
Lexicographic and Depth-Sensitive Margins in Homogeneous and Non-Homogeneous Deep Models

Mor Shpigel Nacson¹ Suriya Gunasekar² Jason D. Lee³ Nathan Srebro² Daniel Soudry¹

Abstract

With an eye toward understanding complexity control in deep learning, we study how infinitesimal regularization or gradient descent optimization lead to margin maximizing solutions in both homogeneous and *non homogeneous* models, extending previous work that focused on infinitesimal regularization only in homogeneous models. To this end we study the limit of loss minimization with a diverging norm constraint (the “constrained path”), relate it to the limit of a “margin path” and characterize the resulting solution. For non-homogeneous ensemble models, which output is a sum of homogeneous sub-models, we show that this solution discards the shallowest sub-models if they are unnecessary. For homogeneous models, we show convergence to a “lexicographic max-margin solution”, and provide conditions under which max-margin solutions are also attained as the limit of unconstrained gradient descent.

1. Introduction

Inductive bias introduced through the learning process plays a crucial role in training deep neural networks and in the generalization properties of the learned models (Neyshabur et al., 2015b;a; Zhang et al., 2017; Keskar et al., 2017; Neyshabur et al., 2017; Wilson et al., 2017; Hoffer et al., 2017). Deep neural networks used in practice are typically highly overparameterized, i.e., have far more trainable parameters than training examples. Thus, using these models, it is usually possible to fit the data perfectly and obtain zero training error (Zhang et al., 2017). However, simply minimizing the training loss does not guarantee good generalization to unseen data – many global min-

ima of the training loss indeed have very high test error (Wu et al., 2017). The inductive bias introduced in our learning process affects which specific global minimizer is chosen as the predictor. Therefore, it is essential to understand the nature of this inductive bias to understand why overparameterized models, and particularly deep neural networks, exhibit good generalization abilities.

A common way to introduce an additional inductive bias in overparameterized models is via small amounts of regularization, or loose constraints. For example, Rosset et al. (2004b;a); Wei et al. (2018) show that, in overparameterized classification models, a vanishing amount of regularization, or a diverging norm constraint can lead to max-margin solutions, which in turn enjoy strong generalization guarantees.

A second and more subtle source of inductive bias is via the optimization algorithm used to minimize the underdetermined training objective (Gunasekar et al., 2017; Soudry et al., 2018b). Common algorithms used in neural network training, such as stochastic gradient descent, iteratively refine the model parameters by making incremental local updates. For different algorithms, the local updates are specified by different geometries in the space of parameters. For example, gradient descent uses an Euclidean ℓ_2 geometry, while coordinate descent updates are specified in the ℓ_1 geometry. The minimizers to which such local search based optimization algorithms converge to are indeed very special and are related to the geometry of the optimization algorithm (Gunasekar et al., 2018b) as well as the choice of model parameterization (Gunasekar et al., 2018a).

In this work we similarly investigate the connection between margin maximization and the limits of

- The “optimization path” of unconstrained, unregularized gradient descent.
- The “constrained path”, where we optimize with a diverging (increasingly loose) constraint on the norm of the parameters.
- The closely related “regularization path”, of solutions with decreasing penalties on the norm.

¹Technion, Israel ²TTI Chicago, USA ³USC Los Angeles, USA. Correspondence to: Mor Shpigel Nacson <morshpigel@google.com>.

To better understand the questions we tackle in this paper, and our contribution toward understanding the inductive bias introduced in training, let us briefly survey prior work.

Equivalence of the regularization or constrained paths and margin maximization: Rosset et al. (2004b;a); Wei et al. (2018) investigated the connection between the regularization and constrained paths and the max-margin solution. Rosset et al. (2004a;b) considered linear (hence homogeneous) models with monotone loss and explicit norm regularization or constraint, and proved convergence to the max-margin solution for certain loss functions (e.g., logistic loss) as the regularization vanishes or the norm constraint diverges. Wei et al. (2018) extended the regularization path result to non-linear but positive-homogeneous prediction functions,

Definition 1 (α -positive homogeneous function). *A function $g(\boldsymbol{\theta}) : \mathbb{R}^d \rightarrow \mathbb{R}$ is α -positive homogeneous if $\forall \rho > 0$ and $\forall \boldsymbol{\theta} \in \mathbb{R}^d : g(\rho\boldsymbol{\theta}) = \rho^\alpha g(\boldsymbol{\theta})$.*

e.g. as obtained by a ReLU network with uniform depth.

These results are thus limited to only positive homogeneous predictors, and do not include deep networks with bias parameters, ensemble models with different depths, ResNets, or other models with skip connections. Here, we extend this connection beyond positive homogeneous predictors.

Furthermore, even for homogeneous or linear predictors, there might be multiple margin maximizing solutions. For linear models, Rosset et al. (2004b) alluded to a refined set of maximum margin classifiers that in addition to maximizing the distance to the closest data point (max-margin), also maximize the distance to the second closest data point, and so on. We formulate such special maximum margin solutions as “lexicographic max-margin” classifiers which we introduce in Section 4.2. We show that for general continuous homogeneous models, the constrained path with diverging norm constraint converges to these more refined “lexicographic max-margin” classifiers.

Equivalence of the optimization path and margin maximization: Another line of works studied the connection between unconstrained, unregularized optimization with a specific algorithm (i.e., the limit of the “optimization path”), and the max-margin solution. For linear prediction with the logistic loss (or other exponential tail losses), we now know gradient descent (Soudry et al., 2018b; Ji & Telgarsky, 2018) as well as SGD (Nacson et al., 2019b) converges in direction to the max-margin solution, while steepest descent with respect to an arbitrary norm converges to the max-margin w.r.t. the corresponding norm (Gunasekar et al., 2018b). All the above results are for linear prediction. Gunasekar et al. (2018a); Nacson et al. (2019a); Ji & Telgarsky (2019) obtained results establishing convergence

to margin maximizing solutions also for certain uniform-depth linear networks (including fully connected networks and convolutional networks), which still implement linear model. Separately, Xu et al. (2019) analyzed a single linear unit with ReLU activation—a limited non-linear but still positive homogeneous model. Lastly, Soudry et al. (2018a) analyzed a non-linear ReLU network where only a single weight layer is optimized.

Here, we extend this relationship to general, non-linear and positive homogeneous predictors for which the loss can be minimized only at infinity. We establish a connection between the limit of unregularized unconstrained optimization and the max-margin solution.

Problems with finite minimizers: We note that the connection between regularization path and optimization path was previously considered in a different settings, where a finite (global) minimum exists. In such settings the questions asked are different than the ones we consider here, and are not about the limit of the paths. E.g., Ali et al. (2018) showed for gradient flow a multiplicative relation between the risk for the gradient flow optimization path and the ridge-regression regularization path. Also, Suggala et al. (2018) showed that for gradient flow and strongly convex and smooth loss function – gradient descent iterates on the unregularized loss function are pointwise close to solutions of a corresponding regularized problem.

Contributions

We examine overparameterized realizable problems (i.e., where it is possible to perfectly classify the training data), when training using monotone decreasing classification loss functions. For simplicity, we focus on the exponential loss. However, using similar techniques as in Soudry et al. (2018a) our results should extend to other exponential-tailed loss functions such as the logistic loss and its multi-class generalization. This is indeed the common setting for deep neural networks used in practice.

We show that in any model,

- As long as the margin attainable by a (unregularized, unconstrained) model is unbounded, then the margin of the constrained path converges to the max-margin. See Corollary 1.
- If additional conditions hold, the constrained path also converges to the “margin path” in parameter space (the path of minimal norm solutions attaining increasingly large margins). See section 3.1.

We then demonstrate that

- If the model is a sum of homogeneous functions of different orders (i.e., it is not homogeneous itself), then we can still characterize the asymptotic solution of both the constrained path and the margin path. See Theorem 3.2.
- This solution implies that in an ensemble of homogeneous neural networks, the ensemble will aim to discard the most *shallow* network. This is in contrast to what we would expect from considerations of optimization difficulty (since deeper networks are typically harder to train (He et al., 2016)).
- This also allows us to represent hard-margin SVM problems **with unregularized bias** using such models. This is in contrast to previous approaches which fail to do so, as pointed out recently (Nar et al., 2019).

Finally, for homogeneous models,

- We find general conditions under which the optimization path converges to stationary points of the margin path or the constrained path. See section 4.1.
- We show that the constrained path converges to a specific type max-margin solution, which we term the “lexicographic max-margin”.¹ See Theorem 4.

2. Preliminaries and Basic Results

In this paper, we will study the following exponential tailed loss function

$$\mathcal{L}(\boldsymbol{\theta}) \triangleq \sum_{n=1}^N \exp(-f_n(\boldsymbol{\theta})), \quad (1)$$

where $f_n : \mathbb{R}^d \rightarrow \mathbb{R}$ is a continuous function, and N is the number of samples. Also, for any norm $\|\cdot\|$ in \mathbb{R}^d we define \mathbb{S}^{d-1} as the unit norm ball in \mathbb{R}^d .

We will use in our results the following basic lemma

Lemma 1. *Let f and g be two functions from \mathbb{R}^d to \mathbb{R} , such that*

$$\phi(\rho) = \min_{\mathbf{w} \in \mathbb{R}^d} f(\mathbf{w}) \text{ s.t. } g(\mathbf{w}) \leq \rho \quad (2)$$

exists and is strictly monotonically decreasing in ρ , $\forall \rho \geq \rho_0$, for some ρ_0 . Then, $\forall \rho \geq \rho_0$, the optimization problem in eq. 2 has the same set of solutions (\mathbf{w}) as

$$\min_{\mathbf{w} \in \mathbb{R}^d} g(\mathbf{w}) \text{ s.t. } f(\mathbf{w}) \leq \phi(\rho), \quad (3)$$

whose minimum is obtained at $g(\mathbf{w}) = \rho$.

Proof. See Appendix A. \square

¹The authors thank Rob Shapire for the suggestion of the nomenclature during initial discussions.

2.1. The Optimization Path

The optimization path in the Euclidean norm $\boldsymbol{\theta}(t)$, is given by the direction of iterates of gradient descent algorithm with initialization $\boldsymbol{\theta}(0)$ and learning rates $\{\eta_t\}_{t=1}^{\infty}$,

$$\begin{aligned} \text{Optimization path: } \bar{\boldsymbol{\theta}}(t) &= \frac{\boldsymbol{\theta}(t)}{\|\boldsymbol{\theta}(t)\|}, \text{ where} \\ \boldsymbol{\theta}(t) &= \boldsymbol{\theta}(t-1) - \eta_t \nabla \mathcal{L}(\boldsymbol{\theta}(t-1)). \end{aligned} \quad (4)$$

2.2. The Constrained Path

The constrained path for the loss in eq. 1 is given by minimizer of the loss at a given norm value $\rho > 0$, i.e.,

$$\text{Constrained path: } \Theta_c(\rho) \triangleq \arg \min_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \mathcal{L}(\rho\boldsymbol{\theta}). \quad (5)$$

The constrained path was previously considered for linear models (Rosset et al., 2004a). However, most previous works (e.g. Rosset et al. (2004b); Wei et al. (2018)) focused on the regularization path, which is the minimizer of the regularized loss. These two paths are closely linked, as we discuss in more detail in Appendix F.

Denote the constrained minimum of the loss as follows:

$$\mathcal{L}^*(\rho) \triangleq \min_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \mathcal{L}(\rho\boldsymbol{\theta}).$$

$\mathcal{L}^*(\rho)$ exists for any finite ρ as the minimum of a continuous function on a compact set.

By Lemma 1, the Assumption

Assumption 1. *There exists ρ_0 such that $\mathcal{L}^*(\rho)$ is strictly monotonically decreasing to zero for any $\rho \geq \rho_0$.*

enables an alternative form of the constrained path

$$\Theta_c(\rho) = \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \|\boldsymbol{\theta}\|^2 \text{ s.t. } \mathcal{L}(\rho\boldsymbol{\theta}) \leq \mathcal{L}^*(\rho).$$

In addition, in the next lemma we show that under this assumption the constrained path minimizers are obtained on the boundary of \mathbb{S}^{d-1} .

Lemma 2. *Under assumption 1, for all $\rho > \rho_0$ and for all $\boldsymbol{\theta}_c \in \Theta_c(\rho)$, we have $\|\boldsymbol{\theta}_c\| = 1$.*

Proof. Let $\rho > 0$. We assume, in contradiction, that $\exists \boldsymbol{\theta}_c \in \Theta_c(\rho)$ so that $\|\boldsymbol{\theta}_c\| = b < 1$. This implies that $\mathcal{L}^*(\rho) = \mathcal{L}^*(\rho b)$ which contradicts our assumption that $\mathcal{L}^*(\rho)$ is strictly monotonically decreasing. \square

2.3. The Margin Path

For prediction functions $f_n : \mathbb{R}^d \rightarrow \mathbb{R}$ on data points indexed as $n = 1, 2, \dots, N$, we define the margin path as:

$$\text{Margin path: } \Theta_m(\rho) \triangleq \arg \max_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \min_n f_n(\rho\boldsymbol{\theta}). \quad (6)$$

For $\boldsymbol{\theta} \in \mathbb{S}^{d-1}$, we denote the margin at scaling $\rho > 0$ as

$$\gamma(\rho, \boldsymbol{\theta}) = \min_n f_n(\rho \boldsymbol{\theta}),$$

and the max-margin at scale of $\rho > 0$ as

$$\gamma^*(\rho) = \max_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \min_n f_n(\rho \boldsymbol{\theta}).$$

Note that for all ρ , this maximum exists as the maximum of a continuous function on a compact set.

Again, we make a simplifying assumption

Assumption 2. *There exist ρ_0 such that $\gamma^*(\rho)$ is strictly monotonically increasing to ∞ for any $\rho \geq \rho_0$.*

Many common prediction functions satisfy this assumption, including the sum of positive-homogeneous prediction functions.

Using Lemma 1 with Assumption 2, we have:

$$\begin{aligned} \Theta_m(\rho) &= \arg \max_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \min_n f_n(\rho \boldsymbol{\theta}) \\ &= \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \|\boldsymbol{\theta}\|^2 \text{ s.t. } \min_n f_n(\rho \boldsymbol{\theta}) \geq \gamma^*(\rho). \end{aligned} \quad (7)$$

3. Non-Homogeneous Models

We first study the constrained path in non-homogeneous models, and relate it to the margin path. To do so, we need to first define the ϵ -ball surrounding a set $\mathcal{A} \subset \mathbb{R}^d$

$$\mathcal{B}_\epsilon(\mathcal{A}) \triangleq \{\boldsymbol{\theta} \in \mathbb{R}^d \mid \exists \boldsymbol{\theta}' \in \mathcal{A} : \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| < \epsilon\},$$

and the notion of set convergence

Definition 2 (Set convergence). *A sequence of sets $\mathcal{A}_t \subset \mathbb{R}^d$ converges to another sequence of sets $\mathcal{A}'_t \subset \mathbb{R}^d$ if $\forall \epsilon > 0 \exists t_0$ such that $\forall t > t_0 \mathcal{A}_t \subset \mathcal{B}_\epsilon(\mathcal{A}'_t)$.*

3.1. Margin of Constrained Path Converges to Maximum Margin

For all ρ , the constrained path margin deviation from the max-margin is bounded, as we prove next.

Lemma 3. *For all ρ , and every $\boldsymbol{\theta}_c(\rho)$ in $\Theta_c(\rho)$*

$$\gamma^*(\rho) - \gamma(\rho, \boldsymbol{\theta}_c(\rho)) \leq \log N. \quad (8)$$

Proof. Note that $\forall \boldsymbol{\theta}$:

$$e^{-\gamma(\rho, \boldsymbol{\theta})} \leq \sum_{n=1}^N \exp(-f_n(\rho \boldsymbol{\theta})) \leq N e^{-\gamma(\rho, \boldsymbol{\theta})}. \quad (9)$$

Since, $\forall \boldsymbol{\theta} \in \mathbb{S}^{d-1}, \mathcal{L}(\rho \boldsymbol{\theta}_c(\rho)) \leq \mathcal{L}(\rho \boldsymbol{\theta})$, we have, $\forall \boldsymbol{\theta}_c(\rho) \in \Theta_c(\rho)$ and $\forall \boldsymbol{\theta}_m(\rho) \in \Theta_m(\rho)$,

$$\begin{aligned} 1 &\leq \frac{\mathcal{L}(\rho \boldsymbol{\theta}_m(\rho))}{\mathcal{L}(\rho \boldsymbol{\theta}_c(\rho))} = \frac{\sum_{n=1}^N \exp(-f_n(\rho \boldsymbol{\theta}_m(\rho)))}{\sum_{n=1}^N \exp(-f_n(\rho \boldsymbol{\theta}_c(\rho)))} \\ &\leq N \exp(-(\gamma^*(\rho) - \gamma(\rho, \boldsymbol{\theta}_c(\rho)))) . \end{aligned}$$

$$\Rightarrow \gamma^*(\rho) - \gamma(\rho, \boldsymbol{\theta}_c(\rho)) \leq \log N. \quad \square$$

Lemma 3 immediately implies that

Corollary 1. *If $\lim_{\rho \rightarrow \infty} \gamma^*(\rho) = \infty$, then for all ρ , and every $\boldsymbol{\theta}_c(\rho)$ in $\Theta_c(\rho)$*

$$\lim_{\rho \rightarrow \infty} \frac{\gamma^*(\rho)}{\gamma(\rho, \boldsymbol{\theta}_c(\rho))} = 1.$$

The last corollary states that the margin of the constrained path converges to the maximum margin. However, this does not necessarily imply convergence in parameter space, i.e., this result does not guaranty that $\Theta_c(\rho)$ converges to $\Theta_m(\rho)$. We analyze some positive and negative examples to demonstrate this claim.

Example 1: homogeneous models

It is straightforward to see that, for α -positive homogeneous prediction functions (Definition 1) the margin path $\Theta_m(\rho)$ in eq. 6 is the same set for any ρ , and is given by

$$\Theta_m^* = \arg \max_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \min_n f_n(\boldsymbol{\theta}).$$

Additionally, as we show next, for such models Lemma 3 implies convergence in parameter space, i.e., $\Theta_c(\rho)$ converges to $\Theta_m(\rho)$. To see this, notice that for α -positive homogeneous functions $f_n, \forall \boldsymbol{\theta}_c(\rho) \in \Theta_c(\rho)$:

$$\begin{aligned} \gamma^*(\rho) - \gamma(\rho, \boldsymbol{\theta}_c(\rho)) &= \\ &= \max_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \min_n f_n(\rho \boldsymbol{\theta}) - \min_n f_n(\rho \boldsymbol{\theta}_c(\rho)) \\ &= \rho^\alpha \left(\max_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \min_n f_n(\boldsymbol{\theta}) - \min_n f_n(\boldsymbol{\theta}_c(\rho)) \right) \\ &\leq \log N. \end{aligned}$$

For $\rho \rightarrow \infty$ we must have

$$\left(\max_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \min_n f_n(\boldsymbol{\theta}) - \min_n f_n(\boldsymbol{\theta}_c(\rho)) \right) \rightarrow 0.$$

By continuity, the last equation implies that $\Theta_c(\rho)$ converges to $\Theta_m(\rho)$. For full details see Appendix D.1.

Connection to previous results: For linear models, Rosset et al. (2004a) connected the L_1 constrained path and maximum L_1 margin solution. In addition, for any norm, Rosset et al. (2004b) showed that the regularization path converges to the limit of the margin path. In a recent work, Wei et al. (2018) extended this result to homogeneous models with cross-entropy loss. Here, for homogeneous models and any norm, we show a connection between the constrained path and the margin path.

Extension: Later, in Theorem 4 we prove a more refined result: the constrained path converges to a *specific subset* of the margin path set (the lexicographic max-margin set).

In contrast, in general models, 8 does not necessarily imply convergence in the parameter space. We demonstrate this result in the next example.

Example 2: log predictor: We denote $\mathbf{z}_n = y_n \mathbf{x}_n$ for some dataset $\{\mathbf{x}_n, y_n\}_{n=1}^N$, with features \mathbf{x}_n and label y_n . We examine the prediction function $f_n(\rho, \boldsymbol{\theta}) = \log(\rho \boldsymbol{\theta}^\top \mathbf{z}_n)$ for $\boldsymbol{\theta}^\top \mathbf{z}_n > 0$. We focus on the loss function tail behaviour and thus only care about the loss function behaviour in $\boldsymbol{\theta}^\top \mathbf{z}_n > 0$ region. We assume that a separator which satisfy this constraint exists since we are focusing on realizable problems.

Since $\log(\cdot)$ is strictly increasing and $\rho > 0$, we have

$$\gamma(\rho, \boldsymbol{\theta}) = \min_n f_n(\rho \boldsymbol{\theta}) = \log\left(\rho \min_n \boldsymbol{\theta}^\top \mathbf{z}_n\right).$$

We denote $\tilde{\gamma}(\boldsymbol{\theta}) = \min_n \boldsymbol{\theta}^\top \mathbf{z}_n$ and $\tilde{\gamma}^* = \max_{\boldsymbol{\theta} \in \mathbb{S}^{d-1}} \tilde{\gamma}(\boldsymbol{\theta})$. Note that $\gamma^*(\rho) = \log(\rho \tilde{\gamma}^*)$. Now consider $\boldsymbol{\theta}_c(\rho) \in \Theta_c(\rho)$ such that for some ρ_0 and $\forall \rho > \rho_0$: $\tilde{\gamma}(\boldsymbol{\theta}_c(\rho)) = \min_n \boldsymbol{\theta}_c(\rho)^\top \mathbf{z}_n \geq \frac{\tilde{\gamma}^*}{N}$. For this case, we still have,

$$\begin{aligned} \gamma^*(\rho) - \gamma(\rho, \boldsymbol{\theta}_c(\rho)) &= \log(\rho \tilde{\gamma}^*) - \log(\rho \tilde{\gamma}(\boldsymbol{\theta}_c(\rho))) \\ &= \log\left(\frac{\tilde{\gamma}^*}{\tilde{\gamma}(\boldsymbol{\theta}_c(\rho))}\right) \leq \log N. \end{aligned}$$

but clearly, $\tilde{\gamma}(\boldsymbol{\theta}_c(\rho)) \not\rightarrow \tilde{\gamma}^*$. Thus, Lemma 3 does not guarantee that $\tilde{\gamma}(\boldsymbol{\theta}_c(\rho)) \rightarrow \tilde{\gamma}^*$ as $\rho \rightarrow \infty$, or that $\Theta_c(\rho)$ converges to $\Theta_m(\rho)$.

Analogies with regularization and optimization paths: This example demonstrates that for the prediction function $\log(\rho \boldsymbol{\theta}^\top \mathbf{z})$ for $\boldsymbol{\theta}^\top \mathbf{z} > 0$, the constrained path does not necessarily converge to the margin path. This is equivalent to *setup A*: linear prediction models with loss function $\exp(-\log(u))$. Rosset et al. (2004b) and Nacson et al. (2019a) state related results for *setup A*. Both works derived conditions on the loss function that ensure convergence to the margin path from the regularization/ optimization path respectively. Rosset et al. (2004b) showed that in *setup A* the regularization path does not necessarily converge to the margin path. (Nacson et al., 2019a) showed a similar result for the optimization path, i.e., that in *setup A* the optimization path does not necessarily converge to the margin path. Both results align with our results for the constrained path.

In contrast, according to the conditions of Rosset et al. (2004b); Nacson et al. (2019a), we know that if the prediction function is $\log^{1+\epsilon}(\rho \boldsymbol{\theta}^\top \mathbf{z})$ for some $\epsilon > 0$ and $\boldsymbol{\theta}^\top \mathbf{z} > 0$, then the regularization path and optimization path *do converge* to the margin path. In the next example, we show that this is also true for the constrained path.

Example 3: $(1+\epsilon)$ -log predictor: We examine the prediction function $f_n(\rho, \boldsymbol{\theta}) = \log^{1+\epsilon}(\rho \boldsymbol{\theta}^\top \mathbf{z}_n)$ for $\boldsymbol{\theta}^\top \mathbf{z}_n > 0$

and some $\epsilon > 0$. Since the log function is strictly increasing and $\epsilon, \rho > 0$, we have

$$\gamma(\rho, \boldsymbol{\theta}) = \min_n f_n(\rho \boldsymbol{\theta}) = \log^{1+\epsilon}\left(\rho \min_n \boldsymbol{\theta}^\top \mathbf{z}_n\right).$$

For all $\boldsymbol{\theta}_c(\rho) \in \Theta_c(\rho)$:

$$\begin{aligned} \gamma^*(\rho) - \gamma(\rho, \boldsymbol{\theta}_c(\rho)) &= (1 + \epsilon) \log^\epsilon(\rho) (\log(\tilde{\gamma}^*) - \log(\tilde{\gamma}(\boldsymbol{\theta}_c(\rho)))) \\ &\quad + o(\log^\epsilon(\rho)) \leq N. \end{aligned}$$

For $\rho \rightarrow \infty$ we must have $(\log(\tilde{\gamma}^*) - \log(\tilde{\gamma}(\boldsymbol{\theta}_c(\rho)))) \rightarrow 0$, which implies, by continuity, that $\Theta_c(\rho)$ converges to $\Theta_m(\rho)$. For details, see Appendix D.2.

3.2. Sum of Positively Homogeneous Functions

Remark: The results in this subsection are specific for the Euclidean or L_2 norm.

Let $f_n(\rho \boldsymbol{\theta})$ be functions that are a finite sum of positively homogeneous functions, i.e., for some finite K :

$$\forall n : f_n(\rho \boldsymbol{\theta}) = \sum_{k=1}^K f_n^{(k)}(\rho \boldsymbol{\theta}_k), \quad (10)$$

where $\boldsymbol{\theta} = [\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K]$ and $f_n^{(k)}(\boldsymbol{\theta}_k)$ are α_k -positive homogeneous functions, where $0 < \alpha_1 < \alpha_2 < \dots < \alpha_K$.

First, we characterize the asymptotic form of the margin path in this setting.

Lemma 4. *Let $f_n(\boldsymbol{\theta})$ be a sum of positively homogeneous functions as in eq. 10. Then, the set of solutions of*

$$\arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \|\boldsymbol{\theta}\|^2 \text{ s.t. } \forall n : f_n(\rho \boldsymbol{\theta}) \geq \gamma^*(\rho). \quad (11)$$

can be written as

$$\boldsymbol{\theta}_k^* = \frac{1}{\rho} (\mathbf{w}_k + o(1)) (\gamma^*(\rho))^{\frac{1}{\alpha_k}} \quad (12)$$

where the $o(1)$ term is vanishing as $\gamma^*(\rho) \rightarrow \infty$, and

$$\mathbf{w}^* = [\mathbf{w}_1^*, \dots, \mathbf{w}_K^*] \in \mathcal{W},$$

where

$$\mathcal{W} = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \|\mathbf{w}_1\|^2 \text{ s.t. } \forall n : f_n(\mathbf{w}) \geq 1. \quad (13)$$

Proof. We write the original optimization problem

$$\arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \sum_{k=1}^K \|\boldsymbol{\theta}_k\|^2 \text{ s.t. } \forall n : \sum_{k=1}^K f_n^{(k)}(\rho \boldsymbol{\theta}_k) \geq \gamma^*(\rho).$$

Dividing by $\gamma^*(\rho)$, using the α_k positive homogeneity of $f_n^{(k)}$, and changing the variables as $\theta_k = \frac{1}{\rho} \mathbf{w}_k (\gamma^*(\rho))^{\frac{1}{\alpha_k}}$, we obtain an equivalent optimization problem

$$\arg \min_{\mathbf{w} \in \mathbb{R}^d} \sum_{k=1}^K \gamma^*(\rho)^{\frac{2}{\alpha_k}} \|\mathbf{w}_k\|^2 \text{ s.t. } \forall n : f_n(\mathbf{w}) \geq 1. \quad (14)$$

We denote the set of solutions of eq. 14 as $\mathcal{W}(\gamma^*(\rho))$. Taking the limit of $\gamma^*(\rho) \rightarrow \infty$ of this optimization problem we find that any solution $\mathbf{w} \in \mathcal{W}(\gamma^*(\rho))$ must minimize the first term in the sum $\|\mathbf{w}_1\|^2$, and only then the other terms. Therefore the asymptotic solution is of the form of eqs. 12 and 13. We prove this reasoning formally in Appendix B, i.e., we show that

Claim 1. *The solution of eq. 14 is the same solution described in Lemma 4, i.e., eqs. 12 and 13.* \square

The following Lemma will be used to connect the constrained path to the characterization of the margin path.

Lemma 5. *Let $f_n(\rho\theta)$ be a sum of positively homogeneous functions as in eq. 10. Any path $\theta(\rho)$ such that*

$$\gamma^*(\rho) - \gamma(\rho, \theta(\rho)) < C. \quad (15)$$

is of the form described in eqs. 12 and 13.

Proof. See Appendix C. \square

Combining Lemma 3, 4 and Lemma 5 we obtain the following Theorem

Theorem 1. *Under Assumption 1 and 2, any solution in $\arg \min_{\|\theta\| \leq 1} \mathcal{L}(\rho\theta)$ converges to*

$$\theta_k^* = \frac{1}{\rho} (\mathbf{w}_k^* + o(1)) (\gamma^*(\rho))^{\frac{1}{\alpha_k}} \quad (16)$$

where the $o(1)$ term is vanishing as $\gamma^(\rho) \rightarrow \infty$, and*

$$\mathbf{w}^* = [\mathbf{w}_1^*, \dots, \mathbf{w}_K^*] \in \mathcal{W},$$

where

$$\mathcal{W} = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \|\mathbf{w}_1\|^2 \text{ s.t. } \forall n : f_n(\mathbf{w}) \geq 1. \quad (17)$$

Theorem 1 implications: An important implication of Theorem 1 is that an ensemble on neural networks will aim to discard the shallowest network in the ensemble. Consider the following setting: for each $k \in \{1, \dots, K\}$, the function $\forall n : f_n^{(k)}(\rho\theta_k)$ represents a prediction function of some feedforward neural network with no bias, all with the same positive-homogeneous activation function $\sigma(\cdot)$ of some degree α (e.g., ReLU activation is positive-homogeneous of degree 1). Note that in this setup, each

of the k prediction functions $f_n^{(k)}(\rho\theta_k)$ is also a positive-homogeneous function. In particular, network k with depth d_k is positive homogeneous with degree $\alpha_k = \alpha d_k$ where α is the activation function degree. Since all the networks have the same activation function, deeper networks will have larger degree. We assume WLOG that $d_1 < d_2 < \dots < d_K$. This implies that $\alpha_1 < \alpha_2 < \dots < \alpha_K$. In this setting, $\forall n : f_n(\rho\theta) = \sum_{k=1}^K f_n^{(k)}(\rho\theta_k)$ represents an ensemble of these networks. From Theorem 1, the solution of the constrained path will satisfy

$$\begin{aligned} f_n(\rho\theta^*) &= \sum_{k=1}^K f_n^{(k)}(\rho\theta_k^*) \\ &= \sum_{k=1}^K f_n^{(k)}\left(\left(\mathbf{w}_k^* + o(1)\right) (\gamma^*(\rho))^{\frac{1}{\alpha_k}}\right) \\ &= \gamma^*(\rho) \sum_{k=1}^K f_n^{(k)}(\mathbf{w}_k^* + o(1)), \end{aligned}$$

where $\mathbf{w}^* \in \mathcal{W}$ and \mathcal{W} is calculated using eq. 13. Examining equation 13, we observe that the network aims to minimize the \mathbf{w}_1 norm. In particular, if the network ensemble can satisfy the constraints $\forall n : f_n(\mathbf{w}) \geq 1$ with $\mathbf{w}_1 = \mathbf{0}$, then the first equation obtained solutions will satisfy $\mathbf{w}_1 = \mathbf{0}$. Thus the ensemble will discard the shallowest network if it is "unnecessary" to satisfy the constraint.

Furthermore, from eq. 14 we conjecture that after discarding the shallowest "unnecessary" network, the ensemble will tend to minimize $\|\mathbf{w}_2\|$, i.e., to discard the second shallowest "unnecessary" network. This will continue until there are no more "unnecessary" shallow networks. In other words, we conjecture that the an ensemble of neural networks will aim to discard the shallowest "unnecessary" networks.

Additionally, using Theorem 1 we can now represent hard-margin SVM problems **with unregularized bias**. Previous results only focused on linear prediction functions without bias. Trying to extend these results to SVM with bias by extending all the input vectors \mathbf{x}_n with an additional '1' component would fail since the obtained solution in the original \mathbf{x} space is the solution of

$$\arg \min_{\mathbf{w} \in \mathbb{R}^d, b \in \mathbb{R}} \|\mathbf{w}\|^2 + b^2 \text{ s.t. } y_n (\mathbf{w}^\top \mathbf{x}_n + b) \geq 1,$$

which is not the L_2 max-margin (SVM) solution, as pointed out by (Nar et al., 2019). However, we can now achieve this goal using Theorem 1. For some dataset $\{\mathbf{x}_n, y_n\}_{n=1}^N$, $\mathbf{x}_n \in \mathbb{R}^d$, $y_n \in \{-1, 1\}$, we use the following prediction function $f_n(\theta) = y_n (\theta_1^\top \mathbf{x}_n + b^2)$ where $\theta = [\theta_1, b]$. From eqs. 12, 13 the asymptotic solution will satisfy $\arg \min_{\theta_1, b} \|\theta_1\|^2 \text{ s.t. } \forall n : y_n (\theta_1^\top \mathbf{x}_n + b^2) \geq 1$.

4. Homogeneous Models

In the previous section we connected the constrained path to the margin path. We would like to refine this characterization and also understand the connection to the optimization path. In this section we are able to do so for prediction functions $f_n(\boldsymbol{\theta})$ which are α -positive homogeneous functions (definition 1).

In the homogeneous case, eq. 7 is equivalent, $\forall \rho$, to

$$\Theta_m^* = \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \|\boldsymbol{\theta}\|^2 \text{ s.t. } \min_n f_n(\boldsymbol{\theta}) \geq \gamma^*(1) \quad (18)$$

since f_n is homogeneous.

4.1. Optimization Path Converges to Stationary Points of the Margin Path and Constrained Path

Remark: The results in this subsection are specific for the Euclidean or L_2 norm, as opposed to many of the results in this paper which are stated for any norm.

In this section, we link the optimization path to the margin path and the constrained path. These results require the following smoothness assumption:

Assumption 3 (Smoothness). *We assume $f_n(\cdot)$ is a \mathcal{C}^2 function.*

Relating optimization path and margin path. The limit of the margin path for homogeneous models is given by eq. 18. In this section we first relate the optimization path to this limit of margin path.

Note that for general homogeneous prediction functions f_n , eq. 18 is a non-convex optimization problem, and thus it is unlikely for an optimization algorithm such as gradient descent to find the global optimum. We can relax the set to $\boldsymbol{\theta}$ that are first-order stationary, i.e., critical points of 18. For $\boldsymbol{\theta} \in \mathbb{S}^{d-1}$, denote the set of support vectors of $\boldsymbol{\theta}$ as

$$S_m(\boldsymbol{\theta}) = \{n : f_n(\boldsymbol{\theta}) = \gamma^*(1)\}. \quad (19)$$

Definition 3 (First-order Stationary Point). *The first-order optimality conditions of 18 are:*

1. $\forall n, f_n(\boldsymbol{\theta}) \geq \gamma^*(1)$
2. *There exists $\boldsymbol{\lambda} \in \mathbb{R}_+^N$ such that $\boldsymbol{\theta} = \sum_n \lambda_n \nabla f_n(\boldsymbol{\theta})$ and $\lambda_n = 0$ for $n \notin S_m(\boldsymbol{\theta})$.*

We denote by Θ_m^s the set of first-order stationary points.

Let $\boldsymbol{\theta}(t)$ be the iterates of gradient descent. Define $\ell_n(t) = \exp(-f_n(\boldsymbol{\theta}(t)))$ and $\boldsymbol{\ell}(t)$ be the vector with entries $\ell_n(t)$. The following two assumptions assume that the limiting direction $\frac{\boldsymbol{\theta}(t)}{\|\boldsymbol{\theta}(t)\|}$ exist and the limiting direction of the losses $\frac{\boldsymbol{\ell}(t)}{\|\boldsymbol{\ell}(t)\|_1}$ exist. Such assumptions are natural in the context

of max-margin problems, since we want to argue that $\boldsymbol{\theta}(t)$ converges to a max-margin direction, and also the losses $\boldsymbol{\ell}(t)/\|\boldsymbol{\ell}(t)\|_1$ converges to an indicator vector of the support vectors. The first step to argue this convergence is to ensure the limits exist.

Assumption 4 (Asymptotic Formulas). *Assume that $\mathcal{L}(\boldsymbol{\theta}(t)) \rightarrow 0$, that is we converge to a global minimizer. Further assume that $\lim_{t \rightarrow \infty} \frac{\boldsymbol{\theta}(t)}{\|\boldsymbol{\theta}(t)\|_2}$ and $\lim_{t \rightarrow \infty} \frac{\boldsymbol{\ell}(t)}{\|\boldsymbol{\ell}(t)\|_1}$ exist. Equivalently,*

$$\ell_n(t) = h(t)a_n + h(t)\epsilon_n(t) \quad (20)$$

$$\boldsymbol{\theta}(t) = g(t)\bar{\boldsymbol{\theta}} + g(t)\boldsymbol{\delta}(t), \quad (21)$$

with $\|\boldsymbol{a}\|_1 = 1$, $\|\bar{\boldsymbol{\theta}}\|_2 = 1$, $\lim_{t \rightarrow \infty} h(t) = 0$, $\lim_{t \rightarrow \infty} \epsilon_n(t) = 0$, and $\lim_{t \rightarrow \infty} \boldsymbol{\delta}(t) = 0$.

Assumption 5 (Linear Independence Constraint Qualification). *Let $\boldsymbol{\theta} \in \mathbb{S}^{d-1}$ be a unit vector. LICQ holds at $\boldsymbol{\theta}$ if the vectors $\{\nabla f_n(\boldsymbol{\theta})\}_{n \in S_m(\boldsymbol{\theta})}$ are linearly independent.*

Remark 1. *Constraint qualifications allow the first-order optimality conditions of Definition 3 to be a necessary condition for optimality. Without constraint qualifications, the global optimum need not satisfy the optimality conditions.*

LICQ is the simplest among many constraint qualification conditions identified in the optimization literature (Nocedal & Wright, 2006).

For example, in linear SVM, LICQ is ensured if the set of support vectors is linearly independent. Consider $f_n(\boldsymbol{\theta}) = \mathbf{x}_n^\top \boldsymbol{\theta}$ and \mathbf{x}_n be the support vectors. Then $\nabla f_n(\boldsymbol{\theta}) = \mathbf{x}_n$, and so linear independence of the support vectors implies LICQ. For data sampled from an absolutely continuous distribution, the SVM solution will always have linearly independent support vectors (Soudry et al., 2018b, Lemma 12), but LICQ may fail when the data is degenerate.

Theorem 2. *Define $\bar{\boldsymbol{\theta}} = \lim_{t \rightarrow \infty} \frac{\boldsymbol{\theta}(t)}{\|\boldsymbol{\theta}(t)\|_2}$. Under Assumptions 3, 4, and constraint qualification at $\bar{\boldsymbol{\theta}}$ (Assumption 5), $\bar{\boldsymbol{\theta}}$ is a first-order stationary point of 18.*

The proof of Theorem 2 can be found in Appendix E.1.

Optimization path and constrained path. Next, we study how the optimization path as $t \rightarrow \infty$ converges to stationary points of the constrained path with $\rho \rightarrow \infty$.

The first-order optimality conditions of the constrained path $\min_{\|\boldsymbol{\theta}\| \leq 1} \mathcal{L}(\rho\boldsymbol{\theta})$, require that the constraints hold, and the gradient of the Lagrangian of the constrained path

$$\rho \nabla_{\boldsymbol{\theta}} \mathcal{L}(\rho\boldsymbol{\theta}) + \lambda(\rho)\boldsymbol{\theta} \quad (22)$$

is equal to zero. In other words,

Remark 2. *Under Assumption 1, $\boldsymbol{\theta}$ is first-order optimal for the problem $\min_{\|\boldsymbol{\theta}\| \leq 1} \mathcal{L}(\rho\boldsymbol{\theta})$ if it satisfies:*

$$\bullet -\frac{\nabla_{\theta} \mathcal{L}(\rho\theta)}{\|\nabla_{\theta} \mathcal{L}(\rho\theta)\|} = \frac{\theta}{\|\theta\|}, \quad \bullet \|\theta\| = 1.$$

On many paths the gradient of the Lagrangian goes to zero as $\rho \rightarrow \infty$. However, we have a faster vanishing rate for the specific optimization paths that follow Definition 4 below. Therefore, these paths better approximate true stationary points:

Definition 4 (First-order optimal for $\rho \rightarrow \infty$). A sequence $\tilde{\theta}(t)$ is first-order optimal for $\min_{\|\theta\| \leq 1} \mathcal{L}(\rho\theta)$ with $\rho \rightarrow \infty$ if

$$\bullet \lim_{t \rightarrow \infty} -\frac{\nabla_{\theta} \mathcal{L}(\rho\tilde{\theta}(t))}{\|\nabla_{\theta} \mathcal{L}(\rho\tilde{\theta}(t))\|} = \lim_{t \rightarrow \infty} \frac{\tilde{\theta}(t)}{\|\tilde{\theta}(t)\|}, \quad \bullet \|\tilde{\theta}(t)\| = 1.$$

To relate the limit points of gradient decent to the constrained path, we will focus on stationary points of the constrained path that minimize the loss.

Theorem 3. Let $\bar{\theta} = \lim_{t \rightarrow \infty} \frac{\theta(t)}{\|\theta(t)\|}$ be the limit direction of gradient descent. Under Assumptions 1, 3, 4, and constraint qualification at $\bar{\theta}$ (Assumption 5), the sequence $\theta(t)/\|\theta(t)\|$ is a first-order optimal point for $\rho \rightarrow \infty$ (Definition 4).

The proof of Theorem 3 can be found in Appendix E.2.

4.2. Lexicographic Max-Margin

Recall that for positive homogeneous prediction functions, the margin path $\Theta_m(\rho)$ in eq. 11 is the same set for any ρ and is given by

$$\Theta_m^* = \arg \max_{\theta \in \mathbb{S}^{d-1}} \min_n f_n(\theta).$$

For non-convex functions f_n or non-Euclidean norms $\|\cdot\|$, the above set need not be unique. In this case, we define the following refined set of maximum margin solution set

Definition 5 (Lexicographic maximum margin set). The lexicographic margin set denoted by $\Theta_{m,N}^*$ is given by the following iterative definition of $\Theta_{m,k}^*$ for $k = 1, 2, \dots, N$:

$$\Theta_{m,0}^* = \mathbb{S}^{d-1},$$

$$\Theta_{m,k}^* = \arg \max_{\theta \in \Theta_{m,k-1}^*} \left(\min_{\{n_{\ell}\}_{\ell=1}^k} \max_{\ell \in [k]} f_{n_{\ell}}(\theta) \right) \subseteq \Theta_{m,k-1}^*.$$

In the above definition, $\Theta_{m,1}^* = \Theta_m^*$ denotes the set of maximum margin solutions, $\Theta_{m,1}^*$ denotes the subset of $\Theta_{m,1}^*$ with second smallest margin, and so on.

For an alternate representation of $\Theta_{m,k}^*$, we introduce the following notation: for $\theta \in \mathbb{S}^{d-1}$, let $n_{\ell}^*(\theta) \in [N]$ denote the index corresponding to the ℓ^{th} smallest margin of θ as defined below by breaking ties in the $\arg \min$ arbitrarily:

$$\begin{aligned} n_1^*(\theta) &= \arg \min_n f_n(\theta) \\ n_k^*(\theta) &= \arg \min_{n \notin \{n_{\ell}^*(\theta)\}_{\ell=1}^{k-1}} f_n(\theta) \quad \text{for } k \geq 2. \end{aligned} \quad (23)$$

Using this notation, we can rewrite $\Theta_{m,k+1}^*$ as

$$\Theta_{m,k+1}^* = \arg \max_{\theta \in \Theta_{m,k}^*} f_{n_{k+1}^*}(\theta).$$

We also define the limit set of constrained path as follows:

Definition 6 (Limit set of constrained path). The limit set of constrained path is defined as follows:

$$\Theta_c^{\infty} = \left\{ \theta : \exists \{\rho_i, \theta_{\rho_i}\}_{i=1}^{\infty} \text{ with } \rho_i \rightarrow \infty, \theta_{\rho_i} \in \Theta_c(\rho_i) \right\}.$$

Theorem 4. For α -positive homogeneous prediction functions the limit set of constrained path is contained in the lexicographic maximum margin set, i.e., $\Theta_c^{\infty} \subseteq \Theta_{m,N}^*$.

The proof of the above Theorem follows from adapting the arguments of (Rosset et al., 2004a) (Theorem 7 in Appendix B.2) for general homogeneous models. We show the complete proof in Appendix E.3.

5. Summary

In this paper we characterized the connections between the constrained, margin and optimization paths. First, in Section 3, we examined general non-homogeneous models. We showed that the margin of the constrained path solution converges to the maximum margin. We further analyzed this result and demonstrated how it implies convergence in parameters, i.e., $\Theta_c(\rho)$ converges to $\Theta_m(\rho)$, for some models. Then, we examined functions that are a finite sum of positively homogeneous functions. These prediction function can represent an ensemble of neural networks with positive homogeneous activation functions. For this model, we characterized the asymptotic constrained path and margin path solution. This implies a surprising result: ensembles of neural networks will aim to discard the most shallow network. In the future work we aim to analyze sum of homogeneous functions with shared variables, such as ResNets.

Second, in Section 4 we focus on homogeneous models. For such models we link the optimization path to the margin and constrained paths. Particularly, we show that the optimization path converges to stationary points of the constrained path and margin path. In future work, we aim to extend this to non-homogeneous models. In addition, we give a more refined characterization of the constrained path limit. It will be interesting to find whether this characterization be further refined to answer whether the weighting of the data point can have any effect on the selection of the asymptotic solution — as (Byrd & Lipton, 2018) observed empirically that it did not.

Acknowledgements

The authors are grateful to C. Zeno, and N. Merlis for helpful comments on the manuscript. This research was supported by the Israel Science foundation (grant No. 31/1031), and by the Taub foundation. SG and NS were partially supported by NSF awards IIS-1302662 and IIS-1764032.

References

- Ali, A., Kolter, J. Z., and Tibshirani, R. J. A continuous-time view of early stopping for least squares regression. *arXiv preprint arXiv:1810.10082*, 2018.
- Byrd, J. and Lipton, Z. C. Weighted risk minimization & deep learning. *arXiv preprint arXiv:1812.03372*, 2018.
- Gunasekar, S., Woodworth, B. E., Bhojanapalli, S., Neyshabur, B., and Srebro, N. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, pp. 6152–6160, 2017.
- Gunasekar, S., Lee, J., Soudry, D., and Srebro, N. Implicit bias of gradient descent on linear convolutional networks. *NIPS*, 2018a.
- Gunasekar, S., Lee, J. D., Soudry, D., and Srebro, N. Characterizing implicit bias in terms of optimization geometry. *ICML*, 2018b.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hoffer, E., Hubara, I., and Soudry, D. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In *NIPS*, 2017.
- Ji, Z. and Telgarsky, M. Risk and parameter convergence of logistic regression. *arXiv preprint arXiv:1803.07300*, 2018.
- Ji, Z. and Telgarsky, M. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations*, 2019.
- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., and Tang, P. T. P. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. *ICLR*, pp. 1–16, 2017.
- Nacson, M. S., Lee, J., Gunasekar, S., Savarese, P. H., Srebro, N., and Soudry, D. Convergence of gradient descent on separable data. *AISTATS*, 2019a.
- Nacson, M. S., Srebro, N., and Soudry, D. Stochastic gradient descent on separable data: Exact convergence with a fixed learning rate. *AISTATS*, 2019b.
- Nar, K., Ocal, O., Sastry, S. S., and Ramchandran, K. Cross-entropy loss leads to poor margins, 2019.
- Neyshabur, B., Salakhutdinov, R. R., and Srebro, N. Path-SGD: Path-normalized optimization in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 2422–2430, 2015a.
- Neyshabur, B., Tomioka, R., and Srebro, N. In search of the real inductive bias: On the role of implicit regularization in deep learning. In *International Conference on Learning Representations*, 2015b.
- Neyshabur, B., Tomioka, R., Salakhutdinov, R., and Srebro, N. Geometry of optimization and implicit regularization in deep learning. *arXiv preprint arXiv:1705.03071*, 2017.
- Nocedal, J. and Wright, S. Numerical optimization. *Springer Science*, 35(67-68), 2006.
- Rosset, S., Zhu, J., and Hastie, T. Boosting as a regularized path to a maximum margin classifier. *Journal of Machine Learning Research*, 2004a.
- Rosset, S., Zhu, J., and Hastie, T. J. Margin maximizing loss functions. In *Advances in neural information processing systems*, pp. 1237–1244, 2004b.
- Soudry, D., Hoffer, E., Shpigel Nacson, M., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *JMLR*, 2018a.
- Soudry, D., Hoffer, E., and Srebro, N. The implicit bias of gradient descent on separable data. *ICLR*, 2018b.
- Suggala, A., Prasad, A., and Ravikumar, P. K. Connecting optimization and regularization paths. In *Advances in Neural Information Processing Systems*, pp. 10631–10641, 2018.
- Wei, C., Lee, J. D., Liu, Q., and Ma, T. On the margin theory of feedforward neural networks. *arXiv preprint arXiv:1810.05369v1*, pp. 1–34, 2018.
- Wilson, A. C., Roelofs, R., Stern, M., Srebro, N., and Recht, B. The marginal value of adaptive gradient methods in machine learning. *arXiv preprint arXiv:1705.08292*, 2017.
- Wu, L., Zhu, Z., and E, W. Towards Understanding Generalization of Deep Learning: Perspective of Loss Landscapes. *arXiv*, 2017.

Xu, T., Zhou, Y., Ji, K., and Liang, Y. When will gradient methods converge to max-margin classifier under reLU models?, 2019.

Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.