 2020; Nacson et al., 2019). They also generalize a previous version of the current paper, which only established bias towards *KKT* points in the case of algorithms whose squared norm is a smooth function (Tsilivis et al., 2025). See the end of Section 3.2 for a technical discussion.

Finally, in Section 4, we train neural networks with the three main steepest descent algorithms (gradient descent, sign gradient descent and coordinate descent). We perform experiments in: (a) teacher-student tasks, to elucidate the theoretical findings and assess the connection between implicit bias and generalization and (b) image classification tasks, to study the relationship between Adam (Kingma and Ba, 2015), Shampoo (Gupta et al., 2018) and steepest descent algorithms.

## 1.1 Related Work

There have been numerous works studying the implicit bias of optimization in supervised learning and their relationship to geometric margin maximization - see Vardi (2023) for a survey.

Steepest descent algorithms with respect to non-Euclidean geometries have been explored before, both in supervised (e.g. Neyshabur et al. (2015a); Large et al. (2024)) and non-supervised (e.g. Carlson et al. (2015)) machine learning problems. The implicit bias of this family of optimization methods was first studied in generality in Gunasekar et al. (2018) in the context of linear models for separable data, where margin maximization was established. Their proof is based on a result on Adaboost due to Telgarsky (2013). Our results generalize the analysis of steepest descent algorithms to any homogeneous neural network. Most related to our paper are the works of Nacson et al. (2019); Lyu and Li (2020) and Ji and Telgarsky (2020). Nacson et al. (2019) studied infinitesimal regularization and its connection to margin maximization in both homogeneous and non-homogeneous deep models, while also proving (directional) convergence of gradient descent to a first order point of an  $\ell_2$ -margin maximization problem for homogeneous models under strong technical assumptions. Lyu and Li (2020); Ji and Telgarsky (2020), whose theoretical setup we mainly follow, significantly weakened the assumptions, under which such a result holds, and Lyu and Li (2020) further demonstrated the experimental benefits of margin maximization in terms of robustness. Kunin et al. (2023) generalized these results to a broader class of networks with varying degree of homogeneity, while Cai et al. (2024) analyzed non-homogeneous 2-layer networks trained with a large learning rate. Vardi et al. (2022) identified cases where the KKT points of the  $\ell_2$  margin maximization problems are not (even locally) optimal.

The implicit bias of Adam (Kingma and Ba, 2015) has been previously studied in Wang et al. (2021, 2022) for homogeneous networks, where it is proven that it shares the same asymptotic properties as gradient descent ( $\ell_2$  margin maximization). Recently, Zhang et al. (2024) analyzed a version of the algorithm in linear models, without a numerical precision constant, which arguably better captures realistic training runs, and found bias towards  $\ell_1$  margin maximization - the same bias as in the case of sign gradient descent. This makes us optimistic that insights from our analysis, which is connected to sign gradient descent (normalized steepest descent with respect to the  $\ell_\infty$  norm), can shed light on the poorly understood implicit bias of Adam in deep neural networks. See also Xie and Li (2024) for a recently established connection between AdamW (Loshchilov and Hutter, 2019) and sign gradient descent. An additional motivation for studying steepest descent algorithms is in improving the robustness of deep neural networks: Tsilivis et al. (2024), recently, provided experimental evidence and theoretical arguments that adversarially trained deep networks exhibit significant differences in their (robust) generalization error, depending on which steepest descent algorithm was used in training.

## 2. Background

**Learning Setup.** We consider binary classification problems with deep, homogeneous, neural networks. Formally, let  $\mathcal{S} = \{\mathbf{x}_i, y_i\}_{i=1}^m$  be a dataset of i.i.d. points sampled from an unknown distribution  $\mathcal{D}$  with  $\mathbf{x}_i \in \mathbb{R}^d$  and  $y_i \in \{\pm 1\}$  for all  $i \in [m]$ , and let  $f(\cdot; \boldsymbol{\theta}) : \mathbb{R}^d \rightarrow \mathbb{R}$  denote a neural network parameterized by  $\boldsymbol{\theta} \in \mathbb{R}^p$ . The vector  $\boldsymbol{\theta}$  contains all the parameters

of the neural network, concatenated into a single vector. We study training under an exponential loss  $\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta})}$ . We focus on this setting for simplicity, but our results generalize to more common losses, such as the logistic loss. See Section A.3 for details and extensions.

**Algorithms.** The family of steepest descent algorithms generalizes gradient descent to different optimization geometries, allowing the update rule to operate under an arbitrary norm (instead of the usual Euclidean one) (Boyd and Vandenberghe, 2014). Formally, the update rule for *steepest descent* with respect to a norm  $\|\cdot\|$  is:

$$\begin{aligned} \boldsymbol{\theta}_{t+1} &= \boldsymbol{\theta}_t + \eta_t \Delta \boldsymbol{\theta}_t, \text{ where } \Delta \boldsymbol{\theta}_t \text{ satisfies} \\ \Delta \boldsymbol{\theta}_t &= \underset{\|\mathbf{u}\| \leq \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|_*}{\operatorname{argmin}} \langle \mathbf{u}, \nabla \mathcal{L}(\boldsymbol{\theta}_t) \rangle, \end{aligned} \tag{1}$$

where the *dual* norm  $\|\cdot\|_*$  of  $\|\cdot\|$  is defined as  $\|\mathbf{z}\|_* = \max_{\mathbf{v}} \{|\langle \mathbf{z}, \mathbf{v} \rangle| : \|\mathbf{v}\| = 1\}$  for any  $\mathbf{z}$ , and  $\eta_t$  is a learning rate. Gradient descent can be derived from Equation 1 with  $\|\cdot\| = \|\cdot\|_2$ . See Appendix C for details on how steepest descent algorithms are closely related to popular adaptive methods, such as Adam (Kingma and Ba, 2015) and Shampoo (Gupta et al., 2018).

**Assumptions & Technical Points.** In order to formally allow for commonly used activation functions, such as the ReLU, we theoretically analyze loss landscapes that are not necessarily differentiable. That is, we consider Clarke’s subdifferentials (Clarke, 1975) in our analysis:

$$\partial f := \operatorname{conv} \left\{ \lim_{k \rightarrow \infty} \nabla f(\mathbf{x}_k) : \mathbf{x}_k \rightarrow \mathbf{x}, f \text{ differentiable at } \mathbf{x}_k \right\}, \tag{2}$$

where  $\operatorname{conv}(\cdot)$  stands for the convex hull of a set.

Furthermore, we analyze steepest descent in the limit of infinitesimal step size, i.e. *steepest flow*:

$$\frac{d\boldsymbol{\theta}}{dt} \in \left\{ \underset{\|\mathbf{u}\| \leq \|\mathbf{g}_t\|_*}{\operatorname{argmin}} \langle \mathbf{u}, \mathbf{g}_t \rangle : \mathbf{g}_t \in \partial \mathcal{L}(\boldsymbol{\theta}_t) \right\}. \tag{3}$$

This choice simplifies the analysis while still capturing the essence of the bias of the algorithms. Finally, we make the following assumptions, which are common in the study of the implicit bias of deep learning methods (e.g. Lyu and Li (2020); Ji and Telgarsky (2020)):

- (A1) Local Lipschitzness: For any  $\mathbf{x}_i \in \mathbb{R}^d$ ,  $f(\mathbf{x}_i; \cdot) : \mathbb{R}^p \rightarrow \mathbb{R}$  is locally Lipschitz (and admits a chain rule - see Theorem 6).
- (A2)  $L$ -Homogeneity: We assume that  $f$  is  $L$ -homogeneous in the parameters, i.e.  $f(\cdot; c\boldsymbol{\theta}) = c^L f(\cdot; \boldsymbol{\theta})$  for any  $c > 0$  and  $\boldsymbol{\theta}$ .
- (A3) Realizability: There is a  $t_0 > 0$ , such that  $\mathcal{L}(\boldsymbol{\theta}_{t_0}) < 1$ .

Assumption (A1) is a minimal assumption on the regularity of the network, while assumption (A2) includes many commonly used architectures. For instance, ReLU networks with an arbitrary number of layers, but without bias terms and skip connections, satisfy (A1),(A2). Homogeneity simplifies the analysis, as Euler’s theorem for homogeneous functions directly links the gradient of the network at a point to its function value at that point. Assumption

(A3) ensures that the algorithm will succeed in classifying the training points and allows us to focus on what happens beyond that point of separation. Indeed, we are particularly interested in understanding the geometric properties of the model  $f(\cdot; \boldsymbol{\theta}_t)$  as  $t \rightarrow \infty$  (at convergence) – the *implicit bias* of the learning algorithms.

### 3. Theory

We analyze the behavior of steepest descent algorithms in the late stage of training and study their geometric properties and how these relate to geometric, algorithm-specific, margins.

#### 3.1 Algorithm-Dependent Margin Increases

In *linear* models, where  $f(\mathbf{x}; \boldsymbol{\theta}) = \langle \boldsymbol{\theta}, \mathbf{x} \rangle$ , the concept of  $\|\cdot\|_*$ -geometric margin<sup>2</sup>,  $\min_{i \in [m]} \frac{y_i \langle \boldsymbol{\theta}, \mathbf{x}_i \rangle}{\|\boldsymbol{\theta}\|}$ , plays a central and fundamental role in the analysis of the convergence of training (Novikoff, 1963) as well as in the generalization of the final model (Vapnik, 1998). Ideally, we would like to track a similar quantity when training general, homogeneous, non-linear networks  $f(\cdot; \boldsymbol{\theta})$  with steepest descent with respect to the  $\|\cdot\|$  norm:

$$\gamma(\boldsymbol{\theta}) = \frac{\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta})}{\|\boldsymbol{\theta}\|^L} = \min_{i \in [m]} y_i f\left(\mathbf{x}_i; \frac{\boldsymbol{\theta}}{\|\boldsymbol{\theta}\|}\right), \tag{4}$$

where recall that  $L$  is the level of homogeneity of the model. As it turns out, it is easier to follow the evolution of the following, *soft*, geometric margin:

$$\tilde{\gamma}(\boldsymbol{\theta}) = -\frac{\log \mathcal{L}(\boldsymbol{\theta})}{\|\boldsymbol{\theta}\|^L}. \tag{5}$$

The characterisation of “soft” comes from the definition of “softmax” (a.k.a. log-sum-exp), which is often used in machine learning. The same idea is used here to approximate the numerator of Equation 4. The soft margin  $\tilde{\gamma}(\boldsymbol{\theta})$  is at most an additive  $\log m$  away from  $\gamma(\boldsymbol{\theta})$  and converges to  $\gamma(\boldsymbol{\theta})$  as  $t \rightarrow \infty$  - see Lemma 11 and Corollary 12.

We show next that, given the algorithm has reached a small value in the loss, the soft margin is non-decreasing. This theorem is similar to part of Lemma 5.1 in (Lyu and Li, 2020), which is the key lemma in their result. Our proof is admittedly simpler, avoiding a beautiful polar decomposition which was crucial in their analysis, yet, unfortunately, pertinent to the  $\ell_2$  case only.

**Theorem 1 (Soft margin increases)** *For almost any  $t > t_0$ , it holds:*

$$\frac{d \log \tilde{\gamma}}{dt} \geq L \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{1}{L\mathcal{L}(\boldsymbol{\theta}_t) \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)}} - \frac{1}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \geq 0.$$

---

2. In this paper, we diverge from the established terminology when it comes to naming margins, by calling it  $\|\cdot\|_*$ -geometric margin (instead of  $\|\cdot\|$ -geometric margin) when it is defined with respect to the  $\|\cdot\|$  norm of the parameters. We believe this is proper, since the  $\|\cdot\|_*$ -geometric margin in linear models maximizes the metric induced by the  $\|\cdot\|_*$  norm (and not its dual,  $\|\cdot\|$ ).

**Proof** [Simplified version] We present a proof for a simplified version of this theorem here, covering differentiable networks  $f$ , while we defer the full proof to Appendix A.2. For differentiable losses, steepest flow corresponds to:

$$\frac{d\boldsymbol{\theta}}{dt} \in \operatorname{argmin}_{\|\mathbf{u}\| \leq \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|_*} \langle \mathbf{u}, \nabla \mathcal{L}(\boldsymbol{\theta}_t) \rangle. \quad (6)$$

By the definition of the dual norm and chain rule, we have for any  $t > 0$ :

$$\left\| \frac{d\boldsymbol{\theta}}{dt} \right\| = \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|_* \quad \text{and} \quad \frac{d\mathcal{L}(\boldsymbol{\theta}_t)}{dt} = - \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2. \quad (7)$$

Let  $\mathbf{n}_t \in \partial \|\boldsymbol{\theta}_t\|$  (recall that a norm  $\|\cdot\|$  might not be differentiable everywhere). For any  $t > t_0$ , we have:

$$\begin{aligned} \frac{d \log \tilde{\gamma}}{dt} &= \frac{d}{dt} \log \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} - L \frac{d}{dt} \log \|\boldsymbol{\theta}_t\| \\ &= \frac{d}{dt} \log \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} - L \left\langle \frac{\mathbf{n}_t}{\|\boldsymbol{\theta}_t\|}, \frac{d\boldsymbol{\theta}}{dt} \right\rangle \quad (\text{Chain rule}) \\ &\geq \frac{d}{dt} \log \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} - L \frac{\left\| \frac{d\boldsymbol{\theta}}{dt} \right\|}{\|\boldsymbol{\theta}_t\|} \quad (\text{def. of dual norm and } \|\mathbf{n}_t\|_* \leq 1, \text{ Lemma 8}) \quad (8) \\ &= - \frac{d\mathcal{L}(\boldsymbol{\theta}_t)}{dt} \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t) \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)}} - L \frac{\left\| \frac{d\boldsymbol{\theta}}{dt} \right\|}{\|\boldsymbol{\theta}_t\|} \quad (\text{Chain rule}) \\ &= \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t) \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)}} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right). \quad (\text{Equation 7}) \end{aligned}$$

The first term inside the parenthesis can be related to the second one via the following calculation:

$$\begin{aligned} \langle \boldsymbol{\theta}_t, -\nabla \mathcal{L}(\boldsymbol{\theta}_t) \rangle &= \left\langle \boldsymbol{\theta}_t, \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} y_i \nabla f(\mathbf{x}_i; \boldsymbol{\theta}_t) \right\rangle \\ &= \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} y_i \langle \boldsymbol{\theta}_t, \nabla f(\mathbf{x}_i; \boldsymbol{\theta}_t) \rangle \quad (9) \\ &= L \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t), \end{aligned}$$

where the last equality follows from Euler's theorem for homogeneous functions. Now, observe that this last term can be lower bounded as:

$$\langle \boldsymbol{\theta}_t, -\nabla \mathcal{L}(\boldsymbol{\theta}_t) \rangle \geq L \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t) \geq L \mathcal{L}(\boldsymbol{\theta}_t) \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)}, \quad (10)$$

where we used the fact  $e^{-\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \leq \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} = \mathcal{L}(\boldsymbol{\theta}_t)$ . We have made the first term of Equation 8 appear. By plugging Equation 10 into Equation 8, we get:

$$\begin{aligned} \frac{d \log \tilde{\gamma}}{dt} &\geq \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{L}{\langle \boldsymbol{\theta}_t, -\nabla \mathcal{L}(\boldsymbol{\theta}_t) \rangle} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \\ &\geq \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{L}{\|\boldsymbol{\theta}_t\| \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|_*} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right). \quad (\text{definition of dual norm}) \end{aligned} \quad (11)$$

Noticing that  $\|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|_* = \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|$  (from Equation 7) concludes the proof.  $\blacksquare$

**Remark 2** Observe that it is the geometric margin induced by the dual norm of the algorithm that is non-decreasing, and not any geometric margin. The proof crucially relies on the fact that  $\|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|_* = \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|$ .

### 3.2 Convergence to KKT Points of the Max-Margin Problem

The previous theorem is a first indication that steepest flow implicitly maximizes the  $\|\cdot\|_*$ -geometric margin in deep neural networks. However, the monotonicity of the (soft) margin alone does not imply anything about its final value and its optimality. In this section, we provide a concrete characterization of the *asymptotic* behavior of steepest flow: we show that any limit point of the iterates produced by steepest flow is along the direction of a Karush-Kuhn-Tucker (KKT) point of the following margin maximization (MM) optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\theta} \in \mathbb{R}^p} \quad & \frac{1}{2} \|\boldsymbol{\theta}\|^2 \\ \text{s.t.} \quad & y_i f(\mathbf{x}_i; \boldsymbol{\theta}) \geq 1, \quad \forall i \in [m]. \end{aligned} \quad (\text{MM})$$

Let us recall the definition of a Karush-Kuhn-Tucker point (Karush, 1939; Kuhn, H. W. and Tucker, A. W., 1951).

**Definition 3 (KKT point)** A feasible point  $\boldsymbol{\theta} \in \mathbb{R}^p$  of (MM) is a Karush-Kuhn-Tucker (KKT) point, if there exist  $\lambda_1, \dots, \lambda_m \geq 0$  such that:

1.  $\partial \frac{1}{2} \|\boldsymbol{\theta}\|^2 + \sum_{i=1}^m \lambda_i \partial (1 - y_i f(\mathbf{x}_i; \boldsymbol{\theta})) \ni 0$ .
2.  $\lambda_i (1 - y_i f(\mathbf{x}_i; \boldsymbol{\theta})) = 0, \quad \forall i \in [m]$ .

Notice that the first *stationarity* condition is defined using set addition, since we are dealing with non-differentiable functions. See Dutta et al. (2013) for more details on optimization problems with non-smooth objectives/constraints. Under some regularity assumptions, the KKT conditions become necessary conditions for global optimality and for non-convex problems like (MM) they might be the best characterization of optimality we can hope for. See Lemma 15 in Appendix A for details.

We are now ready to state our main result.

**Theorem 4** Under assumptions (A1), (A2), (A3), consider steepest flow with respect to a norm  $\|\cdot\|$  (Equation 3) on the exponential loss  $\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta})}$ . Then, any limit point  $\bar{\boldsymbol{\theta}}$  of  $\left\{ \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|} \right\}_{t \geq 0}$  is along the direction of a KKT point of the optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\theta} \in \mathbb{R}^p} \frac{1}{2} \|\boldsymbol{\theta}\|^2 \\ \text{s.t. } y_i f(\mathbf{x}_i; \boldsymbol{\theta}) \geq 1, \forall i \in [m]. \end{aligned} \quad (12)$$

Theorem 4 implies that if the limit of the normalized iterates induced by steepest flow exists, then it is proportional to a KKT point of a margin maximization problem. The proof technique shows that a scaled version of any limit point  $\bar{\boldsymbol{\theta}}$  is an *approximate* KKT point (Def. 16) with an arbitrarily small error. The full proof can be found in Appendix A.

Note that since Theorem 4 is true for *any norm*  $\|\cdot\|$ , the main contribution of Lyu and Li (2020), which characterizes the implicit bias of gradient flow in homogeneous deep networks, can be recovered by our result when  $\|\cdot\| = \|\cdot\|_2$ . Notably, Theorem 4 generalizes the result of Lyu and Li (2020) in the following cases of algorithm norms:

- Any  $\ell_p$  norm with  $p \in [1, \infty]$ . This corresponds, for example, to *coordinate descent* (steepest descent with respect to the  $\ell_1$  norm) and *sign gradient descent* (normalized steepest descent with respect to the  $\ell_\infty$  norm).
- The *modular* norm which was recently introduced by Large et al. (2024) to accommodate scalable neural network training. In particular, let a feed-forward neural network with  $L + 1$  layers be  $f(\mathbf{x}; \boldsymbol{\theta} = \{\mathbf{W}_1, \dots, \mathbf{W}_L, \mathbf{u}\}) = \langle \mathbf{u}, \sigma(\mathbf{W}_L \sigma(\mathbf{W}_{L-1} \dots \sigma(\mathbf{W}_1 \mathbf{x})) \rangle$ , where  $\sigma: \mathbb{R} \rightarrow \mathbb{R}$  is a homogeneous activation applied element-wise. Then, the modular norm induced by the architecture  $f$  is given by  $\|\boldsymbol{\theta}\|_{\mathcal{W}} = \max(\|\mathbf{W}_L\|_L, \dots, \|\mathbf{W}_1\|_1, \|\mathbf{u}\|_u)$  where each of the norms  $\|\cdot\|_L, \dots, \|\cdot\|_1, \|\cdot\|_u$  can be different. For instance, if the norm of each layer is chosen to be the *spectral* norm, then steepest descent with respect to the modular norm  $\|\cdot\|_{\mathcal{W}}$  corresponds to Shampoo (Gupta et al., 2018) without momentum (see Appendix C.1 for details).

Complementary to Theorem 4, we note that we can obtain a geometric characterization of the evolution of the algorithm after  $t_0$ : the (generalized)<sup>3</sup> Bregman divergence induced by the squared norm of the algorithm between the subgradient of the objective function of MM and the subgradient of its constraints reduces at a rate that depends on the *alignment* between the normalized iterates and the negative normalized, minimum norm, subgradient of the loss  $\left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^*}{\|\mathbf{g}_t^*\|_*} \right\rangle$ . In the case of algorithms whose norm squared is a smooth function, this implies that the finite-time iterates are along the direction of an approximate KKT point of MM. In contrast, the proof of Theorem 4 characterizes the limit points directly and, as a result, does not provide any finite-time characterization of the iterates in terms of proximity to stationarity. See Appendix B for details.

**Comparison to a previous version.** An earlier version of the current work that appeared in ICLR 2025 (Tsilivis et al., 2025) proved a weaker version of Theorem 4. In particular, it established that, under the same conditions, the limit points of steepest descent are along

3. See Definition 26.

the direction of a class of *generalized* KKT points of the margin maximization problem that we introduced. The proof proceeded by characterizing the direction of the iterates  $\theta_t$  in terms of their proximity to stationarity. In fact, this weaker notion of stationarity allowed us to explicitly control proximity to (generalized) stationarity, yet it implied convergence to KKT points only in the case of norms whose square is a smooth function. This disqualified interesting and practically relevant cases, such as coordinate descent and sign gradient descent. At a technical level, the proof of Theorem 4 does not rely on sequences of approximate KKT points (unlike our previous version and the proof of Lyu and Li (2020)), but instead shows directly that a scaled version of any limit point is stationary.

To the best of our knowledge, this is a first result about the implicit bias of an algorithm in the parameter space of deep homogeneous neural networks which is not about  $\ell_2$ -geometric margin maximization.

## 4. Experiments

In this section, we train one-hidden layer neural networks with various steepest descent algorithms (gradient descent-GD, coordinate descent-CD, sign descent-SD) to confirm the validity and measure the robustness of the theoretical claims, and to discuss the connection between Adam, Shampoo and steepest descent algorithms. Amongst other quantities, we measure the three relevant geometric margins during training, which, in the context of one-hidden layer neural networks with 1-homogeneous activations and without biases, become:

$$\gamma_1 = \min_{i \in [m]} \frac{y_i f(\mathbf{x}_i; \theta)}{\|\theta\|_\infty^2}, \quad \gamma_2 = \min_{i \in [m]} \frac{y_i f(\mathbf{x}_i; \theta)}{\|\theta\|_2^2}, \quad \gamma_\infty = \min_{i \in [m]} \frac{y_i f(\mathbf{x}_i; \theta)}{\|\theta\|_1^2}. \quad (13)$$

### 4.1 Teacher – Student Experiments

We first perform experiments in a controlled environment, where the generative process consists of Gaussian data passed through a one-hidden layer (“teacher”) neural network, which is sparse. Specifically:

$$\mathbf{x} \sim \mathcal{N}(0, I_d), \quad y = \text{sgn} \left( f_{\text{teacher}} \left( \mathbf{x}; \theta^* = \{u_j^*, \mathbf{w}_j^*\}_{j=1}^k \right) \right) = \text{sgn} \left( \sum_{j=1}^k u_j^* \sigma(\langle \mathbf{w}_j^*, \mathbf{x} \rangle) \right), \quad (14)$$

where  $\sigma(u) = \max(u, 0)$  is the ReLU activation,  $\text{sgn}(\cdot)$  returns the sign of a number, and  $\|\theta^*\|_0$  is assumed to be small. We train (“student”) neural networks of the same architecture, but of larger width and with randomly initialized weights:  $f_{\text{student}}(\mathbf{x}; \theta) = \sum_{j=1}^{k'} u_j \sigma(\langle \mathbf{w}_j, \mathbf{x} \rangle)$ , with width  $k' > k$  and  $w_{jl} \sim \mathcal{U}[-\frac{\alpha}{d}, \frac{\alpha}{d}]$ ,  $j \in [k'], l \in [d]$ ,  $u_j \in \mathcal{U}[-\frac{\alpha}{k'}, \frac{\alpha}{k'}]$  (for CD we use:  $w_{jl} \sim \mathcal{U}[-\frac{\alpha}{k'}, \frac{\alpha}{k'}]$ ,  $j \in [k'], l \in [d]$  in order to keep all the individual parameters to the same scale at initialization). The magnitude of initialization  $\alpha$  can control how fast the implicit bias of the algorithm kicks in, with smaller values entering this “rich” regime faster (Woodworth et al., 2020). We compare the performance of (full batch) GD, CD and SD in minimizing the empirical exponential loss on  $m$  independent points sampled from the generative process of Equation 14. Section D contains full experimental details. According to Theorems 1 and 4, we expect GD to favor solutions with small  $\ell_2$  norm. This is equivalent to a small sum of the product of the magnitude of incoming and outgoing weights across

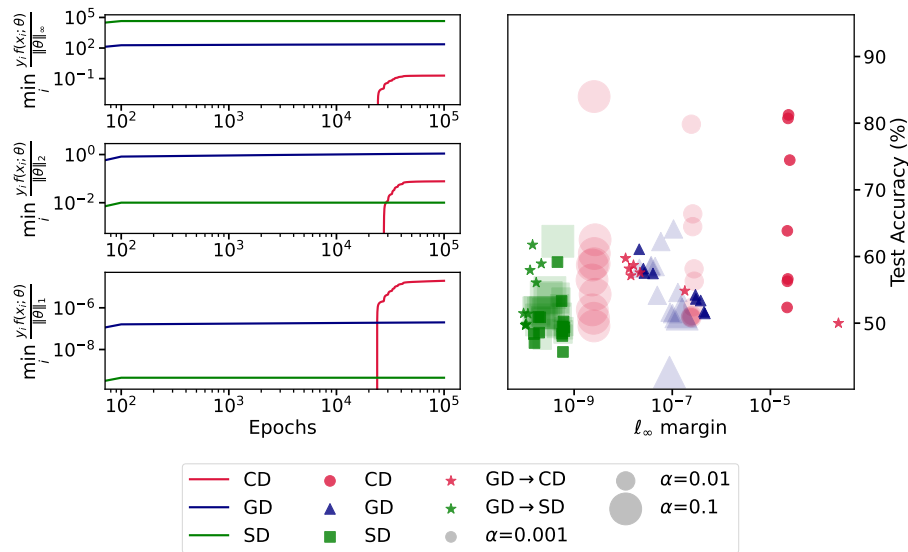


Figure 1: **Evaluation of steepest descent algorithms in a teacher-student setup.** *Left:* Geometric margins ( $\gamma_1, \gamma_2, \gamma_\infty$  in Equation 13) over the course of training (average over 20 different seeds). *Right:* Final test accuracy vs final  $\ell_\infty$  margin ( $\gamma_\infty$ ). Each point in the 2d space corresponds to a different run (only showing runs that did not diverge). Larger points correspond to larger initialization scales  $\alpha$ . The star points are produced by switching from GD to CD (red) or SD (green), right after the point of perfect train accuracy.

all neurons (Theorem 1 in Neyshabur et al. (2015b)). On the other hand, CD will seek to minimize the  $\ell_1$  norm of the parameters, which translates to a narrow network with sparse 1st-layer weights. Finally, SD’s bias towards small  $\|\theta\|_\infty$  solutions does not appear to be useful for generalizing from few samples in this task. Therefore, given that  $\theta^*$  has small  $\ell_0$  “norm”, we expect  $\text{CD} > \text{GD} > \text{SD}$  in terms of generalization.

Figure 1 displays our main results. We summarize our key findings below:

- (i) **Margins increase past  $t_0$ .** As expected from Lemma 1, we observe that, right after the point of separation, each algorithm implicitly increases its corresponding geometric margin (Figure 1 left). Furthermore, we observe that the ordering of the algorithms is as expected for each margin (SD attains larger  $\ell_1$  margin than GD and CD, etc.), despite the fact that Theorem 4 only guarantees convergence to a KKT point of the margin maximization problem - note the log-log plot.
- (ii) **Smaller initialization produces larger geometric margin.** A smaller magnitude of initialization  $\alpha$  causes a larger eventual value of the geometric margin (see Figure 1 right for CD and  $\gamma_\infty$ , where this effect is stronger, and Figure 4 in Appendix D for  $\gamma_1, \gamma_2$ ).
- (iii) **Importance of early-stage dynamics for generalization.** Figure 1, right, shows the final test accuracy of the networks (20 different runs) vs the value of their final  $\ell_\infty$  margin ( $\gamma_\infty$ ). We observe that, while there exist more CD runs with good generalization (red circles), these do not always coincide with larger  $\gamma_\infty$ . Furthermore, intervening

in the algorithms to encourage or discourage  $\gamma_\infty$ -maximization does not result in significant generalization changes: after running GD until the point of perfect train accuracy, we switch to either SD (green stars) or CD (red stars) to directly control the late stage geometric properties of the model. Switching to CD seems to bear marginal benefits in terms of generalization, even though all the switched runs reach smaller values of  $\ell_\infty$  margin compared to the full, no-switching, GD runs. These benefits, however, pale in comparison to the full CD runs. Switching to SD, on the other hand, results in smaller  $\gamma_\infty$  and similar or marginally worse test accuracy. See also Figure 4 in Appendix D for test accuracy vs the other two geometric margins. We conclude that it is unlikely that large generalization benefits can solely and causally be linked to larger geometric margins in this setup, and it appears that the early stage dynamics play an important role for generalization.

## 4.2 Connection between Adam and Sign-GD

Adaptive optimization methods like Adam (Kingma and Ba, 2015) have been popular in deep learning applications, yet theoretically their value has been questioned (Wilson et al., 2017) and their properties remain poorly understood. Wang et al. (2021, 2022) studied the implicit bias of Adam in homogeneous networks and concluded that Adam shares the same asymptotic properties as GD. More recently, this conclusion has been challenged (Zhang et al., 2024), in the sense that this asymptotic property crucially depends on a precision parameter of the algorithm and does not capture realistic runs of the algorithm (see Section C.1 for details). In particular, it was shown that in linear models, Adam, without this precision parameter, implicitly maximizes the  $\ell_1$ -geometric margin (Zhang et al., 2024), a property shared with SD and not GD. Indeed, Adam without momentum, and ignoring the precision parameter, is equivalent to SD (see Section C.1). Setting the precision parameter to 0, on the other hand, is not useful in applications, as small initial values of the gradient result in divergence of the loss. A question arises: what, then, are the relevant geometric properties of Adam *in practice*?

Figure 2 provides some experimental answers to this question, in light of Theorems 1 and 4. On a pair of digits extracted from MNIST we train two-layer neural networks with GD, SD and Adam, with small initialization. See Section D for experimental details. We observe that, as soon as the algorithms reach 100% train accuracy, the margins

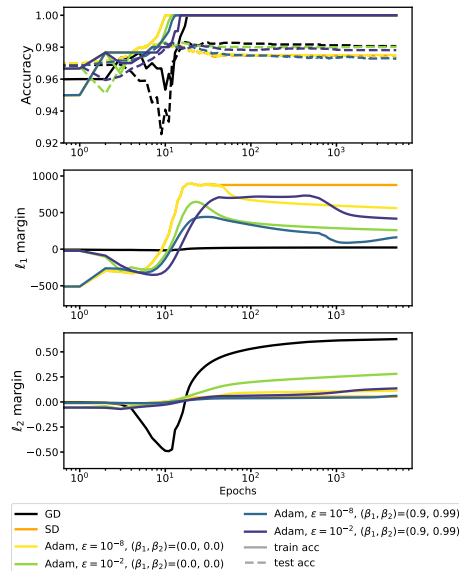


Figure 2: **Accuracy, and  $\ell_1, \ell_2$  margins during training for GD, SD and Adam on MNIST (3 random seeds).** Adam is parameterized by a numerical precision constant  $\epsilon$  and two momentum parameters  $(\beta_1, \beta_2)$  (defaulting to  $10^{-8}$  and  $(0.9, 0.99)$ ). We observe that Adam behaves similarly to SD for the period right after the point of perfect train accuracy.




Figure 3: **Loss, spectral and  $\ell_1$  margin during training for GD, Adam and Shampoo on MNIST with a two-layer neural network with a frozen second layer.** We observe that Shampoo ends up with a larger spectral margin.

start to increase (as Theorem 1 suggests); SD reaches a larger value of  $\gamma_1$ , while GD reaches a larger value of  $\gamma_2$ . Interestingly, Adam with the default hyperparameters (precision  $\epsilon = 10^{-8}$  and non-zero momentum), initially, behaves similarly to SD, increasing  $\gamma_1$ , before it starts decreasing it, in order to slowly start increasing  $\gamma_2$ ! Curiously, larger values of  $\epsilon$  increase  $\gamma_1$  even further and start the second phase slower, but more aggressively. Notice, however, that train and test accuracies have long converged, so it is unlikely that a typical run would have lasted long enough to see the second phase of  $\ell_2$ -margin maximization (in particular, the loss value needs to be smaller than  $10^{-7}$  in order to observe such behavior). Similar observations hold for Adam without momentum (recall that without momentum and for  $\epsilon \rightarrow 0$ , we recover SD). Therefore, it appears that the  $\ell_1$  bias of SD (Theorems 1, 4 for  $\|\cdot\| = \|\cdot\|_\infty$ ) more faithfully describes a typical run of Adam in neural networks.

### 4.3 Shampoo favors a spectral margin

Gupta et al. (2018) introduced Shampoo as an efficient method for deep learning optimization and the algorithm has attracted a lot of interest from practitioners and theorists alike. At a high level, Shampoo is a preconditioning method which respects the structure of the optimization space on which it operates. The algorithm and its variants have shown to speed up convergence in a variety of settings, including language modeling with autoregressive transformers (Jordan et al., 2024). However, little is understood about the generalization properties of the algorithm. Our results let us clarify some of its late-stage properties.

In particular, as observed by Bernstein and Newhouse (2024), Shampoo operating on each layer with accumulation turned off corresponds to steepest descent with respect to the spectral norm of the weight matrix of that layer. Overall, it corresponds to steepest descent with respect to an architecture-dependent norm in parameter space. See Appendix C.2 for more details. For a two-layer neural network with a fully-connected layer, a frozen second layer and scalar output  $f(\mathbf{x}; \mathbf{W}) = \langle \mathbf{u}, \sigma(\mathbf{W}\mathbf{x}) \rangle$ , Shampoo is steepest descent with respect to the spectral norm  $\sigma_{\max}(\mathbf{W})$ . As a result, Theorems 1 and 4 predict that Shampoo implicitly maximizes the *spectral* margin:

$$\gamma_\sigma = \min_{i \in [m]} \frac{y_i f(\mathbf{x}_i; \mathbf{W})}{\sigma_{\max}(\mathbf{W})}. \quad (15)$$

In Figure 3 we provide experimental evidence for the above. We train two-layer neural networks (with a frozen second layer) on a pair of digits extracted from MNIST, starting from small initialization. We compare an off-the-shelf implementation of **Shampoo** (Novik, 2020) to **Adam** and **GD**, all without momentum. As we observe, the spectral margin  $\gamma_\sigma$  increases once the algorithms reach perfect train accuracy, while **Shampoo** achieves the largest value among the three.

## 5. Conclusion

In our work, we considered the large family of steepest descent algorithms with respect to an arbitrary norm  $\|\cdot\|$  and provided a unifying theoretical analysis of their late-stage implicit bias when training homogeneous neural networks. Our main result establishes that homogeneous neural networks implicitly maximize the geometric margin of the training points induced by the norm of the algorithm.

Our results can reinforce several recent efforts that attempt to understand deep learning through the lens of implicit bias. In particular, questions about generalization, robustness, and privacy can now be asked more broadly: (a) What are the important properties of language data distributions which render algorithms such as Adam ( $\approx$  sign gradient descent) and Shampoo ( $\approx$  spectral steepest descent) appealing for training deep networks on such tasks? (b) An interesting connection has been made between implicit bias and the ability to retrieve data points from the parameters of a trained neural network (Haim et al., 2022). In particular, when the parameters converge to a stationary point of a certain margin-maximization problem, they satisfy equations with respect to the training dataset, which allow to reconstruct portions of the training data. Our work leads to the question: Is it also possible to extract training samples from neural networks optimized with Adam or Shampoo? (c) Can we leverage our implicit bias results to design more sample-efficient algorithms for adversarially robust training, as argued by Tsilivis et al. (2024)? (d) A central topic in deep network generalization is their ability to avoid overfitting despite the presence of noise in the data. In prior work, this phenomenon has been attributed to the geometric properties of the networks, which are a consequence of the implicit bias of gradient descent (Frei et al., 2022; Shamir, 2023). Is benign overfitting an inherent feature of first-order methods, or are current results specifically tailored to gradient descent?

## Acknowledgments

NT and JK acknowledge support through the NSF under award 1922658. GV is supported by The Israel Science Foundation (grant No. 2574/25), by a research grant from Mortimer Zuckerman (the Zuckerman STEM Leadership Program), and by research grants from the Center for New Scientists at the Weizmann Institute of Science, and the Shimon and Golde Picker – Weizmann Annual Grant. We would like to thank anonymous referees that helped improving an earlier version of this work. We would like to thank Zhiyuan Li and Kaifeng Lyu for a useful discussion. Early parts of this work were done while NT was visiting the Toyota Technological Institute of Chicago (TTIC) during the winter of 2024, and NT would like to thank everyone at TTIC for their hospitality, which enabled this work. Part of this

work was done while JK and NT were hosted by the Centre Sciences de Données at the École Normale Supérieure (ENS) in 2023/24, whose hospitality we gratefully acknowledge. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise.

## Appendix A. Proof Details

In this section, we provide proofs for the results stated in the main text.

### A.1 Steepest Flow

We first present a series of technical results, which are about steepest flow in the case of non-differentiable loss functions. In what follows, we will denote with  $\mathbf{g}_t^*$  any loss subderivative with minimum  $\|\cdot\|_*$  norm, i.e.  $\mathbf{g}_t^* \in \operatorname{argmin}_{\mathbf{u} \in \partial \mathcal{L}(\theta_t)} \|\mathbf{u}\|_*$ . In the case of subdifferentials, chain rule holds as an inclusion:

**Theorem 5 (Theorem 2.3.9 and 2.3.10 in Clarke (1990))** *Let  $z_1, \dots, z_n : \mathbb{R}^d \rightarrow \mathbb{R}$ ,  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  be locally Lipschitz functions and define  $\mathbf{z} = (z_1, \dots, z_n)$ . Let  $(f \circ \mathbf{z})(\mathbf{x}) = f(z_1(\mathbf{x}), \dots, z_n(\mathbf{x}))$  be the composition of  $\mathbf{z}$  with  $f$ . Then, it holds:*

$$\partial(f \circ \mathbf{z})(\mathbf{x}) \subseteq \operatorname{conv} \left\{ \sum_{i=1}^n \alpha_i \mathbf{h}_i : \alpha \in \partial f(z_1(\mathbf{x}), \dots, z_n(\mathbf{x})), \mathbf{h}_i \in \partial z_i(\mathbf{x}) \right\}. \quad (16)$$

To further analyze steepest flows and to guarantee loss monotonicity, we need a stronger chain rule result. This can be achieved for a large class of locally Lipschitz functions, as per the following theorem which is due to Davis et al. (2020).

**Theorem 6 (Theorem 5.8 in Davis et al. (2020))** *If  $F : \mathbb{R}^k \rightarrow \mathbb{R}$  is locally Lipschitz and Whitney  $C^1$ -stratifiable, then it admits a chain rule: for all arcs (functions which are absolutely continuous on every compact subinterval)  $\mathbf{u} : [0, \infty) \rightarrow \mathbb{R}^k$ , almost all  $t \geq 0$ , and all  $\mathbf{g} \in \partial F(\mathbf{u}(t))$ , it holds:*

$$\frac{dF(\mathbf{u}(t))}{dt} = \left\langle \mathbf{g}, \frac{d\mathbf{u}(t)}{dt} \right\rangle. \quad (17)$$

Whitney  $C^1$ -stratifiability includes a large family of functions, including functions defined in an o-minimal structure which has been a previous assumption in the literature (Ji and Telgarsky, 2020). It excludes some pathological functions - see, for instance, Appendix J in (Lyu and Li, 2020). This version of chain rule allows us to derive the following central properties of steepest flows.

**Proposition 7** *Let  $\mathcal{L} : \mathbb{R}^p \rightarrow \mathbb{R}$  and assume that  $\mathcal{L}$  admits a chain rule. Then, for the steepest flow iterates of Equation 1, it holds for almost any  $t \geq 0$ :*

$$\frac{d\mathcal{L}}{dt} = - \left\| \frac{d\theta}{dt} \right\|^2 \leq 0, \quad (18)$$

and

$$\left\langle \frac{d\theta}{dt}, -\mathbf{g}_t^* \right\rangle = \left\| \frac{d\theta}{dt} \right\|^2 = \|\mathbf{g}_t^*\|_*^2, \quad (19)$$

where  $\mathbf{g}_t^* \in \operatorname{argmin}_{\mathbf{u} \in \partial \mathcal{L}(\theta_t)} \|\mathbf{u}\|_*$ .

**Proof** From Theorem 6, for almost any  $t \geq 0$ , it holds  $\forall \mathbf{g}_t \in \partial \mathcal{L}(\theta_t)$ :

$$\frac{d\mathcal{L}}{dt} = \left\langle \mathbf{g}_t, \frac{d\theta}{dt} \right\rangle. \quad (20)$$

Applying this for the element of  $\partial\mathcal{L}(\boldsymbol{\theta}_t)$ ,  $\mathbf{g}'_t$ , that corresponds to  $\frac{d\boldsymbol{\theta}}{dt}$  from the definition of steepest flow Equation 3, we get:

$$\frac{d\mathcal{L}}{dt} = \left\langle \mathbf{g}'_t, \frac{d\boldsymbol{\theta}}{dt} \right\rangle = - \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2, \quad (21)$$

where the last equality follows from the definition of the dual norm. But, Equation 20 for  $\mathbf{g}_t^* \in \operatorname{argmin}_{\mathbf{u} \in \partial\mathcal{L}(\boldsymbol{\theta}_t)} \|\mathbf{u}\|_*$ , yields:

$$\left| \frac{d\mathcal{L}}{dt} \right| = \left| \left\langle \mathbf{g}_t^*, \frac{d\boldsymbol{\theta}}{dt} \right\rangle \right| \leq \|\mathbf{g}_t^*\|_* \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|. \quad (22)$$

Thus, combining Equation 21, Equation 22, we obtain:

$$\left\| \frac{d\boldsymbol{\theta}}{dt} \right\| \leq \|\mathbf{g}_t^*\|_*, \quad (23)$$

which implies that the update rule Equation 3 is equivalent to:

$$\frac{d\boldsymbol{\theta}}{dt} \in \left\{ \operatorname{argmin}_{\|\mathbf{u}\| \leq \|\mathbf{g}_t^*\|_*} \langle \mathbf{u}, \mathbf{g}_t^* \rangle : \mathbf{g}_t^* \in \operatorname{argmin}_{\mathbf{u} \in \partial\mathcal{L}(\boldsymbol{\theta}_t)} \|\mathbf{u}\|_* \right\}. \quad (24)$$

Therefore, from the definition of the dual norm, we have:

$$\left\langle \frac{d\boldsymbol{\theta}}{dt}, -\mathbf{g}_t^* \right\rangle = \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 = \|\mathbf{g}_t^*\|_*^2. \quad (25)$$

■

Hence, under the mild assumptions of Theorem 6, the loss is non-increasing during training.

## A.2 Late Phase Implicit Bias

A useful standard characterization of the subdifferential of a norm is the following:

### Lemma 8

$$\partial\|\mathbf{x}\| = \{\mathbf{v} : \langle \mathbf{v}, \mathbf{x} \rangle = \|\mathbf{x}\|, \|\mathbf{v}\|_* \leq 1\}$$

We present the proofs for our results about the late stage of training in steepest flow algorithms. The next lemma quantifies the behavior of the smooth margin past the point  $t_0$  (where, recall, zero classification error is achieved).

**Theorem 9 (Soft margin increases - full version)** *For almost any  $t > t_0$ , it holds:*

$$\frac{d \log \tilde{\gamma}}{dt} \geq L \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{1}{L\mathcal{L}(\boldsymbol{\theta}_t) \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)}} - \frac{1}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \geq 0.$$

**Proof** Let  $\mathbf{n}_t \in \partial\|\boldsymbol{\theta}_t\|$ . We have:

$$\begin{aligned}
 \frac{d \log \tilde{\gamma}}{dt} &= \frac{d}{dt} \log \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} - L \frac{d}{dt} \log \|\boldsymbol{\theta}_t\| \\
 &= \frac{d}{dt} \log \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} - L \left\langle \frac{\mathbf{n}_t}{\|\boldsymbol{\theta}_t\|}, \frac{d\boldsymbol{\theta}}{dt} \right\rangle \quad (\text{Chain rule}) \\
 &\geq \frac{d}{dt} \log \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} - L \frac{\|d\boldsymbol{\theta}/dt\|}{\|\boldsymbol{\theta}_t\|} \quad (\text{definition of dual norm and } \|\mathbf{n}_t\|_\star \leq 1) \quad (26) \\
 &= -\frac{d\mathcal{L}(\boldsymbol{\theta}_t)}{dt} \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t) \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)}} - L \frac{\|d\boldsymbol{\theta}/dt\|}{\|\boldsymbol{\theta}_t\|} \quad (\text{Chain rule}) \\
 &= \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t) \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)}} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \quad (\text{Equation 18}).
 \end{aligned}$$

But, the first term inside the parenthesis can be related to the second one via the following calculation. Recall that, by Theorem 6, for any  $\mathbf{g}_t \in \partial\mathcal{L}(\boldsymbol{\theta}_t)$  there exist  $\mathbf{h}_1 \in \partial y_1 f(\mathbf{x}_1; \boldsymbol{\theta}_t), \dots, \mathbf{h}_m \in \partial y_m f(\mathbf{x}_m; \boldsymbol{\theta}_t)$  such that  $\mathbf{g}_t = \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \mathbf{h}_i$ . Thus, for a minimum norm subderivative  $\mathbf{g}_t^\star$ , we have:

$$\begin{aligned}
 \langle \boldsymbol{\theta}_t, -\mathbf{g}_t^\star \rangle &= \left\langle \boldsymbol{\theta}_t, \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \mathbf{h}_i^\star \right\rangle \\
 &= \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \langle \boldsymbol{\theta}_t, \mathbf{h}_i^\star \rangle \quad (27) \\
 &= L \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t),
 \end{aligned}$$

where the last equality follows from Euler's theorem for homogeneous functions (whose generalization for subderivatives can be found in Theorem B.2 in Lyu and Li (2020)). Now, observe that this last term can be lower bounded as:

$$\langle \boldsymbol{\theta}_t, -\mathbf{g}_t^\star \rangle \geq L \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t) \geq L \mathcal{L}(\boldsymbol{\theta}_t) \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)}, \quad (28)$$

where we used the fact  $e^{-\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \leq \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)}$ . We have made the first term of Equation 26 appear. By plugging Equation 28 into Equation 26, we get:

$$\begin{aligned}
 \frac{d \log \tilde{\gamma}}{dt} &\geq \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{L}{\langle \boldsymbol{\theta}_t, -\mathbf{g}_t^\star \rangle} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \\
 &\geq \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{L}{\|\boldsymbol{\theta}_t\| \|\mathbf{g}_t^\star\|_\star} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \quad (\text{definition of dual norm}). \quad (29)
 \end{aligned}$$

Noticing that  $\|\mathbf{g}_t^\star\|_\star = \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|$  (from Proposition 7) concludes the proof.  $\blacksquare$

By extending Lemma B.6 of Lyu and Li (2020), we can further prove that the loss converges to 0 and, thus, the norm of the iterates diverges to infinity.

**Lemma 10** *As  $t \rightarrow \infty$ ,  $\mathcal{L}(\boldsymbol{\theta}_t) \rightarrow 0$  and  $\|\boldsymbol{\theta}_t\| \rightarrow \infty$ .*

**Proof** We suppress the dependence of the loss and the iterates from time  $t$ , when it is obvious from the context.

From the definition of the steepest flow update and chain rule (Equation 18), we have

$$-\frac{d\mathcal{L}}{dt} = \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 = \|\mathbf{g}_t^*\|_*^2 \geq \frac{1}{\|\boldsymbol{\theta}\|^2} \langle \boldsymbol{\theta}, -\mathbf{g}_t^* \rangle^2, \quad (30)$$

where we applied Equation 19, Equation 18 and the definition of the dual norm. But, as we showed in Equation 28, the above inner product can be upper bounded by a function of the loss, so, by plugging in, we get:

$$-\frac{d\mathcal{L}}{dt} \geq \frac{L^2}{\|\boldsymbol{\theta}\|^2} \left( \mathcal{L} \log \frac{1}{\mathcal{L}} \right)^2 = \frac{L^2}{(\log \frac{1}{\mathcal{L}})^{2/L} \tilde{\gamma}^{2/L}(t)} \left( \mathcal{L} \log \frac{1}{\mathcal{L}} \right)^2 \geq \frac{L^2}{(\log \frac{1}{\mathcal{L}})^{2/L} \tilde{\gamma}^{2/L}(t_0)} \left( \mathcal{L} \log \frac{1}{\mathcal{L}} \right)^2, \quad (31)$$

which follows from the definition of the margin (Equation 5) and its monotonicity (Lemma 9). By rearranging:

$$-\frac{d\mathcal{L}}{dt} \frac{1}{\mathcal{L}^2} \left( \log \frac{1}{\mathcal{L}} \right)^{2/L-2} \geq L^2 \tilde{\gamma}(t_0)^{2/L}, \quad (32)$$

and integrating over time from  $t_0$  to  $t > t_0$ , we further have:

$$\int_{t_0}^t \left( \log \frac{1}{\mathcal{L}} \right)^{2/L-2} \frac{d}{dt} \frac{1}{\mathcal{L}} dt \geq L^2 \tilde{\gamma}(t_0)^{2/L} (t - t_0), \quad (33)$$

or, by a change of variables,

$$\int_{1/\mathcal{L}(t_0)}^{1/\mathcal{L}(t)} (\log u)^{2/L-2} du \geq L^2 \tilde{\gamma}(t_0)^{2/L} (t - t_0). \quad (34)$$

The RHS diverges to infinity as  $t \rightarrow \infty$ , hence so does the LHS, which can only happen if  $\mathcal{L} \rightarrow 0$ . In order for  $\mathcal{L}(\boldsymbol{\theta}_t) = \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} = \sum_{i=1}^m e^{-y_i \|\boldsymbol{\theta}_t\|^L f(\mathbf{x}_i; \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|})}$  to go to zero, it must be  $\|\boldsymbol{\theta}_t\| \rightarrow \infty$ .  $\blacksquare$

The following Lemma quantifies the connection between soft and hard margin.

**Lemma 11** *For any  $\boldsymbol{\theta}$ , it holds:*

$$\frac{\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}) - \log m}{\|\boldsymbol{\theta}\|^L} \leq \tilde{\gamma} \leq \frac{\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta})}{\|\boldsymbol{\theta}\|^L}. \quad (35)$$

**Proof** Follows from:

$$e^{-\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta})} \leq \mathcal{L}(\boldsymbol{\theta}) \leq m e^{-\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta})}. \quad (36)$$

From the previous two Lemmata, we deduce that the soft margin converges to the hard margin as  $t \rightarrow \infty$ .

**Corollary 12** For any  $t > t_0$ ,  $\boldsymbol{\theta}_t \in \mathbb{R}^p$ , let  $\gamma(\boldsymbol{\theta}_t) = \frac{\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)}{\|\boldsymbol{\theta}_t\|^L}$ . Then, it holds:

$$\lim_{t \rightarrow \infty} |\tilde{\gamma}(\boldsymbol{\theta}_t) - \gamma(\boldsymbol{\theta}_t)| = 0. \quad (37)$$

**Proof** By taking limits in Equation 35, we have:

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta})}{\|\boldsymbol{\theta}\|^L} - \frac{\log m}{\|\boldsymbol{\theta}\|^L} &\leq \lim_{t \rightarrow \infty} \tilde{\gamma}(\boldsymbol{\theta}_t) \leq \lim_{t \rightarrow \infty} \frac{\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta})}{\|\boldsymbol{\theta}\|^L} \iff \\ \lim_{t \rightarrow \infty} \gamma(\boldsymbol{\theta}_t) - \lim_{t \rightarrow \infty} \frac{\log m}{\|\boldsymbol{\theta}\|^L} &\leq \lim_{t \rightarrow \infty} \tilde{\gamma}(\boldsymbol{\theta}_t) \leq \lim_{t \rightarrow \infty} \gamma(\boldsymbol{\theta}_t). \end{aligned} \quad (38)$$

But, from Lemma 10, we know that  $\|\boldsymbol{\theta}_t\| \rightarrow \infty$ . Thus,

$$\lim_{t \rightarrow \infty} \gamma(\boldsymbol{\theta}_t) \leq \lim_{t \rightarrow \infty} \tilde{\gamma}(\boldsymbol{\theta}_t) \leq \lim_{t \rightarrow \infty} \gamma(\boldsymbol{\theta}_t), \quad (39)$$

which shows the claim.  $\blacksquare$

The last part of the proof consists of characterizing the (directional) convergence of the iterates in relation to stationary points of the following optimization problem (re-introduced here for convenience):

$$\begin{aligned} \min_{\boldsymbol{\theta} \in \mathbb{R}^p} \frac{1}{2} \|\boldsymbol{\theta}\|^2 \\ \text{s.t. } y_i f(\mathbf{x}_i; \boldsymbol{\theta}) \geq 1, \forall i \in [m]. \end{aligned} \quad (40)$$

We first show that we can always construct a feasible point of Equation 40 from a scaled version of  $\boldsymbol{\theta}_t$ .

**Lemma 13** For any  $t > 0$ ,  $\tilde{\boldsymbol{\theta}}_t = \frac{\boldsymbol{\theta}_t}{(\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))^{\frac{1}{L}}}$  is a feasible point of Equation 40.

**Proof** From the homogeneity of  $f$ , we have:

$$y_i f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}_t) = y_i f\left(\mathbf{x}_i; \frac{\boldsymbol{\theta}_t}{(\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))^{\frac{1}{L}}}\right) = \frac{y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)}{\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \geq 1 \quad (41)$$

for all  $i \in [m]$ . So  $\tilde{\boldsymbol{\theta}}_t$  is a feasible point of Equation 40.  $\blacksquare$

Under some regularity assumptions, the KKT conditions (Definition 3) become necessary for global optimality (yet, not sufficient):

**Definition 14** We say that a feasible point of Equation 40 satisfies the Mangasarian-Fromovitz Constraint Qualifications if there exists  $\mathbf{v} \in \mathbb{R}^p$  such that for all  $i \in [m]$  with  $1 - y_i f(\mathbf{x}_i; \boldsymbol{\theta}) = 0$  and for all  $\mathbf{h} \in \partial(1 - y_i f(\mathbf{x}_i; \boldsymbol{\theta}))$ , it holds:

$$\langle \mathbf{v}, \mathbf{h} \rangle > 0. \quad (42)$$

The next Lemma shows that Problem 40 satisfies the Mangasarian-Fromovitz Constraint Qualifications:

**Lemma 15** *Problem 40 satisfies the Mangasarian-Fromovitz Constraint Qualifications at every feasible point  $\boldsymbol{\theta}$ .*

**Proof** Let  $\mathbf{h}_i \in \partial(1 - y_i f(\mathbf{x}_i; \boldsymbol{\theta}))$  and  $\mathbf{v} = -\boldsymbol{\theta}$ , then for all  $i \in [m]$  satisfying  $y_i f(\mathbf{x}_i; \boldsymbol{\theta}) = 1$ , we have from Euler's theorem for homogeneous functions:

$$\langle \mathbf{v}, \mathbf{h}_i \rangle = L y_i f(\mathbf{x}_i; \boldsymbol{\theta}) = L > 0. \quad (43)$$

■

Our proof technique depends on the following relaxed notion of stationarity.

**Definition 16** ( $(\epsilon, \delta)$ -approximate KKT point) *A feasible point  $\boldsymbol{\theta}$  of Problem (40) is called an  $(\epsilon, \delta)$ -approximate KKT point if there exist  $\lambda_1, \dots, \lambda_m \geq 0$ ,  $\mathbf{h}_i \in \partial f(\mathbf{x}_i; \boldsymbol{\theta})$  and  $\mathbf{k} \in \partial \frac{1}{2} \|\boldsymbol{\theta}\|^2$  such that:*

1.  $\|\sum_{i=1}^m \lambda_i y_i \mathbf{h}_i - \mathbf{k}\|_2 \leq \epsilon$ .
2.  $\sum_{i=1}^m \lambda_i (y_i f(\mathbf{x}_i; \boldsymbol{\theta}) - 1) \leq \delta$ .

Before we proceed with the main result, we state and prove three useful Lemmata. The first one lower bounds the alignment between normalized iterates and normalized loss gradients. This Lemma is key for showing that the alignment goes to 1 as  $t \rightarrow \infty$ .

**Lemma 17** *For all  $t_2 > t_1 \geq t_0$ , there exists  $t_\star \in (t_1, t_2)$  such that:*

$$\left( \frac{1}{\left\langle \frac{\boldsymbol{\theta}_{t_\star}}{\|\boldsymbol{\theta}_{t_\star}\|}, \frac{-\mathbf{g}_{t_\star}^\star}{\|\mathbf{g}_{t_\star}^\star\|_\star} \right\rangle} - 1 \right) \leq \frac{1}{L} \frac{\log \frac{\tilde{\gamma}(t_2)}{\tilde{\gamma}(t_1)}}{\int_{t_1}^{t_2} \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|} dt}, \quad (44)$$

for all  $\mathbf{g}_{t_\star}^\star \in \operatorname{argmin}_{\mathbf{u} \in \partial \mathcal{L}(\boldsymbol{\theta}_{t_\star})} \|\mathbf{u}\|_\star$

**Proof** From Lemma 9, we have for all  $\mathbf{g}_t^\star \in \operatorname{argmin}_{\mathbf{u} \in \partial \mathcal{L}(\boldsymbol{\theta}_t)} \|\mathbf{u}\|_\star$ :

$$\begin{aligned} \frac{d \log \tilde{\gamma}}{dt} &\geq L \left\| \frac{d\boldsymbol{\theta}_t}{dt} \right\|^2 \left( \frac{1}{\langle \boldsymbol{\theta}_t, -\mathbf{g}_t^\star \rangle} - \frac{1}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}_t}{dt} \right\|} \right) \\ &= L \frac{\left\| \frac{d\boldsymbol{\theta}_t}{dt} \right\|}{\|\boldsymbol{\theta}_t\|} \left( \frac{1}{\left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^\star(\boldsymbol{\theta}_t)}{\|\mathbf{g}_t^\star\|_\star} \right\rangle} - 1 \right). \end{aligned} \quad (45)$$

We then integrate the two sides from  $t_1$  to  $t_2 > t_1 > t_0$ :

$$\int_{t_1}^{t_2} \left( \frac{1}{\left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^\star}{\|\mathbf{g}_t^\star\|_\star} \right\rangle} - 1 \right) \frac{\left\| \frac{d\boldsymbol{\theta}_t}{dt} \right\|}{\|\boldsymbol{\theta}_t\|} dt \leq \frac{1}{L} \log \frac{\tilde{\gamma}(t_2)}{\tilde{\gamma}(t_1)}. \quad (46)$$

The desired existential statement follows from a proof by contradiction.  $\blacksquare$

Next, we bound the rate of change of the normalized iterates.

**Lemma 18** *For almost any  $t > 0$ , it holds:*

$$\left\| \frac{d \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}}{dt} \right\| \leq 2 \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|}. \quad (47)$$

**Proof** The rate of change of the normalized iterates can be written as follows:

$$\begin{aligned} \frac{d \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}}{dt} &= \frac{1}{\|\boldsymbol{\theta}_t\|} \frac{d\boldsymbol{\theta}_t}{dt} + \boldsymbol{\theta}_t \left( -\frac{1}{\|\boldsymbol{\theta}_t\|^2} \frac{d\|\boldsymbol{\theta}_t\|}{dt} \right) \\ &= \frac{1}{\|\boldsymbol{\theta}_t\|} \frac{d\boldsymbol{\theta}_t}{dt} + \boldsymbol{\theta}_t \left( -\frac{1}{\|\boldsymbol{\theta}_t\|^2} \left\langle \mathbf{n}_t, \frac{d\boldsymbol{\theta}_t}{dt} \right\rangle \right), \quad (\text{Chain rule}) \end{aligned} \quad (48)$$

where  $\mathbf{n}_t \in \partial\|\boldsymbol{\theta}_t\|$ . So, by the triangle inequality, its norm is bounded by:

$$\begin{aligned} \left\| \frac{d \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}}{dt} \right\| &\leq \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|} + \frac{1}{\|\boldsymbol{\theta}_t\|} \left| \left\langle \mathbf{n}_t, \frac{d\boldsymbol{\theta}_t}{dt} \right\rangle \right| \\ &\leq 2 \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|}. \quad (\text{definition of dual norm and } \|\mathbf{n}_t\|_* \leq 1) \end{aligned} \quad (49)$$

$\blacksquare$

Finally, we prove an auxiliary lemma that constructs a convenient sequence.

**Lemma 19** *For any limit point  $\bar{\boldsymbol{\theta}}$  of  $\left\{ \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|} \right\}_{t \geq 0}$ , there exists a sequence  $t_n$  with  $t_n \rightarrow \infty$  so that  $\frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|} \rightarrow \bar{\boldsymbol{\theta}}$  and  $-\frac{\mathbf{g}_{t_n}^*}{\|\mathbf{g}_{t_n}^*\|_*} \rightarrow \mathbf{u}^*$  for some  $\mathbf{u}^* \in \partial\|\bar{\boldsymbol{\theta}}\|$ .*

**Proof** Let  $\epsilon_n = \frac{1}{n}$  for  $n \geq 1$ . We construct a sequence  $\{t_n\}_{n \geq 1}$ , by induction, in the following sense. Suppose  $t_1 < \dots < t_{n-1}$  have been constructed already. Since  $\bar{\boldsymbol{\theta}}$  is a limit point of the normalized iterates and  $\log \tilde{\gamma}_t \rightarrow \log \tilde{\gamma}_\infty < \infty$  (as  $\tilde{\gamma}_t$  is non-decreasing and bounded from above), there exists  $s_n > t_{n-1}$  such that:

$$\left\| \frac{\boldsymbol{\theta}_{s_n}}{\|\boldsymbol{\theta}_{s_n}\|} - \bar{\boldsymbol{\theta}} \right\| \leq \epsilon_n = \frac{1}{n} \quad \text{and} \quad \frac{1}{L} \log \frac{\tilde{\gamma}_\infty}{\tilde{\gamma}_{s_n}} \leq \epsilon_n^2 = \frac{1}{n^2}. \quad (50)$$

Since  $\frac{d \log \|\boldsymbol{\theta}_t\|}{dt} \leq \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|}$  (as in Equation 26), we have that  $\lim_{t \rightarrow \infty} \int_{t_A}^t \frac{\|d\boldsymbol{\theta}_{t'}\|}{\|\boldsymbol{\theta}_{t'}\|} dt' \geq \lim_{t \rightarrow \infty} \log \frac{\|\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_{t_A}\|} = \infty$  for all  $t_A > 0$ . Thus, there exists  $s'_n > s_n$  such that  $\int_{s_n}^{s'_n} \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|} dt = \frac{1}{n}$ .

Now, from Lemma 17, we know there exists  $t_n \in (s_n, s'_n)$  with:

$$\left( \frac{1}{\left\langle \frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|}, \frac{-\mathbf{g}_{t_n}^*}{\|\mathbf{g}_{t_n}^*\|_*} \right\rangle} - 1 \right) \leq \frac{1}{L} \frac{\log \frac{\tilde{\gamma}_{s'_n}}{\tilde{\gamma}_{s_n}}}{\int_{s_n}^{s'_n} \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|} dt} \leq \frac{\frac{1}{n^2}}{\frac{1}{n}} = \frac{1}{n}, \quad (51)$$

which implies  $\left\langle \frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|}, \frac{-\mathbf{g}_{t_n}^*}{\|\mathbf{g}_{t_n}^*\|_*} \right\rangle \geq \frac{1}{1+\frac{1}{n}} \rightarrow 1$  as  $n \rightarrow \infty$ . From the definition of the dual norm,  $\left\langle \frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|}, \frac{-\mathbf{g}_{t_n}^*}{\|\mathbf{g}_{t_n}^*\|_*} \right\rangle \leq 1$ , hence  $\left\langle \frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|}, \frac{-\mathbf{g}_{t_n}^*}{\|\mathbf{g}_{t_n}^*\|_*} \right\rangle \rightarrow 1$  as  $n \rightarrow \infty$ . Also, for the normalized iterates we have:

$$\left\| \frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|} - \bar{\boldsymbol{\theta}} \right\| \leq \left\| \frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|} - \frac{\boldsymbol{\theta}_{s_n}}{\|\boldsymbol{\theta}_{s_n}\|} \right\| + \left\| \frac{\boldsymbol{\theta}_{s_n}}{\|\boldsymbol{\theta}_{s_n}\|} - \bar{\boldsymbol{\theta}} \right\| \stackrel{\text{Equation 50}}{\leq} \left\| \frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|} - \frac{\boldsymbol{\theta}_{s_n}}{\|\boldsymbol{\theta}_{s_n}\|} \right\| + \frac{1}{n} \quad (52)$$

To deal with the first term, we can leverage Lemma 18 which bounds the rate of change of the normalized iterates:

$$\left\| \frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|} - \bar{\boldsymbol{\theta}} \right\| \leq \int_{s_n}^{t_n} \left\| \frac{d}{dt} \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|} \right\| dt + \frac{1}{n} \leq 2 \int_{s_n}^{t_n} \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|} dt + \frac{1}{n} \leq 2 \int_{s_n}^{s'_n} \frac{\|d\boldsymbol{\theta}_t\|}{\|\boldsymbol{\theta}_t\|} dt + \frac{1}{n} = \frac{3}{n} \rightarrow 0 \quad (53)$$

From Bolzano-Weierstrass, let  $t_{n_\ell}$  be a subsequence of  $t_n$  so that  $\frac{-\mathbf{g}_{t_{n_\ell}}^*}{\|\mathbf{g}_{t_{n_\ell}}^*\|_*}$  converges to some limit point  $\mathbf{u}^*$ . This limit point therefore satisfies  $\langle \mathbf{u}^*, \bar{\boldsymbol{\theta}} \rangle = 1, \|\mathbf{u}^*\|_* = 1$  which implies  $\mathbf{u}^* \in \partial\|\bar{\boldsymbol{\theta}}\|$ . The required sequence is  $t_{n_\ell}$ .  $\blacksquare$

We are now ready to state and prove our main result.

**Theorem 20** *For steepest flow (Equation 3) on the exponential loss, under Assumptions (A1), (A2), (A3), any limit point  $\bar{\boldsymbol{\theta}}$  of  $\left\{ \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|} \right\}_{t \geq 0}$  is along the direction of a KKT point of the following optimization problem:*

$$\begin{aligned} & \min_{\boldsymbol{\theta} \in \mathbb{R}^p} \frac{1}{2} \|\boldsymbol{\theta}\|^2 \\ & \text{s.t. } y_i f(\mathbf{x}_i; \boldsymbol{\theta}) \geq 1, \forall i \in [m]. \end{aligned} \quad (54)$$

**Proof**

For convenience, denote

$$q_{\min}(\boldsymbol{\theta}) = \min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}) = \|\boldsymbol{\theta}\|^L \gamma(t) \quad (55)$$

We prove that

$$\hat{\boldsymbol{\theta}} := \frac{\bar{\boldsymbol{\theta}}}{q_{\min}(\bar{\boldsymbol{\theta}})^{1/L}} \quad (56)$$

is a KKT point, by showing that  $\hat{\boldsymbol{\theta}}$  is a  $(\varepsilon_\ell, \delta_\ell)$ -approximate KKT point (Def. 16) for sequences  $\varepsilon_\ell \rightarrow 0, \delta_\ell \rightarrow 0$ .

From Lemma 19, let  $t_n$  be a sequence with  $\frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|} \rightarrow \bar{\boldsymbol{\theta}}$  and  $-\frac{\mathbf{g}_{t_n}^*}{\|\mathbf{g}_{t_n}^*\|_*} \rightarrow \mathbf{u}^*$  for some  $\mathbf{u}^* \in \partial\|\bar{\boldsymbol{\theta}}\|$ . Denote  $\tilde{\boldsymbol{\theta}}_{t_n} = \frac{\boldsymbol{\theta}_{t_n}}{q_{\min}(\boldsymbol{\theta}_{t_n})^{1/L}}$ . Notice that by continuity of  $f$ ,

$$\lim_{n \rightarrow \infty} \tilde{\boldsymbol{\theta}}_{t_n} = \lim_{n \rightarrow \infty} \frac{\boldsymbol{\theta}_{t_n}}{\left(\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_n})\right)^{1/L}} = \lim_{n \rightarrow \infty} \frac{\frac{\boldsymbol{\theta}_{t_n}}{\|\boldsymbol{\theta}_{t_n}\|}}{\frac{1}{\|\boldsymbol{\theta}_{t_n}\|} \left(\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_n})\right)^{1/L}} = \hat{\boldsymbol{\theta}}. \quad (57)$$

By the chain rule, denote  $\mathbf{h}_i^{(t_n)} \in \partial f(\mathbf{x}_i; \boldsymbol{\theta}_{t_n})$  so that

$$\mathbf{g}_{t_n}^* = - \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_n})} y_i \mathbf{h}_i^{(t_n)} \quad (58)$$

Define  $\tilde{\mathbf{h}}_i^{(t_n)} = q_{\min}(\boldsymbol{\theta}_{t_n})^{\frac{1}{L}-1} \mathbf{h}_i^{(t_n)}$ . Since  $\tilde{\boldsymbol{\theta}}_{t_n} = q_{\min}(\boldsymbol{\theta}_{t_n})^{-\frac{1}{L}} \boldsymbol{\theta}_{t_n}$ , it holds that  $\tilde{\mathbf{h}}_i^{(t_n)} = q_{\min}(\boldsymbol{\theta}_{t_n})^{(-\frac{1}{L})(L-1)} \mathbf{h}_i^{(t_n)} \in \partial f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}_{t_n})$  by Theorem B.2(a) (sub-derivatives of homogeneous functions) in (Lyu and Li, 2020). Since  $f$  is locally Lipschitz, it holds that  $\partial f(\mathbf{x}_i; \cdot)$  is bounded and has a closed graph around  $\hat{\boldsymbol{\theta}}$ , so for every  $i \in [m]$  there exists a subsequence  $t_\ell$  of  $t_n$  with  $\tilde{\mathbf{h}}_i^{(t_\ell)} \rightarrow \hat{\mathbf{h}}_i$  for some  $\hat{\mathbf{h}}_i \in \partial f(\mathbf{x}_i; \hat{\boldsymbol{\theta}})$ . By iteratively choosing subsequences for  $i = 1, \dots, m$  we may assume that  $\tilde{\mathbf{h}}_i^{(t_\ell)} \rightarrow \hat{\mathbf{h}}_i$  for all  $i \in [m]$ .

Set  $\lambda_{\ell, i} = \frac{\|\boldsymbol{\theta}_{t_\ell}\|}{\|\mathbf{g}_{t_\ell}^*\|_*} q_{\min}^{1-\frac{2}{L}} e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell})} \geq 0$ . Plugging in,

$$\sum_{i=1}^m \lambda_{\ell, i} y_i \tilde{\mathbf{h}}_i^{(t_\ell)} = \sum_{i=1}^m \lambda_{\ell, i} q_{\min}^{\frac{1}{L}-1} y_i \mathbf{h}_i^{(t_\ell)} = \frac{\|\boldsymbol{\theta}_{t_\ell}\|}{q_{\min}^{1/L} \|\mathbf{g}_{t_\ell}^*\|_*} \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell})} y_i \mathbf{h}_i^{(t_\ell)} = - \|\tilde{\boldsymbol{\theta}}_{t_\ell}\| \frac{\mathbf{g}_{t_\ell}^*}{\|\mathbf{g}_{t_\ell}^*\|_*} \quad (59)$$

Notice that since  $\mathbf{u}^* \in \partial\|\bar{\boldsymbol{\theta}}\| = \partial\|\hat{\boldsymbol{\theta}}\|$ , we have  $\|\hat{\boldsymbol{\theta}}\| \mathbf{u}^* \in \partial \frac{1}{2} \|\hat{\boldsymbol{\theta}}\|^2$  as a primal vector, so we evaluate the following:

$$\varepsilon_\ell := \left\| \sum_{i=1}^m \lambda_{\ell, i} y_i \hat{\mathbf{h}}_i - \|\hat{\boldsymbol{\theta}}\| \mathbf{u}^* \right\|_2 \leq \left\| \sum_{i=1}^m \lambda_{\ell, i} y_i (\hat{\mathbf{h}}_i - \tilde{\mathbf{h}}_i^{(t_\ell)}) \right\|_2 + \left\| \sum_{i=1}^m \lambda_{\ell, i} y_i \tilde{\mathbf{h}}_i^{(t_\ell)} - \|\hat{\boldsymbol{\theta}}\| \mathbf{u}^* \right\|_2 \quad (60)$$

The second term goes to 0 since  $-\frac{\mathbf{g}_{t_\ell}^*}{\|\mathbf{g}_{t_\ell}^*\|_*} \rightarrow \mathbf{u}^*$ . To show the first term goes to 0, we use  $\tilde{\mathbf{h}}_i^{(t_\ell)} \rightarrow \hat{\mathbf{h}}_i$  and show  $\lambda_{\ell, i}$  are bounded. Recall from Theorem 9 that  $\|\mathbf{g}_{t_\ell}^*\|_* \geq \frac{1}{\|\boldsymbol{\theta}_{t_\ell}\|} L \mathcal{L} \log\left(\frac{1}{\mathcal{L}}\right)$ , and by definition  $\log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} = \tilde{\gamma}(\boldsymbol{\theta}_t) \|\boldsymbol{\theta}_t\|^L$ . Also, recall that  $\lim_{t \rightarrow \infty} \gamma(t) = \lim_{t \rightarrow \infty} \tilde{\gamma}(t) = \tilde{\gamma}_\infty$ . Altogether:

$$\begin{aligned} |\lambda_{\ell, i}| &\leq \|\boldsymbol{\theta}_{t_\ell}\|^2 \frac{1}{L \mathcal{L} \log\left(\frac{1}{\mathcal{L}}\right)} \cdot \|\boldsymbol{\theta}_{t_\ell}\|^{L(1-\frac{2}{L})} \cdot \gamma(t_\ell)^{1-\frac{2}{L}} \mathcal{L} = \frac{\gamma(t_\ell)^{1-\frac{2}{L}} \|\boldsymbol{\theta}_{t_\ell}\|^L}{L \log \frac{1}{\mathcal{L}}} \\ &= \frac{\gamma(t_\ell)^{1-\frac{2}{L}}}{L} \cdot \frac{\|\boldsymbol{\theta}_{t_\ell}\|^L}{\|\boldsymbol{\theta}_{t_\ell}\|^L \tilde{\gamma}(t_\ell)} \\ &\xrightarrow{\ell \rightarrow \infty} \frac{\tilde{\gamma}_\infty^{-\frac{2}{L}}}{L} < \infty. \end{aligned} \quad (61)$$

For  $\delta_\ell$ ,

$$\delta_\ell = \left| \sum_{i=1}^m \lambda_{\ell,i} (y_i f(\mathbf{x}_i; \hat{\boldsymbol{\theta}}) - 1) \right| \leq \left| \sum_{i=1}^m \lambda_{\ell,i} (y_i f(\mathbf{x}_i; \hat{\boldsymbol{\theta}}) - y_i f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}_{t_\ell})) \right| + \left| \sum_{i=1}^m \lambda_{\ell,i} (y_i f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}_{t_\ell}) - 1) \right| \quad (62)$$

The first term goes to 0 from continuity of  $f$  and boundedness of  $\lambda_{\ell,i}$ . As for the second term, note that from Lemma 13, each summand is non-negative. Plugging in,

$$\begin{aligned} \sum_{i=1}^m \lambda_{\ell,i} (y_i f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}_{t_\ell}) - 1) &= \frac{\|\boldsymbol{\theta}_{t_\ell}\|}{\|\mathbf{g}_{t_\ell}^*\|_\star} \sum_{i=1}^m q_{\min}^{1-\frac{2}{L}} e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell})} \left( \frac{y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell})}{q_{\min}} - 1 \right) \\ &= \frac{\|\boldsymbol{\theta}_{t_\ell}\|}{q_{\min} \|\mathbf{g}_{t_\ell}^*\|_\star} \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell}) - q_{\min}) \end{aligned} \quad (63)$$

Again we use  $\|\mathbf{g}_{t_\ell}^*\|_\star \geq \frac{L}{\|\boldsymbol{\theta}_{t_\ell}\|} \mathcal{L} \log \frac{1}{\mathcal{L}} \geq \frac{L}{\|\boldsymbol{\theta}_{t_\ell}\|} e^{-q_{\min}} \log \frac{1}{\mathcal{L}}$ :

$$\begin{aligned} \sum_{i=1}^m \lambda_{\ell,i} (y_i f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}_{t_\ell}) - 1) &\leq \frac{\|\boldsymbol{\theta}_{t_\ell}\|^2}{q_{\min}^{2/L} L e^{-q_{\min}} \log \frac{1}{\mathcal{L}}} \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell}) - q_{\min}) \\ &= \frac{\|\boldsymbol{\theta}_{t_\ell}\|^2}{q_{\min}^{2/L} L \log \frac{1}{\mathcal{L}}} \sum_{i=1}^m e^{-(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell}) - q_{\min})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell}) - q_{\min}) \\ &\stackrel{\text{Equation 55}}{=} \frac{1}{\gamma(t_\ell)^{2/L} L \log \frac{1}{\mathcal{L}}} \sum_{i=1}^m e^{-(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell}) - q_{\min})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell}) - q_{\min}) \\ &\stackrel{\text{Lemma 11}}{\leq} \frac{1}{\tilde{\gamma}(t)^{2/L} L \log \frac{1}{\mathcal{L}}} \sum_{i=1}^m e^{-(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell}) - q_{\min})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}_{t_\ell}) - q_{\min}) \\ &\leq \frac{m}{e \tilde{\gamma}(t_0)^{2/L} L \log \frac{1}{\mathcal{L}}} \\ &\rightarrow 0 \end{aligned} \quad (64)$$

In the last inequality we used monotonicity of  $\tilde{\gamma}(t)$ , and the fact  $\forall z \geq 0 : z e^{-z} \leq 1/e$ . The limit holds since  $\mathcal{L} \xrightarrow{t \rightarrow \infty} 0$  and  $t_\ell \rightarrow \infty$ . By Theorem C.4 in Lyu and Li (2020) (which is itself based on a result due to Dutta et al. (2013)), we get that  $\frac{\bar{\boldsymbol{\theta}}}{(\min_{i \in [m]} y_i f(\mathbf{x}_i; \bar{\boldsymbol{\theta}}))^{\frac{1}{L}}}$  is a KKT point of Equation 40.  $\blacksquare$

### A.3 Generalization to other losses

The previous results can be generalized to any loss with exponential tails. In particular, let us proceed to the following definition:

**Definition 21** Let  $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ . Assume that  $\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^m e^{-\Phi(y_i f(\mathbf{x}_i; \boldsymbol{\theta}))}$ ,  $\boldsymbol{\theta} \in \mathbb{R}^p$ , where  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ ,  $y_i \in \{\pm 1\}$ . We call the function  $l : \mathbb{R} \rightarrow \mathbb{R}$ ,  $l(u) := e^{-\Phi(u)}$ , exponentially tailed, if the following conditions hold:

(i)  $\Phi$  is continuously differentiable.

(ii)  $\Phi'(u) > 0$  for all  $u \in \mathbb{R}$ .

(iii) The function  $g(u) = \Phi'(u)u$  is non-decreasing in  $[0, \infty)$ .

Notice that the definition above covers the exponential loss for  $\Phi(u) = u$  and the logistic loss for  $\Phi(u) = -\log \log(1 + e^{-u})$ . To accommodate different loss functions, Assumption A3 needs to be adjusted as follows:

- There is a  $t_0 > 0$ , such that  $\mathcal{L}(\boldsymbol{\theta}_{t_0}) < e^{-\Phi(0)}$ .

See Section A in (Lyu and Li, 2020) for a more general, albeit technical, definition that allows the extension of the full analysis to general exponentially-tailed losses.

Under these conditions, we can define the soft margin as follows:

$$\tilde{\gamma} = \frac{l^{-1}(\mathcal{L})}{\|\boldsymbol{\theta}\|^L} = \frac{\Phi^{-1}\left(\log \frac{1}{\mathcal{L}}\right)}{\|\boldsymbol{\theta}\|^L},$$

and prove a strict generalization of Theorem 1.

**Theorem 22 (Soft margin increases - general loss function)** *For almost any  $t > t_0$ , it holds:*

$$\frac{d \log \tilde{\gamma}}{dt} \geq L \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{(\Phi^{-1})' \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)}{L \mathcal{L}(\boldsymbol{\theta}_t) \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)} - \frac{1}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \geq 0.$$

**Proof** Let  $\mathbf{n}_t \in \partial \|\boldsymbol{\theta}_t\|$ . We have:

$$\begin{aligned} \frac{d \log \tilde{\gamma}}{dt} &= \frac{d}{dt} \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right) - L \frac{d}{dt} \log \|\boldsymbol{\theta}_t\| \\ &= \frac{d}{dt} \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right) - L \left\langle \frac{\mathbf{n}_t}{\|\boldsymbol{\theta}_t\|}, \frac{d\boldsymbol{\theta}}{dt} \right\rangle \quad (\text{Chain rule}) \\ &\geq \frac{d}{dt} \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right) - L \frac{\left\| \frac{d\boldsymbol{\theta}}{dt} \right\|}{\|\boldsymbol{\theta}_t\|} \quad (\text{definition of dual norm and } \|\mathbf{n}_t\|_* \leq 1) \\ &= -\frac{d\mathcal{L}(\boldsymbol{\theta}_t)}{dt} \frac{(\Phi^{-1})' \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)}{\mathcal{L}(\boldsymbol{\theta}_t) \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)} - L \frac{\left\| \frac{d\boldsymbol{\theta}}{dt} \right\|}{\|\boldsymbol{\theta}_t\|} \quad (\text{Chain rule}) \\ &= \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{(\Phi^{-1})' \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)}{\mathcal{L}(\boldsymbol{\theta}_t) \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \quad (\text{Equation 18}). \end{aligned} \tag{65}$$

But, the first term inside the parenthesis can be related to the second one via the following calculation. Recall that, by the chain rule for locally Lipschitz functions (Theorem 6),

for any  $\mathbf{g}_t \in \partial\mathcal{L}(\boldsymbol{\theta}_t)$  there exist  $\mathbf{h}_1 \in \partial y_1 f(\mathbf{x}_1; \boldsymbol{\theta}_t), \dots, \mathbf{h}_m \in \partial y_m f(\mathbf{x}_m; \boldsymbol{\theta}_t)$  such that  $\mathbf{g}_t = \sum_{i=1}^m e^{-\Phi(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))} \Phi'(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)) \mathbf{h}_i$ . Thus, for a minimum norm subderivative  $\mathbf{g}_t^*$ , we have:

$$\begin{aligned} \langle \boldsymbol{\theta}_t, -\mathbf{g}_t^* \rangle &= \left\langle \boldsymbol{\theta}_t, \sum_{i=1}^m e^{-\Phi(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))} \Phi'(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)) \mathbf{h}_i \right\rangle \\ &= \sum_{i=1}^m e^{-\Phi(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))} \Phi'(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)) \langle \boldsymbol{\theta}_t, \mathbf{h}_i^* \rangle \\ &= L \sum_{i=1}^m e^{-\Phi(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))} \Phi'(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)) y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t), \end{aligned} \quad (66)$$

where the last equality follows from Euler's theorem for homogeneous functions (whose generalization for subderivatives can be found in Theorem B.2 in Lyu and Li (2020)). But, now observe that as per assumption,  $u \rightarrow \Phi'(u)u$  is non-decreasing and this last term can be lower bounded as:

$$\langle \boldsymbol{\theta}_t, -\mathbf{g}_t^* \rangle \geq L \sum_{i=1}^m e^{-\Phi(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))} \Phi' \left( \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right) \right) \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right), \quad (67)$$

where we used the fact  $y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t) \leq \Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)$  for all  $i \in [m]$  (by the monotonicity of  $\Phi$  the definition of  $\mathcal{L}$ ). Leveraging the fundamental property between the derivative of a function and its inverse's, we further get:

$$\langle \boldsymbol{\theta}_t, -\mathbf{g}_t^* \rangle \geq L \sum_{i=1}^m e^{-\Phi(y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))} \frac{\Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)}{(\Phi^{-1})' \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)} = L \mathcal{L}(\boldsymbol{\theta}_t) \frac{\Phi^{-1} \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)}{(\Phi^{-1})' \left( \log \frac{1}{\mathcal{L}(\boldsymbol{\theta}_t)} \right)}. \quad (68)$$

We have made the first term of Equation 65 appear. By plugging Equation 68 into Equation 65, we get:

$$\begin{aligned} \frac{d \log \tilde{\gamma}}{dt} &\geq \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{L}{\langle \boldsymbol{\theta}_t, -\mathbf{g}_t^* \rangle} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \\ &\geq \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|^2 \left( \frac{L}{\|\boldsymbol{\theta}_t\| \|\mathbf{g}_t^*\|_*} - \frac{L}{\|\boldsymbol{\theta}_t\| \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|} \right) \quad (\text{definition of dual norm}). \end{aligned} \quad (69)$$

Noticing that  $\|\mathbf{g}_t^*\|_* = \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|$  (from Proposition 7) concludes the proof.  $\blacksquare$

## Appendix B. Approximate KKT points at finite time through Bregman divergences

In this section, we further analyze the late-phase geometric properties of the trajectory and, in particular, explain how the alignment  $\left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^*}{\|\mathbf{g}_t^*\|_*} \right\rangle$  can capture a notion of proximity to stationarity. For algorithms whose norm squared is smooth, this analysis culminates

in showing that the (finite-time) iterates are approximate KKT points of the margin maximization problem.

Our analysis uses core ideas from the theory of conjugate functions and Fenchel's duality.

**Definition 23 (Convex conjugate)** Let  $\psi : \mathbb{R}^p \rightarrow \mathbb{R}$ . We denote by  $\psi^*(\cdot)$  the convex conjugate of  $\psi(\cdot)$ :

$$\psi^*(\boldsymbol{\omega}) = \sup_{\boldsymbol{\theta} \in \mathbb{R}^p} \{\langle \boldsymbol{\omega}, \boldsymbol{\theta} \rangle - \psi(\boldsymbol{\theta})\}. \quad (70)$$

We will make use of the following properties of conjugate functions.

**Proposition 24 (Conjugate subgradient theorem - Theorem 23.5 in Rockafellar (1970), Theorem 4.20 in Beck (2017))** Let  $\psi : \mathbb{R}^p \rightarrow \mathbb{R}$  be convex and closed. For any  $\boldsymbol{\theta}^* \in \partial\psi^*(\boldsymbol{\theta})$ , it holds  $\partial\psi(\boldsymbol{\theta}^*) \ni \boldsymbol{\theta}$ .

**Lemma 25 (Fenchel-Young inequality) (Fenchel, 1949)** For any  $\psi : \mathbb{R}^p \rightarrow \mathbb{R}$  and  $\boldsymbol{\omega}, \boldsymbol{\theta} \in \mathbb{R}^p$ , it holds:

$$\langle \boldsymbol{\theta}, \boldsymbol{\omega} \rangle \leq \psi(\boldsymbol{\theta}) + \psi^*(\boldsymbol{\omega}). \quad (71)$$

We define the following generalized *Bregman divergence*.

**Definition 26 (Generalized Bregman divergence)** Let  $\psi : \mathbb{R}^p \rightarrow \mathbb{R}$  with  $\psi(\boldsymbol{\theta}) = \frac{1}{2}\|\boldsymbol{\theta}\|_\star^2$  for all  $\boldsymbol{\theta} \in \mathbb{R}^p$ . We define the (generalized) Bregman divergence  $D_{\frac{1}{2}\|\cdot\|_\star^2}^{\mathbf{m}}(\cdot, \cdot) : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$  induced by  $\psi$  as follows:

$$D_{\frac{1}{2}\|\cdot\|_\star^2}^{\mathbf{m}}(\mathbf{y}, \mathbf{z}) = \frac{1}{2}\|\mathbf{y}\|_\star^2 - \frac{1}{2}\|\mathbf{z}\|_\star^2 - \langle \mathbf{m}, \mathbf{y} - \mathbf{z} \rangle, \quad (72)$$

where  $\mathbf{m} \in \partial_{\frac{1}{2}}\|\mathbf{z}\|_\star^2$ .

**Remark 27** Notice that if the function  $\psi(\boldsymbol{\theta}) = \frac{1}{2}\|\boldsymbol{\theta}\|_\star^2$  is differentiable, then the subdifferential defined at any point collapses to a single element: the gradient of  $\psi$ . If, further,  $\psi$  is strictly convex, then Equation 72 coincides with the usual Bregman divergence induced by  $\psi$ , defined as  $D_\psi(\mathbf{y}, \mathbf{z}) = \psi(\mathbf{y}) - \psi(\mathbf{z}) - \langle \nabla\psi(\mathbf{z}), \mathbf{y} - \mathbf{z} \rangle$ . Bregman divergences (Bregman, 1967) generalize the Euclidean squared distance in different geometries and have found numerous applications in machine learning (A. Nemirovskii and D. Yudin, 1983; Banerjee et al., 2005).

Next, we derive guarantees for the finite-time iterates, akin to those of Definition 16, with the difference that the following ones are with respect to a generalized Bregman divergence.

**Proposition 28** For any  $t > t_0$ , let  $\tilde{\boldsymbol{\theta}}_t = \frac{\boldsymbol{\theta}_t}{(\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))^{\frac{1}{L}}}$ . There exist  $\lambda_i \geq 0, \tilde{\mathbf{h}}_i^* \in \partial f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}_t)$ , such that for any  $\tilde{\mathbf{k}} \in \partial_{\frac{1}{2}}\|\tilde{\boldsymbol{\theta}}_t\|^2$ , it holds:

$$\begin{aligned} D_{\frac{1}{2}\|\cdot\|_\star^2}^{\tilde{\boldsymbol{\theta}}_t} \left( \sum_{i=1}^m \lambda_i y_i \tilde{\mathbf{h}}_i, \tilde{\mathbf{k}} \right) &\leq \frac{1}{\tilde{\gamma}(t_0)^{\frac{2}{L}}} \left( 1 - \left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^*}{\|\mathbf{g}_t^*\|_\star} \right\rangle \right), \\ \sum_{i=1}^m \lambda_i \left( y_i f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}) - 1 \right) &\leq \frac{m}{e \tilde{\gamma}(t_0)^{\frac{2}{L}} L \log \frac{1}{L}}, \end{aligned} \quad (73)$$

with  $\mathbf{g}_t^* \in \operatorname{argmin}_{\mathbf{u} \in \partial\mathcal{L}(\boldsymbol{\theta}_t)} \|\mathbf{u}\|_\star$ .

**Proof** To simplify the notation, let  $q_{\min} = \min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)$ . Let  $\tilde{\mathbf{k}} \in \partial_{\frac{1}{2}} \|\tilde{\boldsymbol{\theta}}_t\|^2$ ,  $\mathbf{g}_t^* \in \operatorname{argmin}_{\mathbf{u} \in \partial \mathcal{L}(\boldsymbol{\theta}_t)} \|\mathbf{u}\|_*$  and  $\mathbf{h}_i^* \in \partial f(\mathbf{x}_i; \boldsymbol{\theta}_t)$ ,  $i \in [m]$ , such that  $\mathbf{g}_t^* = -\sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} y_i \mathbf{h}_i^*$  (whose existence is guaranteed from chain rule – Theorem 5). Finally, we define  $\tilde{\mathbf{h}}_i^* = q_{\min}^{\frac{1}{L}-1} \mathbf{h}_i^*$ ,  $i \in [m]$ , for which it holds:  $\tilde{\mathbf{h}}_i^* \in \partial f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}_t)$  from Theorem B.2(a) in Lyu and Li (2020). We set  $\lambda_i = \frac{\|\boldsymbol{\theta}_t\|}{\|\mathbf{g}_t^*\|_*} q_{\min}^{1-\frac{2}{L}} e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} \geq 0$ . Then, it holds:

$$\begin{aligned} \sum_{i=1}^m \lambda_i y_i \tilde{\mathbf{h}}_i^* &= \sum_{i=1}^m \lambda_i q_{\min}^{\frac{1}{L}-1} y_i \mathbf{h}_i^* \quad (\text{Thm B.2(a) in Lyu and Li (2020)}) \\ &= \frac{\|\boldsymbol{\theta}_t\|}{q_{\min}^{\frac{1}{L}} \|\mathbf{g}_t^*\|_*} \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t)} y_i \mathbf{h}_i^* = -\frac{\|\boldsymbol{\theta}_t\| \mathbf{g}_t^*}{q_{\min}^{\frac{1}{L}} \|\mathbf{g}_t^*\|_*}, \end{aligned} \quad (74)$$

which is a scaled version of the (minimum norm) subderivative of the loss.

Let  $\psi(\boldsymbol{\theta}_t) = \frac{1}{2} \|\boldsymbol{\theta}_t\|_*^2$  be the potential function that we shall use in order to define our divergence. For this specific  $\psi$ , it holds:  $\psi^*(\boldsymbol{\omega}) = \frac{1}{2} \|\boldsymbol{\omega}\|^2$  (see for instance Example 3.27 in Boyd and Vandenberghe (2014) for a derivation). Recall that in the definition of  $D_{\frac{1}{2} \|\cdot\|_*}^{\mathbf{m}}$  (Equation 72) there is an extra choice that we have to make; the one of the subderivative  $\mathbf{m}$ . In what follows, we will specifically measure “distance” between  $\sum_{i=1}^m \lambda_i y_i \tilde{\mathbf{h}}_i^*$  and  $\tilde{\mathbf{k}}$  using  $D_{\frac{1}{2} \|\cdot\|_*}^{\tilde{\boldsymbol{\theta}}_t}(\cdot, \cdot)$ , i.e. by picking  $\mathbf{m} = \tilde{\boldsymbol{\theta}}_t$ . This is possible, since from Proposition 24 it holds that  $\tilde{\boldsymbol{\theta}}_t \in \partial_{\frac{1}{2}} \|\tilde{\mathbf{k}}\|_*^2$ . Finally, let  $\mathbf{r} \in \partial \|\tilde{\boldsymbol{\theta}}_t\|$  be the subgradient of  $\|\cdot\|$  that stems from the chain rule of  $\frac{1}{2} \|\cdot\|_*^2$  evaluated at  $\tilde{\boldsymbol{\theta}}_t$ . We calculate the divergence between the two vectors:

$$\begin{aligned} &D_{\frac{1}{2} \|\cdot\|_*}^{\tilde{\boldsymbol{\theta}}_t} \left( \sum_{i=1}^m \lambda_i y_i \tilde{\mathbf{h}}_i^*, \tilde{\mathbf{k}} \right) \\ &= \frac{1}{2} \left\| -\frac{\|\boldsymbol{\theta}_t\| \mathbf{g}_t^*}{q_{\min}^{\frac{1}{L}} \|\mathbf{g}_t^*\|_*} \right\|_*^2 - \frac{1}{2} \|\tilde{\mathbf{k}}\|_*^2 - \left\langle \frac{\boldsymbol{\theta}_t}{q_{\min}^{\frac{1}{L}}}, -\frac{\|\boldsymbol{\theta}_t\| \mathbf{g}_t^*}{q_{\min}^{\frac{1}{L}} \|\mathbf{g}_t^*\|_*} - \tilde{\mathbf{k}} \right\rangle \\ &= \frac{1}{2} \frac{\|\boldsymbol{\theta}_t\|^2}{q_{\min}^{\frac{2}{L}}} - \frac{1}{2} \|\tilde{\boldsymbol{\theta}}_t\|_*^2 - \left\langle \frac{\boldsymbol{\theta}_t}{q_{\min}^{\frac{1}{L}}}, \frac{-\|\boldsymbol{\theta}_t\| \mathbf{g}_t^*}{q_{\min}^{\frac{1}{L}} \|\mathbf{g}_t^*\|_*} - \|\tilde{\boldsymbol{\theta}}_t\| \mathbf{r} \right\rangle \quad (\text{Chain rule}) \\ &= \frac{\|\boldsymbol{\theta}_t\|^2}{q_{\min}^{\frac{2}{L}}} \left( \frac{1}{2} - \frac{1}{2} \|\mathbf{r}\|_*^2 - \left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^*}{\|\mathbf{g}_t^*\|_*} \right\rangle + \left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \mathbf{r} \right\rangle \right) \\ &\leq \frac{\|\boldsymbol{\theta}_t\|^2}{q_{\min}^{\frac{2}{L}}} \left( \frac{1}{2} - \frac{1}{2} \|\mathbf{r}\|_*^2 - \left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^*}{\|\mathbf{g}_t^*\|_*} \right\rangle + \frac{1}{2} \left\| \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|} \right\|^2 + \frac{1}{2} \|\mathbf{r}\|_*^2 \right) \quad (\text{Equation 71}) \\ &= \frac{\|\boldsymbol{\theta}_t\|^2}{q_{\min}^{\frac{2}{L}}} \left( 1 - \left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^*}{\|\mathbf{g}_t^*\|_*} \right\rangle \right) \\ &\leq \frac{1}{\tilde{\gamma}^{\frac{2}{L}}} \left( 1 - \left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^*}{\|\mathbf{g}_t^*\|_*} \right\rangle \right) \leq \frac{1}{\tilde{\gamma}(t_0)^{\frac{2}{L}}} \left( 1 - \left\langle \frac{\boldsymbol{\theta}_t}{\|\boldsymbol{\theta}_t\|}, \frac{-\mathbf{g}_t^*}{\|\mathbf{g}_t^*\|_*} \right\rangle \right), \end{aligned} \quad (75)$$

where the last 2 inequalities follow from the relation between soft and hard margin (Lemma 11), and the monotonicity of the former.

For the second condition, we have:

$$\begin{aligned} \sum_{i=1}^m \lambda_i \left( y_i f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}) - 1 \right) &= \frac{\|\boldsymbol{\theta}\|}{\|\mathbf{g}_t^*\|_\star} \sum_{i=1}^m q_{\min}^{1-\frac{2}{L}} e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta})} \left( \frac{y_i f(\mathbf{x}_i; \boldsymbol{\theta})}{q_{\min}} - 1 \right) \\ &= \frac{\|\boldsymbol{\theta}\|}{q_{\min}^{\frac{2}{L}} \|\mathbf{g}_t^*\|_\star} \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}) - q_{\min}). \end{aligned} \quad (76)$$

From Equation 30 and Equation 28, we can lower bound the dual norm of the subderivate:

$$\|\mathbf{g}_t^*\|_\star \geq \frac{L}{\|\boldsymbol{\theta}\|} \mathcal{L} \log \frac{1}{\mathcal{L}} \geq \frac{L}{\|\boldsymbol{\theta}\|} e^{-q_{\min}} \log \frac{1}{\mathcal{L}}. \quad (77)$$

By plugging in back to Equation 76, we obtain

$$\begin{aligned} \sum_{i=1}^m \lambda_i \left( y_i f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}}) - 1 \right) &\leq \frac{\|\boldsymbol{\theta}\|^2}{q_{\min}^{\frac{2}{L}} L e^{-q_{\min}} \log \frac{1}{\mathcal{L}}} \sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}) - q_{\min}) \\ &= \frac{\|\boldsymbol{\theta}\|^2}{q_{\min}^{\frac{2}{L}} L \log \frac{1}{\mathcal{L}}} \sum_{i=1}^m e^{-(y_i f(\mathbf{x}_i; \boldsymbol{\theta}) - q_{\min})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}) - q_{\min}) \\ &\leq \frac{1}{\tilde{\gamma}(t_0)^{\frac{2}{L}} L \log \frac{1}{\mathcal{L}}} \sum_{i=1}^m e^{-(y_i f(\mathbf{x}_i; \boldsymbol{\theta}) - q_{\min})} (y_i f(\mathbf{x}_i; \boldsymbol{\theta}) - q_{\min}) \quad (\text{Lemmata 11, 9}) \\ &\leq \frac{m}{e \tilde{\gamma}(t_0)^{\frac{2}{L}} L \log \frac{1}{\mathcal{L}}}, \end{aligned} \quad (78)$$

since the function  $u \mapsto e^{-u}u, u > 0$  has a maximum value of  $e^{-1}$ . ■

Building on the previous Proposition, we can provide an explicit characterization of proximity to stationarity at finite time for some algorithmic norms. The proof relies on a fundamental relationship between smoothness of a function and strong convexity of its convex conjugate.

**Proposition 29** (*Conjugate Correspondence Theorem - Thm. 5.26 in Beck (2017)*) Let  $\sigma > 0$ . If  $\psi$  is a  $\frac{1}{\sigma}$ -smooth convex function, then its conjugate  $\psi^*$  is  $\sigma$ -strongly convex.

Indeed, we show the following corollary for a special class of steepest flows.

**Corollary 30** *For steepest flow (Equation 3) with respect to a norm  $\|\cdot\|$ , whose square is a smooth function, on the exponential loss, under assumptions A1, A2, A3, parameter iterates  $\boldsymbol{\theta}_t$  are along the direction of an approximate KKT point of optimization problem MM for any  $t \geq t_0$ .*

**Proof** From Proposition 29, if  $\frac{1}{2}\|\cdot\|^2$  is  $\frac{1}{\sigma}$ -smooth w.r.t.  $\|\cdot\|$ , then the function  $\frac{1}{2}\|\cdot\|_\star^2$  is  $\sigma$ -strongly convex w.r.t.  $\|\cdot\|_\star$ . Thus, the function  $D_{\frac{1}{2}\|\cdot\|_\star^2}^\theta$  is defined with respect to a strongly convex function and it becomes a proper Bregman divergence. Hence, from Theorem 5.24 in Beck (2017), for  $\tilde{\mathbf{h}}_i^\star = q_{\min}^{\frac{1}{L}-1} \mathbf{h}_i^\star$ , where  $\tilde{\mathbf{h}}_i^\star \in \partial f(\mathbf{x}_i; \tilde{\boldsymbol{\theta}})$ ,  $i \in [m]$  such that  $\tilde{\mathbf{g}}_t^\star = -\sum_{i=1}^m e^{-y_i f(\mathbf{x}_i; \boldsymbol{\theta})} y_i \mathbf{h}_i^\star$  and  $\tilde{\mathbf{k}} \in \partial \frac{1}{2}\|\tilde{\boldsymbol{\theta}}\|^2$ , it holds:

$$D_{\frac{1}{2}\|\cdot\|_\star^2}^{\tilde{\boldsymbol{\theta}}} \left( \sum_{i=1}^m \lambda_i y_i \tilde{\mathbf{h}}_i^\star, \tilde{\mathbf{k}} \right) \geq \sigma \left\| \sum_{i=1}^m \lambda_i y_i \tilde{\mathbf{h}}_i^\star - \tilde{\mathbf{k}} \right\|_\star. \quad (79)$$

From the equivalence of the norms, this implies that the normalized iterates  $\tilde{\boldsymbol{\theta}}_t = \frac{\boldsymbol{\theta}_t}{(\min_{i \in [m]} y_i f(\mathbf{x}_i; \boldsymbol{\theta}_t))^{\frac{1}{L}}}$  satisfy all the conditions of an approximate KKT point. ■

**Remark 31** Leveraging the proof of Theorem 4, and in particular Lemma 19, one can show that there exists a subsequence of the iterates that (a) converges to a limit point, and (b) has vanishing approximation errors  $\epsilon(t)$  and  $\delta(t)$  as  $t \rightarrow \infty$ . This is because the alignment in the RHS of Equation 73 goes to 1 for such a subsequence. Therefore, the previous Corollary defines well-motivated approximate KKT points.

## Appendix C. Relationship to Adam and Shampoo

The family of steepest descent algorithms includes simplified versions (momentum turned-off) of two adaptive methods, Adam and Shampoo, which have been very popular for training deep neural networks.

### C.1 Adam

Adam (Kingma and Ba, 2015) is a popular adaptive optimization method, which is frequently used in deep learning. Following our previous notation, the update rule of Adam amounts to:

$$\begin{aligned} \mathbf{m}_t &= \beta_1 \mathbf{m}_{t-1} + (1 - \beta_1) \nabla \mathcal{L}(\boldsymbol{\theta}_{t-1}) \\ \mathbf{v}_t &= \beta_2 \mathbf{v}_{t-1} + (1 - \beta_2) \nabla \mathcal{L}(\boldsymbol{\theta}_{t-1})^2 \\ \hat{\mathbf{m}}_t &= \frac{\mathbf{m}_t}{1 - \beta_1^t}; \hat{\mathbf{v}}_t = \frac{\mathbf{v}_t}{1 - \beta_2^t} \\ \boldsymbol{\theta}_t &= \boldsymbol{\theta}_{t-1} - \eta_t \frac{\hat{\mathbf{m}}_t}{\sqrt{\hat{\mathbf{v}}_t} + \epsilon}, \end{aligned} \quad (80)$$

where the  $\sqrt{\cdot}$ ,  $^2$ ,  $\div$  operations are overloaded to operate elementwise in vectors. Parameters  $\beta_1, \beta_2$  control the memory of the update rule, while  $\epsilon$  is a numerical precision parameter. Notice that for  $\beta_1 = \beta_2 = \epsilon = 0$ , we recover sign-gradient descent, a version of Equation 1 with  $\|\cdot\| = \|\cdot\|_\infty$  where  $\Delta \boldsymbol{\theta}_t$  is normalized to satisfy  $\|\Delta \boldsymbol{\theta}_t\| = 1$ .

Wang et al. (2022) studied the implicit bias of (80) for  $\epsilon > 0$  in linear networks establishing bias towards  $\ell_2$  margin maximization, while Zhang et al. (2024) analyzed the case of  $\epsilon = 0$




Figure 4: **Geometric margins vs test accuracy in a teacher-student setup.** *Left:*  $\ell_1$  margin. *Right:*  $\ell_2$  margin. Each point corresponds to a different run (different random seed).

and generic  $\beta_1, \beta_2 \in [0, 1)$  also in linear networks and showed bias towards  $\ell_\infty$  margin maximization.

## C.2 Shampoo

Shampoo (Gupta et al., 2018) is an adaptive optimization algorithm, which has recently gained popularity in deep learning applications. For each weight matrix  $\mathbf{W}_t$  and its corresponding gradient matrix  $\mathbf{G}_t$ , the update rule amounts to:

$$\begin{aligned} \mathbf{L}_t &= \mathbf{L}_{t-1} + \mathbf{G}_t \mathbf{G}_t^T \\ \mathbf{R}_t &= \mathbf{R}_{t-1} + \mathbf{G}_t^T \mathbf{G}_t \\ \mathbf{W}_{t+1} &= \mathbf{W}_t - \eta \mathbf{L}_t^{-1/4} \mathbf{G}_t \mathbf{R}_t^{-1/4}. \end{aligned} \tag{81}$$

As Bernstein and Newhouse (2024) recently observed, with momentum turned off, this simplifies to:

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \eta \mathbf{U}_t \mathbf{V}_t^T, \tag{82}$$

where  $\mathbf{U}_t, \mathbf{V}_t$  contain the left and right singular vectors of  $\mathbf{G}_t$ , i.e.,  $\mathbf{G}_t = \mathbf{U}_t \mathbf{\Sigma}_t \mathbf{V}_t^T$ . Bernstein and Newhouse (2024) further remarked that this update corresponds to steepest descent (in matrix space) with respect to the *spectral norm*  $\sigma_{\max}(\cdot)$ . This is equivalent to an architecture-dependent norm in parameter space. For instance, if  $\boldsymbol{\theta} = (\mathbf{W}_1, \dots, \mathbf{W}_L)$ , then Shampoo without momentum corresponds to steepest descent with respect to the norm  $\|\boldsymbol{\theta}\|_S := \max(\sigma_{\max}(\mathbf{W}_1), \dots, \sigma_{\max}(\mathbf{W}_L))$ .

## Appendix D. Experimental Details

All experiments are implemented in PyTorch (Paszke et al., 2017).

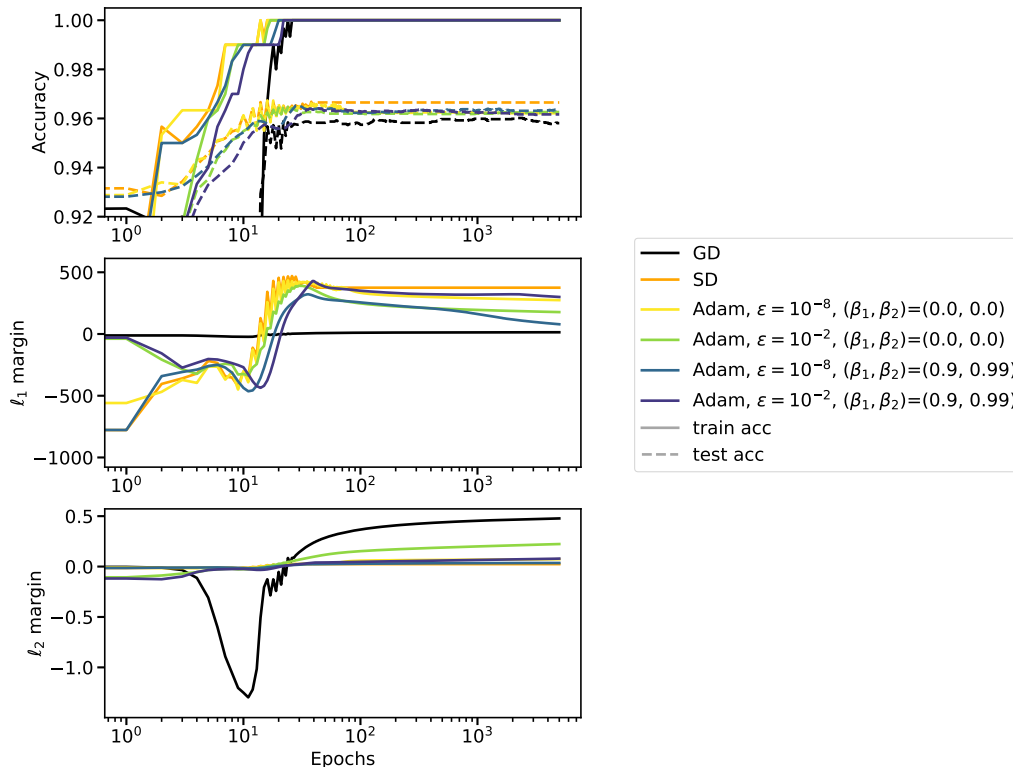


Figure 5: **Relationship between Adam and steepest descent algorithms.** Digits '2' and '7'.

**Teacher-student experiments** We use the following hyperparameters:  $d = 2^{32}$ ,  $k = 64$ ,  $k' = 1024$ ,  $m = 250$ , learning rate  $\eta = 6 \times 10^{-3}$  and density  $\frac{\|\theta^*\|_0}{k'(d+1)} = 0.0001$  (3 coordinates active per neuron). We vary the scale of initialization in  $\{0.1, 0.01, 0.001\}$  and we train for  $10^5$  epochs. Each random seed affects the draw of the datasets and the initialization of the parameters of the network. Test accuracy is estimated using 20,000 unseen data drawn from the same generative process.

**Adam experiments** We use a constant learning rate of  $3 \times 10^{-3}$  and 1-hidden layer neural networks of width 128, optimizing the logistic loss. The digits that we extract are '3' and '6' (100 training points). Each random seed corresponds to a different draw of the training dataset and different initialization. Sign gradient descent runs were very effective in minimizing the training loss, and we stopped the training early after the loss reached value smaller than  $10^{-7}$  in order to avoid numerical issues. We depict the final value, repeated for as many epochs as shown in the figures (as if the model has indeed converged).

Figure 5 shows accuracy and margins for a different pair of digits ("2" vs "7").

**Shampoo experiments** We trained 1-hidden layer neural networks of width 128 with the second layer frozen at initialization. The scale of initialization was set to  $10^{-2}$ . We used a learning rate of  $10^{-2}$  for Adam and 10, 20 for GD and Shampoo, respectively (the unusually

large learning rate is due to a frozen second layer – the models were still away from the EOS regime).

## References

- A. Nemirovskii and D. Yudin. *Problem complexity and method efficiency in optimization*. Wiley, 1983.
- Arindam Banerjee, Srujana Merugu, Inderjit S. Dhillon, and Joydeep Ghosh. Clustering with Bregman Divergences. *J. Mach. Learn. Res.*, 6:1705–1749, 2005.
- Amir Beck. *First-Order Methods in Optimization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017.
- Jeremy Bernstein and Laker Newhouse. Old Optimizer, New Norm: An Anthology. abs/2409.20325, 2024.
- Stephen P. Boyd and Lieven Vandenberghhe. *Convex Optimization*. Cambridge University Press, 2014. ISBN 978-0-521-83378-3.
- L.M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967. ISSN 0041-5553.
- Yuhang Cai, Jingfeng Wu, Song Mei, Michael Lindsey, and Peter L. Bartlett. Large stepsize gradient descent for non-homogeneous two-layer networks: Margin improvement and fast optimization. abs/2406.08654, 2024.
- David E. Carlson, Edo Collins, Ya-Ping Hsieh, Lawrence Carin, and Volkan Cevher. Pre-conditioned Spectral Descent for Deep Learning. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2971–2979, 2015.
- Frank H. Clarke. Generalized gradients and applications. *Transactions of the American Mathematical Society*, 205:247–262, 1975.
- Frank H. Clarke. *Optimization and Nonsmooth Analysis*. Society for Industrial and Applied Mathematics, 1990.
- Damek Davis, Dmitriy Drusvyatskiy, Sham M. Kakade, and Jason D. Lee. Stochastic Subgradient Method Converges on Tame Functions. *Found. Comput. Math.*, 20(1):119–154, 2020.
- Joydeep Dutta, Kalyanmoy Deb, Rupesh Tulshyan, and Ramnik Arora. Approximate KKT points and a proximity measure for termination. *J. Glob. Optim.*, 56(4), 2013.
- W. Fenchel. On Conjugate Convex Functions. *Canadian Journal of Mathematics*, 1(1): 73–77, 1949.

- Spencer Frei, Niladri S. Chatterji, and Peter L. Bartlett. Benign Overfitting without Linearity: Neural Network Classifiers Trained by Gradient Descent for Noisy Linear Data. In *Conference on Learning Theory, 2-5 July 2022, London, UK*, volume 178 of *Proceedings of Machine Learning Research*, pages 2668–2703. PMLR, 2022.
- Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nathan Srebro. Characterizing Implicit Bias in Terms of Optimization Geometry. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 1827–1836. PMLR, 2018.
- Vineet Gupta, Tomer Koren, and Yoram Singer. Shampoo: Preconditioned Stochastic Tensor Optimization. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 1837–1845. PMLR, 2018.
- Niv Haim, Gal Vardi, Gilad Yehudai, Ohad Shamir, and Michal Irani. Reconstructing Training Data From Trained Neural Networks. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022.
- Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- Ziwei Ji and Matus Telgarsky. Directional convergence and alignment in deep learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Keller Jordan, Yuchen Jin, Vlado Boza, Jiacheng You, Franz Cesista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks, 2024. URL <https://web.archive.org/web/20250122060345/https://kellerjordan.github.io/posts/muon/>.
- William Karush. Minima of Functions of Several Variables with Inequalities as Side Conditions. Master’s thesis, Department of Mathematics, University of Chicago, 1939.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Kuhn, H. W. and Tucker, A. W. Nonlinear programming. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950*, pages 481–492. University of California Press, 1951.
- Daniel Kunin, Atsushi Yamamura, Chao Ma, and Surya Ganguli. The Asymmetric Maximum Margin Bias of Quasi-Homogeneous Neural Networks. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023.

- Frederik Kunstner, Jacques Chen, Jonathan Wilder Lavington, and Mark Schmidt. Noise Is Not the Main Factor Behind the Gap Between Sgd and Adam on Transformers, But Sign Descent Might Be. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023.
- Tim Large, Yang Liu, Minyoung Huh, Hyojin Bahng, Phillip Isola, and Jeremy Bernstein. Scalable Optimization in the Modular Norm. abs/2405.14813, 2024. URL <https://doi.org/10.48550/arXiv.2405.14813>.
- Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- Kaifeng Lyu and Jian Li. Gradient Descent Maximizes the Margin of Homogeneous Neural Networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020.
- Llew Mason, Jonathan Baxter, Peter Bartlett, and Marcus Frean. Boosting Algorithms as Gradient Descent. In *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.
- Mor Shpigel Nacson, Suriya Gunasekar, Jason D. Lee, Nathan Srebro, and Daniel Soudry. Lexicographic and Depth-Sensitive Margins in Homogeneous and Non-Homogeneous Deep Models. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 4683–4692. PMLR, 2019.
- Behnam Neyshabur, Ruslan Salakhutdinov, and Nathan Srebro. Path-SGD: Path-Normalized Optimization in Deep Neural Networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2422–2430, 2015a.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In Search of the Real Inductive Bias: On the Role of Implicit Regularization in Deep Learning. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Workshop Track Proceedings*, 2015b.
- Andrew Y. Ng. Feature selection, L1 vs. L2 regularization, and rotational invariance. In *Proceedings of the Twenty-First International Conference on Machine Learning, ICML '04*, page 78, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581138385.
- Mykola Novik. torch-optimizer – collection of optimization algorithms for PyTorch., 1 2020.
- Albert B. J. Novikoff. On Convergence Proofs For Perceptrons. 1963.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS 2017 Workshop on Autodiff*, 2017.

- R. Tyrrell Rockafellar. *Convex Analysis*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, 1970.
- Ohad Shamir. The Implicit Bias of Benign Overfitting. *J. Mach. Learn. Res.*, 24:113:1–113:40, 2023. URL <https://jmlr.org/papers/v24/22-0784.html>.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The Implicit Bias of Gradient Descent on Separable Data. *J. Mach. Learn. Res.*, 19:70:1–70:57, 2018.
- Matus Telgarsky. Margins, Shrinkage, and Boosting. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 307–315, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- Nikolaos Tsilivis, Natalie Frank, Nathan Srebro, and Julia Kempe. The Price of Implicit Bias in Adversarially Robust Generalization. abs/2406.04981, 2024. URL <https://arxiv.org/abs/2406.04981>.
- Nikolaos Tsilivis, Gal Vardi, and Julia Kempe. Flavors of Margin: Implicit Bias of Steepest Descent in Homogeneous Neural Networks. In *13th International Conference on Learning Representations, ICLR 2025, Singapore, Singapore, April 24-28, 2025*, 2025.
- Vladimir Vapnik. *Statistical learning theory*. Wiley, 1998. ISBN 978-0-471-03003-4.
- Gal Vardi. On the Implicit Bias in Deep-Learning Algorithms. *Commun. ACM*, 66(6):86–93, 2023.
- Gal Vardi, Ohad Shamir, and Nati Srebro. On Margin Maximization in Linear and ReLU Networks. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022.
- Bohan Wang, Qi Meng, Wei Chen, and Tie-Yan Liu. The Implicit Bias for Adaptive Optimization Algorithms on Homogeneous Neural Networks. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 10849–10858. PMLR, 2021.
- Bohan Wang, Qi Meng, Huishuai Zhang, Ruoyu Sun, Wei Chen, Zhi-Ming Ma, and Tie-Yan Liu. Does Momentum Change the Implicit Regularization on Separable Data? In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022.
- Ashia C. Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The Marginal Value of Adaptive Gradient Methods in Machine Learning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4148–4158, 2017.

- Blake E. Woodworth, Suriya Gunasekar, Jason D. Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and Rich Regimes in Overparametrized Models. In *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pages 3635–3673. PMLR, 2020.
- Shuo Xie and Zhiyuan Li. Implicit Bias of AdamW:  $\ell_\infty$ -Norm Constrained Optimization. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, 2024.
- Chenyang Zhang, Difan Zou, and Yuan Cao. The Implicit Bias of Adam on Separable Data. [abs/2406.10650](https://arxiv.org/abs/2406.10650), 2024.
- Jingzhao Zhang, Sai Praneeth Karimireddy, Andreas Veit, Seungyeon Kim, Sashank J. Reddi, Sanjiv Kumar, and Suvrit Sra. Why are Adaptive Methods Good for Attention Models? In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.