ON THE MARGIN THEORY OF FEEDFORWARD NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Past works have shown that, somewhat surprisingly, over-parametrization can help generalization in neural networks. Towards explaining this phenomenon, we adopt a margin-based perspective. We establish: 1) for multi-layer feedforward relu networks, the global minimizer of a weakly-regularized cross-entropy loss has the maximum normalized margin among all networks, 2) as a result, increasing the over-parametrization improves the normalized margin and generalization error bounds for two-layer networks. In particular, an infinite-size neural network enjoys the best generalization guarantees. The typical infinite feature methods are kernel methods; we compare the neural net margin with that of kernel methods and construct natural instances where kernel methods have much weaker generalization guarantees. We validate this gap between the two approaches empirically. Finally, this infinite-neuron viewpoint is also fruitful for analyzing optimization. We show that a perturbed gradient flow on infinite-size networks finds a global optimizer in polynomial time.

1 INTRODUCTION

In deep learning, over-parametrization refers to the widely-adopted technique of using more parameters than necessary (Krizhevsky et al., 2012; Livni et al., 2014). Both computationally and statistically, over-parametrization is crucial for learning neural nets. Controlled experiments demonstrate that over-parametrization eases optimization by smoothing the non-convex loss surface (Livni et al., 2014; Sagun et al., 2017). Statistically, increasing model size without any regularization still *improves* generalization even after the model interpolates the data perfectly (Neyshabur et al., 2017b). This is surprising given the conventional wisdom on the trade-off between model capacity and generalization.

In the absence of an explicit regularizer, algorithmic regularization is likely the key contributor to good generalization. Recent works have shown that gradient descent finds the minimum norm solution fitting the data for problems including logistic regression, linearized neural networks, and matrix factorization (Soudry et al., 2018; Gunasekar et al., 2018b; Li et al., 2018; Gunasekar et al., 2018a; Ji & Telgarsky, 2018). Many of these proofs require a delicate analysis of the algorithm's dynamics, and some are not fully rigorous due to assumptions on the iterates. To the best of our knowledge, it is an open question to prove analogous results for even two-layer relu networks. (For example, the technique of Li et al. (2018) on two-layer neural nets with quadratic activations still falls within the realm of linear algebraic tools, which apparently do not suffice for other activations.)

We propose a different route towards understanding generalization: *making the regularization explicit*. The motivations are: 1) with an explicit regularizer, we can analyze generalization without fully understanding optimization; 2) it is unknown whether gradient descent provides additional implicit regularization beyond what ℓ_2 regularization already offers; 3) on the other hand, with a sufficiently weak ℓ_2 regularizer, we can prove stronger results that apply to multi-layer neural nets with relu activations. Additionally, explicit regularization is perhaps more relevant because ℓ_2 regularization is typically used in practice.

Concretely, we add a norm-based regularizer to the cross entropy loss of a multi-layer feedforward neural network with relu activations. We show that the global minimizer of the regularized objective achieves the maximum normalized margin among all the models with the same architecture, if the regularizer is sufficiently weak (Theorem 2.1). Informally, for models with norm 1 that perfectly classify the data, the margin is the smallest difference across all datapoints between the classifier

score for the true label and the next best score. We are interested in normalized margin because its inverse bounds the generalization error (see recent work (Bartlett et al., 2017; Neyshabur et al., 2017a; 2018) and our Theorem 3.1). Our work explains why optimizing the training loss can lead to parameters with a large margin and thus, better generalization error.

At a first glance, it might seem counterintuitive that *decreasing* the regularizer is the right approach. At a high level, we show that the regularizer only serves as a tiebreaker to steer the model towards choosing the largest normalized margin. Our proofs are simple, oblivious to the optimization procedure, and apply to any norm-based regularizer. We also show that an exact global minimum is unnecessary: if we approximate the minimum loss within a constant, we obtain the max-margin within a constant (Theorem 2.2).

We further study the margin of two-layer networks: let $\gamma^{\star,m}$ be the max normalized margin of a neural net with m hidden units (formally defined in Section 3.1). Let $\gamma^{\star,\infty} \triangleq \sup_m \gamma^{\star,m}$ be the largest possible margin of an infinite two-layer network. We will show three properties of the margins:

1. In Theorem 3.2, we show that the optimal normalized margin of two-layer networks is non-decreasing as the width of the architecture grows, so the generalization error bound only improves with a wider network. Thus, even if the dataset is already separable, it could still be useful to increase the width to achieve larger margin and better generalization. More formally, let n be the number of training examples. We additionally approach the maximum possible margin $\gamma^{\star,\infty}$ after over-parameterizing with $m \geq n$ neurons: $\forall m \geq n, \gamma^{\star,m} = \gamma^{\star,\infty}$.
2. The max-margin of infinite-size nets, $\gamma^{\star,\infty}$, equals half the margin of the ℓ_1-norm SVM (Zhu et al., 2004) over the lifted feature space defined by the activation function applied to all possible hidden units. (See Theorem 3.3.)
3. We compare the neural net margin $\gamma^{\star,\infty}$ to the standard margin for the kernel SVM on the same features. We design a simple data distribution (Figure 1) where neural net margin $\gamma^{\star,\infty}$ is large but the kernel margin is small. This translates to an $\Omega(\sqrt{d})$ factor gap between the generalization error bounds for the two approaches and demonstrates the power of neural nets compared to kernel methods. We experimentally confirm that a gap does indeed exist.

In the context of bullet 2, our work is closely related to that of Rosset et al. (2007) and Neyshabur et al. (2014), who show that optimizing the loss over the parameters of a two-layer relu network is equivalent to optimizing the loss of a "convex neural net" parametrized by a distribution over hidden units. We go one step further and connect the weakly regularized training loss to the ℓ_1 SVM.

We will also adopt this view of infinite-size neural networks to study how over-parametrization helps optimization. Prior works (Mei et al., 2018; Chizat & Bach, 2018) show that gradient descent on two-layer networks becomes Wasserstein gradient flow over parameter distributions in the limit of infinite neurons. For this setting, we prove that perturbed Wasserstein gradient flow finds a global optimizer in polynomial time.

Finally, we empirically validate several of the claims made in this paper. First, we train a two-layer network on a one-dimensional classification task that is simple to visualize. In one dimension, it is possible to brute-force approximate the maximum neural network margin and we show that training with an progressively smaller regularizer results in convergence to this margin. Second, we compare the generalization performance of neural networks and kernel methods and confirm that neural networks do achieve better generalization, as our theory predicts.

1.1 Additional Related Work

Zhang et al. (2016) and Neyshabur et al. (2017b) show that neural network generalization defies conventional explanations and requires new ones. One proposed explanation is the inductive bias of the training algorithm. Recent papers (Hardt et al., 2015; Brutzkus et al., 2017; Chaudhari et al., 2016) study inductive bias through training time and sharpness of local minima. Neyshabur et al. (2015a) propose a new steepest descent algorithm in a geometry invariant to weight rescaling and show that this improves generalization. Morcos et al. (2018) relate generalization in deep nets to the number of "directions" in the neurons. Other papers (Gunasekar et al., 2017; Soudry et al., 2018; Gunasekar et al., 2018b; Li et al., 2018; Gunasekar et al., 2018a) study implicit regularization towards a specific solution. Ma et al. (2017) show that implicit regularization can help gradient descent avoid

overshooting optima. Rosset et al. (2004) study logistic regression with a weak regularization and show convergence to the max margin solution. We adopt their techniques and extend their results.

Recent works have also derived tighter Rademacher complexity bounds for deep neural networks (Neyshabur et al., 2015b; Bartlett et al., 2017; Neyshabur et al., 2017a; Golowich et al., 2017) and new compression based generalization properties (Arora et al., 2018b). Dziugaite & Roy (2017) manage to compute non-vacuous generalization bounds from PAC-Bayes bounds. Neyshabur et al. (2018) investigate the Rademacher complexity of two-layer networks. Liang & Rakhlin (2018) and Belkin et al. (2018) study the generalization of kernel methods.

On the optimization side, Soudry & Carmon (2016) explain why over-parametrization can remove bad local minima. Safran & Shamir (2016) show that over-parametrization can improve the quality of the random initialization. Haeffele & Vidal (2015), Nguyen & Hein (2017), and Venturi et al. (2018) show that for sufficiently overparametrized networks, all local minima are global, but do not show how to find these minima via gradient descent. Du & Lee (2018) show that for two-layer networks with quadratic activations, all second-order stationary points are global minimizers. Arora et al. (2018a) interpret over-parametrization as a means of implicit acceleration during optimization. Mei et al. (2018), Chizat & Bach (2018), and Sirignano & Spiliopoulos (2018) take a distributional view of over-parametrized networks. Chizat & Bach (2018) show that Wasserstein gradient flow converges to global optimizers under structural assumptions. We extend this to a polynomial-time result.

1.2 Notation

Let $\mathbb{R}$ denote the set of real numbers. We will use $\|\cdot\|$ to indicate a general norm, with $\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty$ denoting the $\ell_1, \ell_2, \ell_\infty$ norms on finite dimensional vectors, respectively, and $\|\cdot\|_F$ denoting the Frobenius norm on a matrix. In general, we use $\bar{\ }$ on top of a symbol to denote a unit vector: when applicable, $\bar{u} \triangleq u/\|u\|$, where the norm $\|\cdot\|$ will be clear from context. Let $\mathbb{S}^{d-1} \triangleq \{\bar{u} \in \mathbb{R}^d : \|\bar{u}\|_2 = 1\}$ be the unit sphere in d dimensions. Let $\mathcal{L}^p(\mathbb{S}^{d-1})$ be the space of functions on $\mathbb{S}^{d-1}$ for which the p-th power of the absolute value is Lebesgue integrable. For $\alpha \in \mathcal{L}^p(\mathbb{S}^{d-1})$, we overload notation and write $\|\alpha\|_p \triangleq \left(\int_{\mathbb{S}^{d-1}} |\alpha(\bar{u})|^p d\bar{u}\right)^{1/p}$. Additionally, for $\alpha_1 \in \mathcal{L}^1(\mathbb{S}^{d-1})$ and $\alpha_2 \in \mathcal{L}^\infty(\mathbb{S}^{d-1})$ or $\alpha_1, \alpha_2 \in \mathcal{L}^2(\mathbb{S}^{d-1})$, we can define $\langle \alpha_1, \alpha_2 \rangle \triangleq \int_{\mathbb{S}^{d-1}} \alpha_1(\bar{u})\alpha_2(\bar{u}) d\bar{u} < \infty$. Furthermore, we will use $\mathrm{Vol}(\mathbb{S}^{d-1}) \triangleq \int_{\mathbb{S}^{d-1}} 1 d\bar{u}$.

Throughout this paper, we reserve the symbol $X = [x_1, \ldots, x_n]$ to denote the collection of datapoints (as a matrix), and $Y = [y_1, \ldots, y_n]$ to denote labels. We use d to denote the dimension of our data. We often use Θ to denote the parameters of a prediction function f, and $f(\Theta; x)$ to denote the prediction of f on datapoint x.

We will use the notation $\lesssim, \gtrsim$ to mean less than or greater than up to a universal constant, respectively. Unless stated otherwise, we use $O(\cdot), \Omega(\cdot)$ as a placeholders for some universal constant in upper and lower bounds, respectively. We will use poly to denote some universal constant-degree polynomial in the arguments.

2 Weak Regularizer Guarantees Max Margin Solutions

In this section, we will show that when we add a weak regularizer to cross-entropy loss with a positive-homogeneous prediction function, the normalized margin of the optimum converges to some max-margin solution. As a concrete example, feedforward relu networks are positive-homogeneous.

Let l be the number of labels, so the i-th example has label $y_i \in [l]$. We work with a family $\mathcal{F}$ of prediction functions $f(\Theta; \cdot) : \mathbb{R}^d \to \mathbb{R}^l$ that are a-positive-homogeneous in their parameters for some $a > 0$: $f(c\Theta; x) = c^a f(\Theta; x), \forall c > 0$. We additionally require that f is continuous in Θ. For some general norm $\|\cdot\|$, we study the λ-regularized cross-entropy loss L_λ, defined as

$$L_\lambda(\Theta) \triangleq \sum_{i=1}^{n} -\log \frac{\exp(f_{y_i}(\Theta; x_i))}{\sum_{j=1}^{l} \exp(f_j(\Theta; x_i))} + \lambda \|\Theta\|^r \tag{2.1}$$

for fixed $r > 0$. Let $\Theta_\lambda \in \arg\min L_\lambda(\Theta)$.[1] We define the normalized margin of Θ_λ as:

$$\gamma_\lambda \triangleq \min_i \left(f_{y_i}(\bar{\Theta}_\lambda; x_i) - \max_{j \neq y_i} f_j(\bar{\Theta}_\lambda; x_i) \right) \tag{2.2}$$

Define the $\|\cdot\|$-max normalized margin as

$$\gamma^\star \triangleq \max_{\|\Theta\| \leq 1} \left[\min_i \left(f_{y_i}(\Theta; x_i) - \max_{j \neq y_i} f_j(\Theta; x_i) \right) \right]$$

and let $\Theta^\star$ be a parameter achieving this maximum. We show that with sufficiently small regularization level λ, the normalized margin γ_λ approaches the maximum margin $\gamma^\star$.

Theorem 2.1. *Assume the training data is separable by a network $f(\Theta^\star; \cdot) \in \mathcal{F}$ with an optimal normalized margin $\gamma^\star > 0$. Then, the normalized margin of the global optimum of the weakly-regularized objective (equation 2.1) converges to $\gamma^\star$ as the strength of the regularizer goes to zero. Mathematically, let γ_λ be defined in equation 2.2. Then*

$$\gamma_\lambda \to \gamma^\star \text{ as } \lambda \to 0$$

An intuitive explanation for our result is as follows: because of the homogeneity, the loss $L(\Theta_\lambda)$ roughly satisfies the following (for small λ, and ignoring problem parameters such as n):

$$L_\lambda(\Theta_\lambda) \approx \exp(-\|\Theta_\lambda\|^a \gamma_\lambda) + \lambda \|\Theta_\lambda\|^r$$

Thus, the loss focuses on choosing parameters with larger margin, and the regularization term biases the loss to select parameters with a smaller norm. The full proof of the theorem is deferred to Section A.1.

We can also provide an analogue of Theorem 2.1 for the binary classification setting. For this setting, our prediction is now a single real output and we train using logistic loss. We provide formal definitions and results in Section A.2. Our theory for two-layer neural networks (see Section 3) is based in this setting.

2.1 Optimization Accuracy

Since L_λ is typically hard to optimize exactly for neural nets, it would be ideal to relax the condition that Θ_λ minimizes L_λ. Thus, we ask, how accurately do we need to optimize L_λ to obtain a margin that approximates $\gamma^\star$ up to a constant? The following theorem shows that if suffices to find Θ' achieving a constant factor multiplicative approximation of $L_\lambda(\Theta_\lambda)$, where λ is some sufficiently small polynomial in $n, l, \gamma^\star$. Though our theorem is stated for the general multi-class setting, our result applies for binary classification as well. We provide the proof in Section A.3.

Theorem 2.2. *In the setting of Theorem 2.1, suppose that we choose $\lambda = \frac{(\gamma^\star)^{r/a}}{n^c (l-1)^c}$ for sufficiently large c (that only depends on r/a). Let Θ' denote a 2-approximate [2] minimizer of L_λ, so $L_\lambda(\Theta') \leq 2L_\lambda(\Theta_\lambda)$. Denote the normalized margin of Θ' by γ'. Then*

$$\gamma' \geq \frac{\gamma^\star}{2 \cdot 4^{a/r}}$$

.

3 Margins of Over-parameterized Two-layer Homogeneous Neural Nets

In Section 2 we showed that a weakly-regularized logistic loss leads to the maximum normalized margin. In this section, we analyze the properties of the max-margin of neural nets more closely. We will contrast neural networks with kernel methods, for which margins have already been extensively studied. Towards a first-cut understanding, we focus on two-layer networks for binary classification.

[1] We formally show that L_λ has a minimizer in Claim A.1 of Section A.

[2] The exact approximation constant is not important, so we choose 2 for simplicity.

First, in Section 3.1 we provide a bound stating that the generalization error is roughly linear in the inverse of the margin, establishing that a larger margin implies better generalization. In Section 3.2, we show that the maximum normalized margin is non-decreasing with the hidden layer size and stays constant as soon as there are more hidden units than data points. This suggests that increasing the size of the network improves the generalization of the solution.

Second, in Section 3.3, we draw an analogy to classical kernel methods by proving that the maximum ℓ_2-normalized margin of an over-parameterized neural net is equal to half the maximum possible ℓ_1-normalized margin of linear functionals on a lifted feature space. In other words, we establish an equivalence between neural networks and the 1-norm SVM (Zhu et al., 2004) on the lifted features. These features are constructed by applying the activation function on all possible hidden layer weights.

Third, continuing this analogy, we will compare the generalization power of a two-layer neural network to that of a kernel method on the lifted space. This kernel method corresponds to fixing random weights for the hidden layer and solving a 2-norm max-margin problem on the top layer weights. We demonstrate instances where two layer neural networks give better generalization error guarantees than the kernel method.

3.1 Setup and Margin-based Generalization Error

In the rest of the paper, we work with two-layer neural networks with a single output for binary classification. We use m to denote the number of hidden units, $w_1, \ldots, w_m \in \mathbb{R}^d$ for the weight vectors on the first layer, and $u_1, \ldots, u_m \in \mathbb{R}$ for the weights on the second layer. We let $\theta_j \triangleq (w_j, u_j)$, and we use Θ to denote the collection of all the parameters. We assume in this section that the activation $\phi : \mathbb{R} \to \mathbb{R}$ is 1-homogeneous and 1-Lipschitz. The network thus computes a single score

$$f(\Theta; x) \triangleq \sum_{j=1}^{m} w_j \phi(u_j^\top x)$$

We consider ℓ_2 regularization from here on. The regularized logistic loss of the architecture with m hidden units is therefore

$$L_{\lambda,m} \triangleq \frac{1}{n} \sum_{i=1}^{n} \log(1 + \exp(-y_i f(\Theta; x_i))) + \lambda \|\Theta\|_2^2 \tag{3.1}$$

where $\|\Theta\|_2$ denotes the Euclidean norm of all the parameters in Θ. We note that f and the regularizer are both 2-homogeneous in Θ, so the results of Section 2 apply to $L_{\lambda,m}$.[3]

Following our conventions from Section 2, we denote the optimizer of $L_{\lambda,m}$ by $\Theta_{\lambda,m}$, the normalized margin of $\Theta_{\lambda,m}$ by $\gamma_{\lambda,m}$, the max-margin solution by $\Theta^{\star,m}$, and the max-margin by $\gamma^{\star,m}$. We emphasize the size of the network in our notation. Since our classifier f now predicts a single real value, we need to redefine

$$\gamma_{\lambda,m} \triangleq \min_i y_i f(\bar{\Theta}_{\lambda,m}; x_i)$$
$$\gamma^{\star,m} \triangleq \max_{\|\Theta\| \le 1} \min_i y_i f(\Theta; x_i)$$

When the data is not separable by a m-unit neural net, $\gamma^{\star,m}$ is zero by definition.

Recall that $X = [x_1, \ldots, x_n]$ denotes the matrix with all the data points as columns, and $Y = [y_1, \ldots, y_n]$ denotes the labels. We sample X and Y i.i.d. from the data generating distribution p_{data}, which is supported on $\mathcal{X} \times \{-1, +1\}$. We can define the population 0-1 loss and the training 0-1 loss of the network Θ as

$$L(\Theta) = \Pr_{(x,y)\sim p_{\text{data}}} [y f(\Theta; x) \le 0]$$

We will let $D \triangleq \frac{\|X\|_F^2}{n}$ be the average norm squared of the data and $C \triangleq \sup_{x \in \mathcal{X}} \|x\|_2$ be an upper bound on the norm of a single datapoint. The following theorem shows that the generalization error only depends on the parameters through the inverse of the margin on the training data. We provide a proof in Section C.1.

[3] Although Theorem 2.1 is written in the language of multi-class prediction where the classifier outputs $l \ge 2$ scores, the results translate to single-output binary classification. See Section A.2.

Theorem 3.1. *Suppose ϕ is 1-Lipschitz and 1-homogeneous. Then for any Θ that separates the data with margin $\gamma \triangleq \min_i y_i f(\bar{\Theta}; x_i) > 0$, with probability at least $1-\delta$ over the draw of X, Y,*

$$L(\Theta) \leq \frac{6}{\gamma}\sqrt{\frac{D}{n}} + \epsilon(\gamma) \tag{3.2}$$

where $\epsilon(\gamma) \triangleq \sqrt{\frac{\log \log_2 \frac{4C}{\gamma}}{n}} + \sqrt{\frac{\log(1/\delta)}{2n}}$. Note that $\epsilon(\gamma)$ is typically small, and thus the above bound mainly scales with $\frac{1}{\gamma}\sqrt{\frac{D}{n}}$. As a corollary, with probability $1-\delta$,[4]

$$\lim_{\lambda \to 0} L(\Theta_{\lambda,m}) \leq \frac{6}{\gamma^{\star,m}}\sqrt{\frac{D}{n}} + \epsilon(\gamma^{\star,m}) \tag{3.3}$$

Above we implicitly assume $\gamma^{\star,m} > 0$, since otherwise the right hand side of the bound is vacuous.

One consequence of the above theorem and Theorem 2.2 is that if λ is polynomially small in $\gamma^{\star,m}$ and n, we only need to optimize $L_{\lambda,m}$ up to a constant multiplicative factor to obtain parameters with generalization bounds roughly as good as those for $\Theta^{\star,m}$.

3.2 The max margin is non-decreasing in the hidden layer size

Now we show that the maximum normalized margin is nondecreasing with the hidden layer size and stays constant once we have more hidden units than examples.

Theorem 3.2. *In the setting of Section 3.1, recall that $\gamma^{\star,m}$ denotes the max normalized margin of a two-layer neural network with hidden layer size m. Then,*

$$\gamma^{\star,1} \leq \gamma^{\star,2} \cdots \leq \gamma^{\star,n} = \gamma^{\star,n+1} = \gamma^{\star,n+2} = \cdots \tag{3.4}$$

We note that $\gamma^{\star,n}$ will be positive when ϕ is a sufficiently powerful activation such as relu or sigmoid and the data points are not repetitive, so the neural network can fit any function of the data. We prove Theorem 3.2 in Section B. Theorem 3.2 can explain why additional over-parametrization has been observed to improve generalization in two-layer networks Neyshabur et al. (2017b). Our margin does not decrease with a larger network size, and therefore Theorem 3.1 gives a better generalization bound. We precisely characterize the value of $\gamma^{\star,n}$ in the following section.

3.3 The max margin of neural nets is equivalent to ℓ_1 SVM in lifted space

We link infinite-size neural networks to the ℓ_1 SVM over a lifted space, defined via a lifting function $\varphi : \mathbb{R}^d \to \mathcal{L}^\infty(\mathbb{S}^{d-1})$ mapping data to an infinite feature vector:

$$x \in \mathbb{R}^d \to \varphi(x) \in \mathcal{L}^\infty(\mathbb{S}^{d-1}) \text{ satisfying } \varphi(x)[\bar{u}] = \phi(\bar{u}^\top x) \tag{3.5}$$

We look at the margin of linear functionals corresponding to $\alpha \in \mathcal{L}^1(\mathbb{S}^{d-1})$. The 1-norm SVM over the lifted feature $\varphi(x)$ solves for the maximum margin:

$$\begin{aligned} \gamma_{\ell_1} &\triangleq \max_{\alpha} \min_{i \in [n]} y_i \langle \alpha, \varphi(x_i) \rangle \\ &\text{subject to } \|\alpha\|_1 \leq 1 \end{aligned} \tag{3.6}$$

where we rely on the inner product and 1-norm defined in Section 1.2. A priori, it is unclear how to optimize this since the kernel trick does not work for ℓ_1 norm. Here we will show that optimizing two-layer neural networks with weak regularization is equivalent to solving equation 3.6.

Theorem 3.3. *Let γ_{ℓ_1} be defined in equation 3.6, and $\gamma^{\star,m}$ be defined in Section 3.1. For any $m \geq n$,*

$$\gamma^{\star,m} = \frac{\gamma_{\ell_1}}{2} \tag{3.7}$$

Rosset et al. (2007) and Neyshabur et al. (2014) show a similar equivalence, but between a lifted logistic regression problem and equation 3.1. In contrast, the above theorem, proved in Section B, shows the equivalence between equation 3.1 and the 1-norm SVM when the regularizer is small.

3.4 Comparison to kernel methods

[4]The quantity $\lim_{\lambda \to 0} L(\Theta_{\lambda,m})$ does not necessarily exist, but here we take it to mean $\forall \delta > 0, \exists \lambda_0 > 0$ with $L(\Theta_{\lambda,m})$ at most the RHS of equation 3.3 plus δ for all $\lambda < \lambda_0$.

We compare the ℓ_1 SVM margin, attainable by a finite neural network, to the ℓ_2 margin attainable via kernel methods. Following the setup of Section 3.3, we define the kernel problem over $\alpha \in \mathcal{L}^2(\mathbb{S}^{d-1})$:

$$\gamma_{\ell_2} \triangleq \max_{\alpha} \min_{i \in [n]} y_i \langle \alpha, \varphi(x_i) \rangle \quad \text{subject to } \sqrt{\kappa}\|\alpha\|_2 \leq 1 \tag{3.8}$$

where $\kappa \triangleq \mathrm{Vol}(\mathbb{S}^{d-1})$. (We scale $\|\alpha\|_2$ by $\sqrt{\kappa}$ to make the lemma statement below cleaner.) First, γ_{ℓ_2} can be used to obtain a standard upper bound on the generalization error of the kernel SVM. Following the notation of Section 3.1, we will let $L_{\ell_2\text{-svm}}$ denote the 0-1 population classification error for the optimizer of equation 3.8.

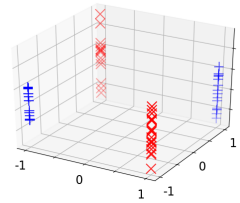

Figure 1: A visualization of 60 sampled points from $\mathcal{D}$ in 3 dimensions. Red points denote negative examples and blue points denote positive examples.

Lemma 3.4. *In the setting of Theorem 3.1, with probability at least $1-\delta$, the generalization error of the standard kernel SVM with* relu *feature (defined in equation 3.8) is bounded by*

$$L_{\ell_2\text{-svm}} \lesssim \frac{1}{\gamma_{\ell_2}} \sqrt{\frac{D}{dn}} + \epsilon_{\ell_2} \tag{3.9}$$

where $\epsilon_{\ell_2} \triangleq \sqrt{\frac{\log \max\left\{\log_2 \frac{C\sqrt{\kappa/d}}{\gamma_{\ell_2}}, 2\right\}}{n}} + \sqrt{\frac{\log(1/\delta)}{n}}$ *is typically a lower-order term.*

The bound above follows from standard techniques (Bartlett & Mendelson, 2002), and we provide a full proof in Section C.1. We construct a data distribution for which this lemma does not give a good bound for kernel methods, but Theorem 3.1 does imply good generalization for two-layer networks.

Theorem 3.5. *There exists a data distribution p_{data} such that the ℓ_1 SVM with* relu *features has a good margin:*

$$\gamma_{\ell_1} \gtrsim 1$$

and with probability $1-\delta$ over the choice of i.i.d. samples from p_{data}, obtains generalization error

$$L_{\ell_1\text{-svm}} \lesssim \sqrt{\frac{d}{n}} + \epsilon_{\ell_1}$$

where $\epsilon_{\ell_1} \triangleq \sqrt{\frac{\log\log(d\log n)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}}$ *is typically a lower order term. Meanwhile, with high probability the ℓ_2 SVM has a small margin:*

$$\gamma_{\ell_2} \lesssim \max\left\{\sqrt{\frac{\log n}{n}}, 1/d\right\}$$

and therefore the generalization upper bound from Lemma 3.4 is at least

$$\Omega\left(\min\left\{\frac{1}{\log n}, \frac{d}{\sqrt{n}}\right\}\right)$$

We briefly overview the construction of p_{data} here and defer the full proof of Theorem 3.5 to Section D.1.

Proof sketch for Theorem 3.5. We base p_{data} on the distribution $\mathcal{D}$ of examples (x, y) described below. Here e_i is the i-th standard basis vector and we use $x^\top e_i$ to represent the i-coordinate of x (since the subscript is reserved to index training examples).

$$\begin{bmatrix} e_3^\top x \\ \vdots \\ e_d^\top x \end{bmatrix} \sim \mathcal{N}(0, I_{d-2}), \text{ and } \begin{cases} y=+1, & x^\top e_1 = +1, & x^\top e_2 = +1 & \text{w/ prob. } 1/4 \\ y=+1, & x^\top e_1 = -1, & x^\top e_2 = -1 & \text{w/ prob. } 1/4 \\ y=-1, & x^\top e_1 = +1, & x^\top e_2 = -1 & \text{w/ prob. } 1/4 \\ y=-1, & x^\top e_1 = -1, & x^\top e_2 = +1 & \text{w/ prob. } 1/4 \end{cases}$$

Figure 1 shows samples from $\mathcal{D}$ when there are 3 dimensions. From the visualization, it is clear that there is no linear separator for $\mathcal{D}$. As Lemma D.1 shows, a relu network with four neurons can fit this relatively complicated decision boundary. On the other hand, for kernel methods, we prove that the symmetries in $\mathcal{D}$ induce cancellation in feature space. The following lemmas, proved in Section D.1, formalize this cancellation and show that it results in a small margin for kernel methods.

Lemma 3.6 (Margin upper bound tool)**.** *In the setting of Theorem 3.5, we have*

$$\gamma_{\ell_2} \leq \frac{1}{\sqrt{\kappa}} \cdot \left\| \frac{1}{n} \sum_{i=1}^{n} \varphi(x_i) y_i \right\|_2$$

Lemma 3.7. *In the setting of Theorem 3.5, let $(x_i, y_i)_{i=1}^n$ be n i.i.d samples and corresponding labels from $\mathcal{D}$. Let φ be defined in equation 3.5 with $\phi =$ relu. With high probability (at least $1 - dn^{-10}$), we have*

$$\left\| \frac{1}{n} \sum_{i=1}^{n} \varphi(x_i) y_i \right\|_2 \lesssim \sqrt{\kappa/n} \log n + \sqrt{\kappa}/d$$

Combining these lemmas gives us the desired bound on γ_{ℓ_2}.

Gap in regression setting: We are able to prove an even larger $\Omega(\sqrt{n/d})$ gap between neural networks and kernel methods in the regression setting where we wish to interpolate continuous labels. Analogously to the classification setting, optimizing a regularized squared error loss on neural networks is equivalent to solving a minimum 1-norm regression problem (see Theorem D.3). Furthermore, kernel methods correspond to a minimum 2-norm problem. We construct distributions p_{data} where the 1-norm solution will have a generalization error bound of $O(\sqrt{d/n})$, whereas the 2-norm solution will have a generalization error bound that is $\Omega(1)$ and thus vacuous. In Section D.2, we define the 1-norm and 2-norm regression problems. In Theorem D.6 we formalize our construction.

4 Perturbed Wasserstein gradient flow finds global optimizers in polynomial time

In the prior section, we studied the limiting behavior of the generalization of a two-layer network as its width goes to infinity. In this section, we will now study the limiting behavior of the optimization algorithm, gradient descent. Prior work (Mei et al., 2018; Chizat & Bach, 2018) has shown that as the hidden layer size grows to infinity, gradient descent for a finite neural network approaches the Wasserstein gradient flow over distributions of hidden units (defined in equation 4.1). Chizat & Bach (2018) also prove that Wasserstein gradient flow converges to a global optimizer in this setting but do not specify a convergence rate.

We show that a perturbed version of Wasserstein gradient flow converges in *polynomial time*. The informal take-away of this section is that a perturbed version of gradient descent converges in polynomial time on infinite-size neural networks (for the right notion of infinite-size.)

Formally, we optimize the following functional over distributions ρ on $\mathbb{R}^{d+1}$:

$$L[\rho] \triangleq R\left(\int \Phi d\rho\right) + \int V d\rho$$

where $\Phi : \mathbb{R}^{d+1} \to \mathbb{R}^k$, $R : \mathbb{R}^k \to \mathbb{R}$, and $V : \mathbb{R}^{d+1} \to \mathbb{R}$. In this work, we consider 2-homogeneous Φ and V. We will additionally require that R is nonnegative and V is positive on the unit sphere. Finally, we need standard regularity assumptions on R, Φ, and V:

Assumption 4.1 (Regularity conditions on Φ, R, V)**.** Φ and V are differentiable as well as upper bounded and Lipschitz on the unit sphere. R is Lipschitz and its Hessian has bounded operator norm.

We provide more details on the specific parameters (for boundedness, Lipschitzness, etc.) in Section E.1. We note that relu networks satisfy every condition but differentiability of Φ.[5] We can fit a neural network under our framework as follows:

[5]The relu activation is non-differentiable at 0 and hence the gradient flow is not well-defined. Chizat & Bach (2018) acknowledge this same difficulty with relu.

Example 4.2 (Logistic loss for neural networks)**.** We interpret ρ as a distribution over the parameters of the network. Let $k \triangleq n$ and $\Phi_i(\theta) \triangleq w\phi(u^\top x_i)$ for $\theta = (w, u)$. In this case, $\int \Phi d\rho$ is a distributional neural network that computes an output for each of the n training examples (like a standard neural network, it also computes a weighted sum over hidden units). We can compute the distributional version of the regularized logistic loss in equation 3.1 by setting $V(\theta) \triangleq \lambda\|\theta\|_2^2$ and $R(a_1, \ldots, a_n) \triangleq \sum_{i=1}^n \log(1 + \exp(-y_i a_i))$.

We will define $L'[\rho] : \mathbb{R}^{d+1} \to \mathbb{R}$ with $L'[\rho](\theta) \triangleq \langle R'(\int \Phi d\rho), \Phi(\theta)\rangle + V(\theta)$ and $v[\rho](\theta) \triangleq -\nabla_\theta L'[\rho](\theta)$. Informally, $L'[\rho]$ is the gradient of L with respect to ρ, and v is the induced velocity field. For the standard Wasserstein gradient flow dynamics, ρ_t evolves according to

$$\frac{d}{dt}\rho_t = -\nabla \cdot (v[\rho_t]\rho_t) \tag{4.1}$$

where $\nabla\cdot$ denotes the divergence of a vector field. For neural networks, these dynamics formally define continuous-time gradient descent when the hidden layer has infinite size (see Theorem 2.6 of Chizat & Bach (2018), for instance).

We propose the following modification of the Wasserstein gradient flow dynamics:

$$\frac{d}{dt}\rho_t = -\sigma\rho_t + \sigma U^d - \nabla \cdot (v[\rho_t]\rho_t) \tag{4.2}$$

where U^d is the uniform distribution on $\mathbb{S}^d$. In our perturbed dynamics, we add uniform noise over U^d. For infinite-size neural networks, one can informally interpret this as re-initializing a very small fraction of the neurons at every step of gradient descent. We prove convergence to a global optimizer in time polynomial in $1/\epsilon, d$, and the regularity parameters.

Theorem 4.3 (Theorem E.4 with regularity parameters omitted)**.** *Suppose that Φ and V are 2-homogeneous and the regularity conditions of Assumption 4.1 are satisfied. Also assume that from starting distribution ρ_0, a solution to the dynamics in equation 4.2 exists. Define $L^\star \triangleq \inf_\rho L[\rho]$. Let $\epsilon > 0$ be a desired error threshold and choose $\sigma \triangleq \exp(-d\log(1/\epsilon)\text{poly}(k, L[\rho_0] - L^\star))$ and $t_\epsilon \triangleq \frac{d^2}{\epsilon^4}\text{poly}(\log(1/\epsilon), k, L[\rho_0] - L^\star)$, where the regularity parameters for Φ, V, and R are hidden in the* poly$(\cdot)$. *Then, perturbed Wasserstein gradient flow converges to an ϵ-approximate global minimum in t_ϵ time:*

$$\min_{0 \le t \le t_\epsilon} L[\rho_t] - L^\star \le \epsilon.$$

We provide a theorem statement that includes regularity parameters in Section E.1. We prove the theorem in Section E.2.

As a technical detail, Theorem 4.3 requires that a solution to the dynamics exists. We can remove this assumption by analyzing a discrete-time version of equation 4.2:

$$\rho_{t+1} \triangleq \rho_t + \eta(-\sigma\rho_t + \sigma U^d - \nabla \cdot (v[\rho_t]\rho_t))$$

and additionally assuming Φ and V have Lipschitz gradients. In this setting, a polynomial time convergence result also holds. We state the result in Section E.3.

5 Simulations

We first verify the normalized margin convergence on a two-layer networks with one-dimensional input. A single hidden unit computes the following: $x \mapsto a_j\text{relu}(w_j x + b_j)$. We add $\|\cdot\|_2^2$-regularization to a, w, and b and compare the resulting normalized margin to that of an approximate solution of the ℓ_1 SVM problem with features $\text{relu}(wx_i + b)$ for $w^2 + b^2 = 1$. Writing this feature vector is intractable, so we solve an approximate version by choosing 1000 evenly spaced values of (w, b). Our theory predicts that with decreasing regularization, the margin of the neural network converges to the ℓ_1 SVM objective. In Figure 2, we plot this margin convergence and visualize the final networks and ground truth labels. The network margin approaches the ideal one as $\lambda \to 0$, and the visualization shows that the network and ℓ_1 SVM functions are extremely similar.

Next, we experiment on synthetic data in a higher-dimensional setting. For classification and regression, we compare the generalization error and predicted generalization upper bounds[6] (from

[6]We compute the leading term that is linear in the norm or inverse margin.

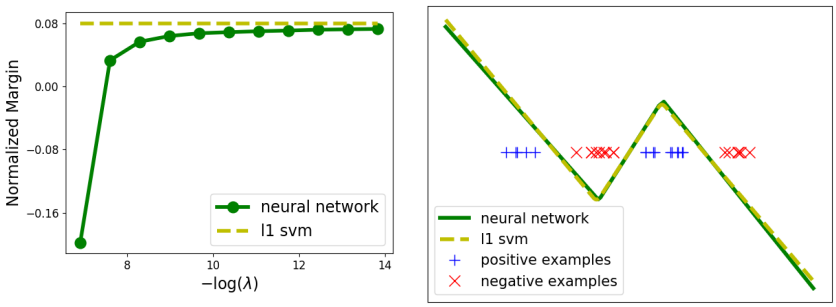

Figure 2: Neural network with input dimension 1. **Left:** Normalized margin as we decrease λ. **Right:** Visualization of the normalized functions computed by the neural network and ℓ_1 SVM solution for $\lambda \approx 10^{-14}$.

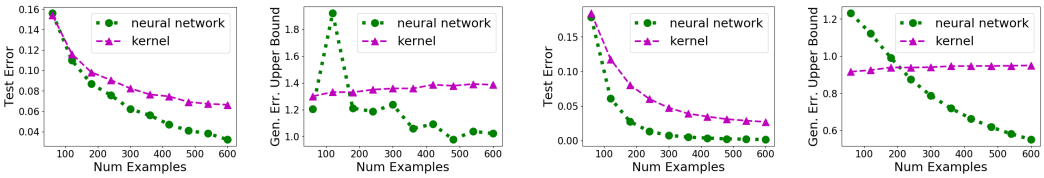

Figure 3: Comparing neural networks and kernel methods. **Left:** Classification. **Right:** Regression.

Theorem 3.1 and Lemmas 3.4, D.4, and D.5) of a trained neural network against a ℓ_2 kernel SVM with relu features as we vary n. For classification we plot 0-1 error, whereas for regression we plot squared error. Our ground truth comes from a random neural network with 6 hidden units. For classification, we used rejection sampling to obtain datapoints with unnormalized margin of at least 0.1 on the ground truth network. We use a fixed dimension of $d = 20$. For all experiments, we train the network for 20000 steps with $\lambda = 10^{-8}$ and average over 100 trials for each plot point.

The plots in Figure 3 show that two-layer networks clearly outperform kernel methods in test error as n grows. However, there seems to be looseness in the upper bounds for kernel methods: the kernel generalization bound appears to stay constant with n (as predicted by our theory for regression), but the kernel test error decreases. There is also some variance in the neural network generalization bound for classification. This occured likely because we did not tune learning rate and training time, so the optimization failed to find the best margin.

In Section F, we include additional experiments training modified WideResNet architectures on CIFAR10 and CIFAR100. Although ResNet is not homogeneous, we still report interesting increases in generalization performance from annealing the weight decay during training, versus staying at a fixed decay rate.

6 Conclusion

We have made the case that maximizing margin is one of the inductive biases of relu networks with cross-entropy loss. We show that we can obtain a maximum normalized margin by training with a weak regularizer. We also prove that larger ℓ_2-normalized margin indicates better generalization for two-layer nets. Our work leaves open the question of how the ℓ_2-normalized margin relates to generalization in much deeper neural networks. This is a fascinating theoretical and empirical question for future work. On the optimization side, we make progress towards understanding over-parametrized gradient descent by analyzing infinite-size neural networks. A natural direction for future work is to apply our theory to optimize the margin of finite-sized neural networks.

REFERENCES

Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. *arXiv preprint arXiv:1802.06509*, 2018a.

Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. *arXiv preprint arXiv:1802.05296*, 2018b.

Keith Ball et al. An elementary introduction to modern convex geometry. *Flavors of geometry*, 31: 1–58, 1997.

Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.

Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pp. 6240–6249, 2017.

Mikhail Belkin, Siyuan Ma, and Soumik Mandal. To understand deep learning we need to understand kernel learning. *arXiv preprint arXiv:1802.01396*, 2018.

Alon Brutzkus, Amir Globerson, Eran Malach, and Shai Shalev-Shwartz. Sgd learns over-parameterized networks that provably generalize on linearly separable data. *arXiv preprint arXiv:1710.10174*, 2017.

Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. *arXiv preprint arXiv:1611.01838*, 2016.

Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. *arXiv preprint arXiv:1805.09545*, 2018.

Simon S Du and Jason D Lee. On the power of over-parametrization in neural networks with quadratic activation. *arXiv preprint arXiv:1803.01206*, 2018.

Simon S Du, Jason D Lee, Yuandong Tian, Barnabas Poczos, and Aarti Singh. Gradient descent learns one-hidden-layer cnn: Don't be afraid of spurious local minima. *arXiv preprint arXiv:1712.00779*, 2017.

Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*, 2017.

Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. *arXiv preprint arXiv:1712.06541*, 2017.

Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, pp. 6151–6159, 2017.

Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. *arXiv preprint arXiv:1802.08246*, 2018a.

Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Implicit bias of gradient descent on linear convolutional networks. *arXiv preprint arXiv:1806.00468*, 2018b.

Benjamin D Haeffele and René Vidal. Global optimality in tensor factorization, deep learning, and beyond. *arXiv preprint arXiv:1506.07540*, 2015.

Moritz Hardt, Benjamin Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. *arXiv preprint arXiv:1509.01240*, 2015.

Ziwei Ji and Matus Telgarsky. Risk and parameter convergence of logistic regression. *arXiv preprint arXiv:1803.07300*, 2018.

Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in neural information processing systems*, pp. 793–800, 2009.

Vladimir Koltchinskii, Dmitry Panchenko, et al. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics*, 30(1):1–50, 2002.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*, pp. 2–47, 2018.

T. Liang and A. Rakhlin. Just Interpolate: Kernel "Ridgeless" Regression Can Generalize. *ArXiv e-prints*, August 2018.

Roi Livni, Shai Shalev-Shwartz, and Ohad Shamir. On the computational efficiency of training neural networks. In *Advances in Neural Information Processing Systems*, pp. 855–863, 2014.

Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. *arXiv preprint arXiv:1711.10467*, 2017.

Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layers neural networks. *arXiv preprint arXiv:1804.06561*, 2018.

Ari S Morcos, David GT Barrett, Neil C Rabinowitz, and Matthew Botvinick. On the importance of single directions for generalization. *arXiv preprint arXiv:1803.06959*, 2018.

Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.

Behnam Neyshabur, Ruslan R Salakhutdinov, and Nati Srebro. Path-sgd: Path-normalized optimization in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 2422–2430, 2015a.

Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pp. 1376–1401, 2015b.

Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nathan Srebro. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1707.09564*, 2017a.

Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pp. 5947–5956, 2017b.

Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. Towards understanding the role of over-parametrization in generalization of neural networks. *arXiv preprint arXiv:1805.12076*, 2018.

Quynh Nguyen and Matthias Hein. The loss surface of deep and wide neural networks. *arXiv preprint arXiv:1704.08045*, 2017.

Saharon Rosset, Ji Zhu, and Trevor Hastie. Boosting as a regularized path to a maximum margin classifier. *Journal of Machine Learning Research*, 5(Aug):941–973, 2004.

Saharon Rosset, Grzegorz Swirszcz, Nathan Srebro, and Ji Zhu. l1 regularization in infinite dimensional feature spaces. In *International Conference on Computational Learning Theory*, pp. 544–558. Springer, 2007.

Itay Safran and Ohad Shamir. On the quality of the initial basin in overspecified neural networks. In *International Conference on Machine Learning*, pp. 774–782, 2016.

Levent Sagun, Utku Evci, V Ugur Guney, Yann Dauphin, and Leon Bottou. Empirical analysis of the hessian of over-parametrized neural networks. *arXiv preprint arXiv:1706.04454*, 2017.

Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks. *arXiv preprint arXiv:1805.01053*, 2018.

Daniel Soudry and Yair Carmon. No bad local minima: Data independent training error guarantees for multilayer neural networks. *arXiv preprint arXiv:1605.08361*, 2016.

Daniel Soudry, Elad Hoffer, and Nathan Srebro. The implicit bias of gradient descent on separable data. In *International Conference on Learning Representations*, 2018. URL `https://openreview.net/forum?id=r1q7n9gAb`.

Ryan J Tibshirani et al. The lasso problem and uniqueness. *Electronic Journal of Statistics*, 7: 1456–1490, 2013.

Twan van Laarhoven. L2 regularization versus batch and weight normalization. *arXiv preprint arXiv:1706.05350*, 2017.

Luca Venturi, Afonso Bandeira, and Joan Bruna. Neural networks with finite intrinsic dimension have no spurious valleys. *arXiv preprint arXiv:1802.06384*, 2018.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

Ji Zhu, Saharon Rosset, Robert Tibshirani, and Trevor J Hastie. 1-norm support vector machines. In *Advances in neural information processing systems*, pp. 49–56, 2004.

A Missing Proofs in Section 2

We first show that L_λ does indeed have a global minimizer.

Claim A.1. *In the setting of Theorems 2.1 and A.3,* $\arg\min_\Theta L_\lambda(\Theta)$ *exists.*

Proof. We will argue in the setting of Theorem 2.1 where L_λ is the multi-class cross entropy loss, because the logistic loss case is analogous. We first note that L_λ is continuous in Θ because f is continuous in Θ and the term inside the logarithm is always positive. Next, define $b \triangleq \inf_\Theta L_\lambda(\Theta) > 0$. Then we note that for $\|\Theta\| > (b/\lambda)^{1/r} \triangleq M$, we must have $L_\lambda(\Theta) > b$. It follows that $\inf_{\|\Theta\|\le M} L_\lambda(\Theta) = \inf_\Theta L_\lambda(\Theta)$. However, there must be a value Θ_λ which attains $\inf_{\|\Theta\|\le M} L_\lambda(\Theta)$, because $\{\Theta : \|\Theta\| \le M\}$ is a compact set and L_λ is continuous. Thus, $\inf_\Theta L_\lambda(\Theta)$ is attained by some Θ_λ. □

A.1 Missing Proofs for Multi-class Setting

Towards proving Theorem 2.1, we first show as we decrease λ, the norm of the solution $\|\Theta_\lambda\|$ grows.

Lemma A.2. *In the setting of Theorem 2.1, as* $\lambda \to 0$*, we have* $\|\Theta_\lambda\| \to \infty$.

To prove Theorem 2.1, we rely on the exponential scaling of the cross entropy: L_λ can be lower bounded roughly by $\exp(-\|\Theta_\lambda\|\gamma_\lambda)$, but also has an upper bound that scales with $\exp(-\|\Theta_\lambda\|\gamma^\star)$. By Lemma A.2, we can take large $\|\Theta_\lambda\|$ so the gap $\gamma^\star - \gamma_\lambda$ vanishes. This proof technique is inspired by that of Rosset et al. (2004).

Proof of Theorem 2.1. For any $M > 0$ and Θ with $\gamma_\Theta \triangleq \min_i \left(f(\bar{\Theta}; x_i) - \max_{j \neq y_i} f(\bar{\Theta}; x_i)\right)$,

$$\begin{aligned} L_\lambda(M\Theta) &= \frac{1}{n}\sum_{i=1}^{n} -\log \frac{\exp(M^a f_{y_i}(\Theta; x_i))}{\sum_{j=1}^{l} \exp(M^a f_j(\Theta; x_i))} + \lambda M^r \|\Theta\|^r \qquad \text{(by the homogeneity of } f\text{)} \\ &= \frac{1}{n}\sum_{i=1}^{n} -\log \frac{1}{1 + \sum_{j \neq y_i} \exp(M^a (f_j(\Theta; x_i) - f_{y_i}(\Theta; x_i)))} + \lambda M^r \|\Theta\|^r \qquad \text{(A.1)} \\ &\leq \log(1 + (l-1)\exp(-M^a \gamma_\Theta)) + \lambda M^r \|\Theta\|^r \qquad \text{(A.2)} \end{aligned}$$

We can also apply $\sum_{j \neq y_i} \exp(M^a (f_j(\Theta; x_i) - f_{y_i}(\Theta; x_i))) \geq \max \exp(M^a (f_j(\Theta; x_i) - f_{y_i}(\Theta; x_i))) = \exp \gamma_\Theta$ in order to lower bound equation A.1 and obtain

$$L_\lambda(M\Theta) \geq \frac{1}{n}\log(1 + \exp(-M^a \gamma_\Theta)) + \lambda M^r \|\Theta\|^r \qquad \text{(A.3)}$$

Applying equation A.2 with $M = \|\Theta_\lambda\|$ and $\Theta = \Theta^\star$, noting that $\|\Theta^\star\| \leq 1$, we have:

$$L_\lambda(\Theta^\star \|\Theta_\lambda\|) \leq \log(1 + (l-1)\exp(-\|\Theta_\lambda\|^a \gamma^\star)) + \lambda \|\Theta_\lambda\|^r \qquad \text{(A.4)}$$

Next we lower bound $L_\lambda(\Theta_\lambda)$ by applying equation A.3,

$$L_\lambda(\Theta_\lambda) \geq \frac{1}{n}\log(1 + \exp(-\|\Theta_\lambda\|^a \gamma_\lambda)) + \lambda \|\Theta_\lambda\|^r \qquad \text{(A.5)}$$

Combining equation A.4 and equation A.5 with the fact that $L_\lambda(\Theta_\lambda) \leq L_\lambda(\Theta^\star \|\Theta_\lambda\|)$ (by the global optimality of Θ_λ), we have

$$\forall \lambda > 0, n\log(1 + (l-1)\exp(-\|\Theta_\lambda\|^a \gamma^\star)) \geq \log(1 + \exp(-\|\Theta_\lambda\|^a \gamma_\lambda))$$

Recall that by Lemma A.2, as $\lambda \to 0$, we have $\|\Theta_\lambda\| \to \infty$. Therefore, $\exp(-\|\Theta_\lambda\|^a \gamma^\star), \exp(-\|\Theta_\lambda\|^a \gamma_\lambda) \to 0$. Thus, we can apply Taylor expansion to the equation above with respect to $\exp(-\|\Theta_\lambda\|^a \gamma^\star)$ and $\exp(-\|\Theta_\lambda\|^a \gamma_\lambda)$. If $\max\{\exp(-\|\Theta_\lambda\|^a \gamma^\star), \exp(-\|\Theta_\lambda\|^a \gamma_\lambda)\} < 1$, then we obtain

$$n(l-1)\exp(-\|\Theta_\lambda\|^a \gamma^\star) \geq \exp(-\|\Theta_\lambda\|^a \gamma_\lambda) - O(\max\{\exp(-\|\Theta_\lambda\|^a \gamma^\star)^2, \exp(-\|\Theta_\lambda\|^a \gamma_\lambda)^2\})$$

We claim this implies that $\gamma^\star \leq \liminf_{\lambda \to 0} \gamma_\lambda$. If not, we have $\liminf_{\lambda \to 0} \gamma_\lambda < \gamma^\star$, which implies that the equation above is violated with sufficiently large $\|\Theta_\lambda\|$ ($\|\Theta_\lambda\| \gg \log(2(\ell-1)n)^{1/a}$ would suffice). By Lemma A.2, $\|\Theta_\lambda\| \to \infty$ as $\lambda \to 0$ and therefore we get a contradiction.

Finally, we have $\gamma_\lambda \leq \gamma^\star$ by definition of $\gamma^\star$. Hence, $\lim_{\lambda \to 0} \gamma_\lambda$ exists and equals $\gamma^\star$. □

Now we fill in the proof of Lemma A.2.

Proof of Lemma A.2. For the sake of contradiction, we assume that $\exists C > 0$ such that for any $\lambda_0 > 0$, there exists $0 < \lambda < \lambda_0$ with $\|\Theta_\lambda\| \leq C$. We will determine the choice of λ_0 later and pick λ such that $\|\Theta_\lambda\| \leq C$. Then the logits (the prediction $f_j(\Theta, x_i)$ before softmax) are bounded in absolute value by some constant (that depends on C), and therefore the loss function $-\log \frac{\exp(f_{y_i}(\Theta; x_i))}{\sum_{j=1}^{l} \exp(f_j(\Theta; x_i))}$ for every example is bounded from below by some constant $D > 0$ (depending on C but not λ.)

Let $M = \lambda^{-1/(r+1)}$, we have that

$$\begin{aligned} 0 < D \leq L_\lambda(\Theta_\lambda) &\leq L_\lambda(M\Theta^\star) \qquad \text{(by the optimality of } \Theta_\lambda\text{)} \\ &\leq -\log \frac{1}{1 + (l-1)\exp(-M^a \gamma^\star)} + \lambda M^r \qquad \text{(by equation A.2)} \\ &= \log(1 + (l-1)\exp(-\lambda^{-a/(r+1)} \gamma^\star)) + \lambda^{1/(r+1)} \\ &\leq \log(1 + (l-1)\exp(-\lambda_0^{-a/(r+1)} \gamma^\star)) + \lambda_0^{1/(r+1)} \end{aligned}$$

Taking a sufficiently small λ_0, we obtain a contradiction and complete the proof. □

A.2 FULL BINARY CLASSIFICATION SETTING

For completeness, we state and prove our max-margin results for the setting where we fit binary labels $y_i \in \{-1, +1\}$ (as opposed to indices in $[l]$) and redefining $f(\Theta; \cdot)$ to assign a single real-valued score (as opposed to a score for each label). This lets us work with the simpler λ-regularized logistic loss:

$$L_\lambda(\Theta) \triangleq \frac{1}{n}\sum_{i=1}^{n} \log(1 + \exp(-y_i f(\Theta; x_i))) + \lambda\|\Theta\|^r$$

As before, let $\Theta_\lambda \in \arg\min L_\lambda(\Theta)$, and define the normalized margin γ_λ by $\gamma_\lambda \triangleq \min_i y_i f(\bar{\Theta}_\lambda; x_i)$. Define the maximum possible normalized margin

$$\gamma^\star \triangleq \max_{\|\Theta\|\le 1} \min_i y_i f(\Theta; x_i) \tag{A.6}$$

Theorem A.3. *Assume $\gamma^\star > 0$ in the binary classification setting with logistic loss. Then as $\lambda \to 0$, $\gamma_\lambda \to \gamma^\star$.*

The proof follows via simple reduction to the multi-class case.

Proof of Theorem A.3. We prove this theorem via reduction to the multi-class case with $l = 2$. Construct $\tilde{f} : \mathbb{R}^d \to \mathbb{R}^2$ with $\tilde{f}_1(\Theta; x_i) = -\frac{1}{2} f(\Theta; x_i)$ and $\tilde{f}_2(\Theta; x_i) = \frac{1}{2} f(\Theta; x_i)$. Define new labels $\tilde{y}_i = 1$ if $y_i = -1$ and $\tilde{y}_i = 2$ if $y_i = 1$. Now note that $\tilde{f}_{\tilde{y}_i}(\Theta; x_i) - \tilde{f}_{j\neq\tilde{y}_i}(\Theta; x_i) = y_i f(\Theta; x_i)$, so the multi-class margin for Θ under $\tilde{f}$ is the same as binary margin for Θ under f. Furthermore, defining

$$\tilde{L}_\lambda(\Theta) \triangleq \frac{1}{n}\sum_{i=1}^{n} -\log \frac{\exp(\tilde{f}_{\tilde{y}_i}(\Theta; x_i))}{\sum_{j=1}^{2}\exp(\tilde{f}_j(\Theta; x_i))} + \lambda\|\Theta\|^r$$

we get that $\tilde{L}_\lambda(\Theta) = L_\lambda(\Theta)$, and in particular, $\tilde{L}_\lambda$ and L_λ have the same set of minimizers. Therefore we can apply Theorem 2.1 for the multi-class setting and conclude $\gamma_\lambda \to \gamma^\star$ in the binary classification setting. □

A.3 MISSING PROOF FOR OPTIMIZATION ACCURACY

Proof of Theorem 2.2. Choose $B \triangleq \left(\frac{1}{\gamma^\star}\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}\right)^{1/a}$. We can upper bound $L_\lambda(\Theta')$ by computing

$$\begin{aligned}
L_\lambda(\Theta') \le 2L\lambda(\Theta_\lambda) &\le 2L_\lambda(B\Theta^\star) \\
&\le 2\log(1 + (l-1)\exp(-B^a\gamma^\star)) + 2\lambda B^r && \text{(by equation A.2)} \\
&\le 2(l-1)\exp(-B^a\gamma^\star) + 2\lambda B^r && \text{(using } \log(1+x) \le x\text{)} \\
&\le 2\frac{\lambda}{(\gamma^\star)^{r/a}} + 2\lambda\left(\frac{1}{\gamma^\star}\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}\right)^{r/a} \\
&\le 4\lambda\left(\frac{1}{\gamma^\star}\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}\right)^{r/a} \triangleq L^{(UB)}
\end{aligned}$$

Furthermore, it holds that $\|\Theta'\|^r \le \frac{L^{(UB)}}{\lambda}$. Now we note that

$$L_\lambda(\Theta') \le L^{(UB)} \le 4\frac{[(c+1)\log(n(l-1))]^{r/a}}{n^c(l-1)^c} \le \frac{1}{2n}$$

for sufficiently large c. Now using the fact that $\log(x) \ge \frac{x}{1+x}\ \forall x \ge -1$, we additionally have the lower bound $L_\lambda(\Theta') \ge \frac{1}{n}\log(1 + \exp(-\gamma'\|\Theta'\|^a)) \ge \frac{1}{n}\frac{\exp(-\gamma'\|\Theta'\|^a)}{1+\exp(-\gamma'\|\Theta'\|^a)}$. Since $L^{(UB)} \le 1$, we can rearrange to get

$$\gamma' \ge \frac{-\log\frac{nL_\lambda(\Theta')}{1-nL_\lambda(\Theta')}}{\|\Theta'\|^a} \ge \frac{-\log\frac{nL^{(UB)}}{1-nL^{(UB)}}}{\|\Theta'\|^a} \ge \frac{-\log(2nL^{(UB)})}{\|\Theta'\|^a}$$

The middle inequality followed because $\frac{x}{1-x}$ is increasing in x for $0 \leq x < 1$, and the last because $L^{(UB)} \leq \frac{1}{2n}$. Since $-\log 2nL^{(UB)} > 0$ we can also apply the bound $\|\Theta'\|^r \leq \frac{L^{(UB)}}{\lambda}$ to get

$$\begin{aligned}
\gamma' &\geq \frac{-\lambda^{a/r}\log 2nL^{(UB)}}{(L^{(UB)})^{a/r}} \\
&= \frac{-\log\left(8n\lambda\left(\frac{1}{\gamma^\star}\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}\right)^{r/a}\right)}{\frac{4^{a/r}}{\gamma^\star}\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}} \qquad \text{(by definition of } L^{(UB)}\text{)} \\
&\geq \frac{\gamma^\star}{4^{a/r}}\left(\frac{\log(\frac{(\gamma^\star)^{r/a}}{8n\lambda})}{\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}} - \frac{\frac{r}{a}\log\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}}{\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}}\right)
\end{aligned}$$

By assumption, $\log\frac{(\gamma^\star)^{r/a}}{\lambda} \geq c\log n(l-1)$, and so $\log(\frac{(\gamma^\star)^{r/a}}{8n\lambda}) \geq \frac{c-2}{c+1}\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}$. Thus, if $c \geq 4$,

$$\frac{-\lambda^{a/r}\log L^{(UB)}}{(L^{(UB)})^{a/r}} \geq \frac{\gamma^\star}{4^{a/r}}\left(\frac{3}{5} - \frac{\frac{r}{a}\log\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}}{\log\frac{(l-1)(\gamma^\star)^{r/a}}{\lambda}}\right) \geq \frac{\gamma^\star}{2\cdot 4^{a/r}}$$

for c sufficiently large. □

B Missing Proofs in Sections 3.2 and 3.3

We first show Theorem 3.3 and then complete the proof of Theorem 3.2. The proof of Theorem 3.3 will consist of two steps: first, show that equation 3.6 has an optimal solution with sparsity n, and second, show that sparse solutions to equation 3.6 can be mapped to a neural network with the same margin, and vice versa. The following lemma and proof are based on Lemma 14 of Tibshirani et al. (2013).

Lemma B.1. *Let supp*$(\alpha) \triangleq \{\bar{u} : |\alpha(\bar{u})| > 0\}$. *There exists an optimal solution* $\alpha^\star$ *to equation 3.6 with* $|$*supp*$(\alpha^\star)| \leq n$.

For the proof of this lemma, we find it convenient to work with a minimum norm formulation which we show is equivalent to equation 3.6:

$$\begin{aligned}
&\min_{\alpha} \|\alpha\|_1 \\
&\text{subject to } y_i\langle\alpha, \varphi(x_i)\rangle \geq 1 \ \forall i
\end{aligned} \tag{B.1}$$

Claim B.2. *Let* $S \subset \mathcal{L}^1(\mathbb{S}^{d-1})$ *be the set of optimizers for equation* 3.6*, and let* $S' \subset \mathcal{L}^1(\mathbb{S}^{d-1})$ *be the set of optimizers for equation B.1. If equation B.1 is feasible, for any* $\alpha \in S$, $\frac{\alpha}{\gamma_{\ell_1}} \in S'$*, and for any* $\alpha' \in S'$, $\frac{\alpha'}{\|\alpha'\|_1} \in S$.

Proof. Let opt' denote the optimal objective for equation B.1. We note that $\frac{\alpha'}{\|\alpha'\|_1}$ is feasible for equation 3.6 with objective $\frac{1}{\mathrm{opt}'}$, and therefore $\gamma_{\ell_1} \geq \frac{1}{\mathrm{opt}'}$. Furthermore, $\frac{1}{2\gamma_{\ell_1}} y_i \int_{\bar{u}\in\mathbb{S}^{d-1}} \alpha(\bar{u})\phi(\bar{u}^\top x_i)d\bar{u} \geq 1 \ \forall i$, and so $\frac{\alpha}{\gamma_{\ell_1}}$ is feasible for equation B.1 with objective $\frac{1}{\gamma_{\ell_1}}$. Therefore, $\mathrm{opt}' \leq \frac{1}{\gamma_{\ell_1}}$. As a result, it must hold that $\mathrm{opt}' = \frac{1}{\gamma_{\ell_1}}$, which means that $\frac{\alpha'}{\|\alpha'\|_1}$ is optimal for equation 3.6, and $\frac{\alpha}{\gamma_{\ell_1}}$ is optimal for equation B.1, as desired. □

First, note that if equation B.1 is not feasible, then $\gamma_{\ell_1} = 0$ and equation 3.6 has a trivial sparse solution, the all zeros function. Thus, it suffices to show that an optimal solution to equation B.1 exists that is n-sparse, since by Lemma B.2 equation B.1 and equation 3.6 have equivalent solutions up to a scaling. We begin by taking the dual of equation B.1.

Claim B.3. *The dual of equation B.1 has form*

$$\max_{\lambda \in \mathbb{R}^n} \lambda^\top \vec{1}$$
$$\textit{subject to} \left| \sum_{i=1}^n \lambda_i y_i \phi(\bar{u}^\top x_i) \right| \leq 1 \; \forall \bar{u} \in \mathbb{S}^{d-1}$$
$$\lambda_i \geq 0$$

For any primal optimal solution $\alpha^\star$ and dual optimal solution $\lambda^\star$, it must hold that

$$\sum_{i=1}^n \lambda_i^\star y_i \phi(\bar{u}^\top x_i) = \text{sign}(\alpha^\star(\bar{u})) \iff \alpha^\star(\bar{u}) \neq 0 \tag{B.2}$$

Proof. The dual form can be solved for by computation. By strong duality, equation B.2 must follow from the KKT conditions. □

Now define the mapping $v : \mathbb{S}^{d-1} \to \mathbb{R}^n$ with $v_i(\bar{u}) \triangleq y_i \phi(\bar{u}^\top x_i)$. We will show a general result about linearly dependent $v(\bar{u})$ for $\bar{u} \in \text{supp}(\alpha^\star)$, after which we can reduce directly to the proof of Tibshirani et al. (2013).

Claim B.4. *Let $\alpha^\star$ be any optimal solution. Suppose that there exists $S \subseteq supp(\alpha^\star)$ such that $\{v(\bar{u}) : \bar{u} \in S\}$ forms a linearly dependent set, i.e.*

$$\sum_{\bar{u} \in S} c_{\bar{u}} v(\bar{u}) = \vec{0} \tag{B.3}$$

for coefficients c. Then $\sum_{\bar{u} \in S} c_{\bar{u}} \text{sign}(\alpha^\star(\bar{u})) = 0$.

Proof. Let $\lambda^\star$ be any dual optimal solution, then ${\lambda^\star}^\top v(\bar{u}) = \text{sign}(\alpha^\star(\bar{u})) \; \forall \bar{u} \in \text{supp}(\alpha^\star)$ by Claim B.3. Thus, we apply ${\lambda^\star}^\top$ to both sides of equation B.3 to get the desired statement. □

Proof of Lemma B.1. The rest of the proof follows Lemma 14 in Tibshirani et al. (2013). The lemma argues that if the conclusion of Claim B.4 holds and an optimal solution $\alpha^\star$ has $S \subseteq \text{supp}(\alpha^\star)$ with $\{v(\bar{u}) : \bar{u} \in S\}$ linearly dependent, we can construct a new α' with $\|\alpha'\|_1 = \|\alpha^\star\|_1$ and $\text{supp}(\alpha') \subset \text{supp}(\alpha^\star)$ (where the inclusion is strict). Thus, if we consider an optimal $\alpha^\star$ with minimal support, it must follow that $\{v(\bar{u}) : \bar{u} \in \text{supp}(\alpha^\star)\}$ is a linearly independent set, and therefore $|\text{supp}(\alpha^\star)| \leq n$. □

We can now complete the proof of Theorem 3.3.

Proof of Theorem 3.3. We first apply Lemma B.1 to conclude that equation 3.6 admits a n-sparse optimal solution $\alpha^\star$. Because of sparsity, we can now abuse notation and treat $\alpha^\star$ as a real-valued function such that $\sum_{\bar{u} \in \text{supp}(\alpha^\star)} |\alpha^\star(\bar{u})| \leq 1$. We construct Θ with normalized margin at least $\frac{\gamma_{\ell_1}}{2}$. For every $\bar{u} \in \text{supp}(\alpha)$, we let Θ have a corresponding hidden unit j with $(w_j, u_j) = \left(\text{sign}(\alpha^\star(\bar{u})) \sqrt{\frac{|\alpha^\star(\bar{u})|}{2}}, \sqrt{\frac{|\alpha^\star(\bar{u})|}{2}} \bar{u} \right)$, and set the remaining hidden units to $\vec{0}$. This is possible because $m \geq n$. Now

$$f(\Theta; x) = \sum_{j=1}^m w_j \phi(u_j^\top x) = \frac{1}{2} \sum_{\bar{u} \in \text{supp}(\alpha^\star)} \alpha^\star(\bar{u}) \phi(\bar{u}^\top x)$$

Furthermore,

$$\|\Theta\|_2^2 = \sum_{j=1}^m w_j^2 + \|u_j\|_2^2 = \sum_{\bar{u} \in \text{supp}(\alpha)} \frac{|\alpha^\star(\bar{u})|}{2} + \frac{|\alpha^\star(\bar{u})|}{2} \|\bar{u}\|_2^2 = \sum_{\bar{u} \in \text{supp}(\alpha)} |\alpha^\star(\bar{u})| \leq 1$$

Thus it follows that Θ has normalized margin at least $\gamma_{\ell_1}/2$, so $\gamma^{\star,m} \geq \gamma_{\ell_1}/2$.

To conclude, we show that $\gamma^{\star,m} \leq \gamma_{\ell_1}/2$. Let $\Theta^{\star,m}$ denote the parameters obtaining optimal m-unit margin $\gamma^{\star,m}$ with hidden units $(w_j^{\star,m}, u_j^{\star,m})$ for $j \in [m]$. We can construct α to put a scaled delta mass of $2w_j^{\star,m}\|u_j^{\star,m}\|_2$ on $\bar{u}_j^{\star,m}$ for $j \in [m]$. It follows that

$$\|\alpha\|_1 = \sum_{j=1}^{m} 2|w_j^{\star,m}|\|u_j^{\star,m}\|_2 \leq \sum_{j=1}^{m} w_j^{\star,m2} + \|u_j^{\star,m}\|_2^2 = \|\Theta^{\star,m}\|_2^2 \leq 1$$

Furthermore,

$$\int_{\mathbb{S}^{d-1}} \alpha(\bar{u})\phi(\bar{u}^\top x) = 2\sum_{j=1}^{m} w_j^{\star,m}\|u_j^{\star,m}\|_2 \phi((\bar{u}_j^{\star,m})^\top x) = 2\sum_{j=1}^{m} w_j^{\star,m}\phi(u_j^{\star,m\top} x) = 2f(\Theta^{\star,m}; x)$$

Thus, α is a feasible solution to equation 3.6 with objective value at least $2\gamma^{\star,m}$. Therefore, $\gamma_{\ell_1} \geq 2\gamma^{\star,m}$, so $\gamma^{\star,m} = \gamma_{\ell_1}/2$. □

Theorem 3.1 follows almost immediately.

Proof of Theorem 3.1. Let $\Theta^{\star,m}$ denote the parameters obtaining optimal m-unit margin $\gamma^{\star,m}$. It is simple to see that $\gamma^{\star,m+1} \geq \gamma^{\star,m}$: we can set m units of our $m+1$-hidden unit network to $\Theta^{\star,m}$, and set the remaining unit to $\vec{0}$. Then this network will compute $f(\Theta^{\star,m}; \cdot)$ using $m+1$ units and have the same parameter norm.

Finally, by Theorem 3.3 we have $\gamma_{\ell_1}/2 = \gamma^{\star,n} = \gamma^{\star,n+1} = \cdots$ □

C Rademacher Complexity and Generalization Error

We prove the generalization error bounds stated in Theorem 3.1 and Lemma 3.4 via Rademacher complexity. In this section, we will introduce the margin-based generalization error bounds for Rademacher complexity and prove Rademacher complexity bounds for the ℓ_1 and ℓ_2 norms.

Assume that our data X, Y are drawn i.i.d. from ground truth distribution p_{data} supported on $\mathcal{X} \times \mathcal{Y}$. For some hypothesis class $\mathcal{F}$ of real-valued functions, we define the empirical Rademacher complexity $\hat{\mathfrak{R}}(\mathcal{F})$ as follows:

$$\hat{\mathfrak{R}}(\mathcal{F}) \triangleq \frac{1}{n}\mathbb{E}_{\epsilon_i}\left[\sup_{f\in\mathcal{F}} \sum_{i=1}^{n} \epsilon_i f(x_i)\right]$$

where ϵ_i are independent Rademacher random variables. Specifically, we consider the hypothesis classes of linear functionals in lifted feature space with bounded norm: for activation ϕ, define $f^\phi(\alpha; \cdot) : \mathbb{R}^d \to \mathbb{R}$ with $f^\phi(\alpha; x) \triangleq \langle \varphi(x), \alpha \rangle$, and let $\mathcal{F}_B^{p,\phi} \triangleq \{f^\phi(\alpha; \cdot) : \alpha \in \mathcal{L}^p, \|\alpha\|_p \leq B\}$. Recall that we defined $\varphi(x) \in \mathcal{L}^\infty(\mathbb{S}^{d-1})$ in equation 3.5.

For the binary classification setting, we will use the following ramp loss, which can be thought of a as a Lipschitz version of the standard 0-1 loss:

$$l_\gamma(a) \triangleq \begin{cases} 1 & a \leq 0) \\ 1 - a/\gamma & 0 < a \leq \gamma \\ 0 & a > \gamma \end{cases}$$

Following Section 3.1, we will use $L(\alpha) \triangleq \Pr_{(x,y)\sim p_{\text{data}}}(yf^\phi(\alpha; x) \leq 0)$ to denote the population 0-1 loss of the classifier $f^\phi(\alpha; \cdot)$. The following theorem follows from classical results in the literature (Koltchinskii et al., 2002), (Kakade et al., 2009).

Theorem C.1. *Let $(x_i, y_i)_{i=1}^n$ be drawn iid from p_{data}. We work in the binary classification setting, so $\mathcal{Y} = \{-1, 1\}$. Let $q \triangleq \frac{p}{p-1}$, $C \triangleq \sup_{x\in\mathcal{X}} \|\varphi(x)\|_q$. Then with probability at least $1-\delta$ over the random draws of the data, every $f^\phi(\alpha; \cdot) \in \mathcal{F}_1^{p,\phi}$ satisfies*

$$L(\alpha) \leq \frac{1}{n}\sum_{i=1}^{n} l_\gamma(y_i f^\phi(\alpha; x_i)) + \frac{4\hat{\mathfrak{R}}(\mathcal{F}_1^{p,\phi})}{\gamma} + \sqrt{\frac{\log\log_2 \frac{4C}{\gamma}}{n}} + \sqrt{\frac{\log(1/\delta)}{2n}} \tag{C.1}$$

In particular, for any α that classifies all training data correctly with p-norm normalized margin $\gamma_\alpha \triangleq \min_i y_i f^\phi(\frac{\alpha}{\|\alpha\|_p}; x_i) > 0$, with probability at least $1-\delta$,

$$L(\alpha) \leq \frac{4\hat{\mathfrak{R}}(\mathcal{F}_1^{p,\phi})}{\gamma_\alpha} + \sqrt{\frac{\log\log_2 \frac{4C}{\gamma_\alpha}}{n}} + \sqrt{\frac{\log(1/\delta)}{2n}} \tag{C.2}$$

Proof. Our starting point is Theorem 2 of (Kakade et al., 2009), which states the same claim for $C \triangleq \sup_{f^\phi(\alpha;\cdot)\in\mathcal{F}_1^{p,\phi}} \sup_{x\in\mathcal{X}} f^\phi(\alpha; x)$. We note that

$$\begin{aligned}
\sup_{f^\phi(\alpha;\cdot)\in\mathcal{F}_1^{p,\phi}} \sup_{x\in\mathcal{X}} f^\phi(\alpha; x) &= \sup_{f^\phi(\alpha;\cdot)\in\mathcal{F}_1^{p,\phi}} \sup_{x\in\mathcal{X}} \langle \alpha, \varphi(x)\rangle \\
&\leq \sup_{f^\phi(\alpha;\cdot)\in\mathcal{F}_1^{p,\phi}} \sup_{x\in\mathcal{X}} \|\alpha\|_p \|\varphi(x)\|_q \qquad \text{(by the duality of } p \text{ and } q \text{ norms)} \\
&\leq C
\end{aligned}$$

For the second statement, we simply apply equation C.1 to $\bar{\alpha} \triangleq \alpha/\|\alpha\|_p$ and use the fact that $L(\alpha) = L(\bar{\alpha})$. □

Next, we focus on analyzing the Rademacher complexity $\hat{\mathfrak{R}}(\mathcal{F}_B^{p,\phi})$, specifically for the cases $p = 1, 2$. Since our general analysis only involves duality of norms, similar derivations have been done in the past (Bartlett & Mendelson, 2002). We include our derivations here for completeness.

Lemma C.2. $\hat{\mathfrak{R}}(\mathcal{F}_B^{1,\phi}) \leq \frac{1}{n} B \mathbb{E}_{\epsilon_i}\left[\|\sum_{i=1}^n \epsilon_i \varphi(x_i)\|_\infty\right]$.

Lemma C.3. $\hat{\mathfrak{R}}(\mathcal{F}_B^{2,\phi}) \leq \frac{1}{n} B \sqrt{\sum_{i=1}^n \|\varphi(x_i)\|_2^2}$.

Proof of Lemmas C.2 and C.3. For $p = 1, 2$, let $\|\cdot\|_{p,*}$ denote the dual norm of $\|\cdot\|_p$, so $\|\cdot\|_{1,*} = \|\cdot\|_\infty$ and $\|\cdot\|_{2,*} = \|\cdot\|_2$. We write

$$\begin{aligned}
\hat{\mathfrak{R}}(\mathcal{F}_B^{p,\phi}) &= \frac{1}{n}\mathbb{E}_{\epsilon_i}\left[\sup_{\alpha\in\mathcal{F}_B^k} \langle \alpha, \sum_{i=1}^n \epsilon_i \varphi(x_i)\rangle\right] \\
&\leq \frac{1}{n}\mathbb{E}_{\epsilon_i}\left[\sup_{\alpha\in\mathcal{F}_B^k} \|\alpha\|_k \left\|\sum_{i=1}^n \epsilon_i \varphi(x_i)\right\|_{k,*}\right] \\
&\leq \frac{1}{n} B \cdot \mathbb{E}_{\epsilon_i}\left[\left\|\sum_{i=1}^n \epsilon_i \varphi(x_i)\right\|_{k,*}\right]
\end{aligned}$$

Lemma C.2 immediately follows. To obtain Lemma C.3, we have

$$\begin{aligned}
\hat{\mathfrak{R}}(\mathcal{F}_B^{2,\phi}) &\leq \frac{1}{n} B \mathbb{E}_{\epsilon_i}\left[\left\|\sum_{i=1}^n \epsilon_i \varphi(x_i)\right\|_2\right] \\
&\leq \frac{1}{n} B \sqrt{\mathbb{E}_{\epsilon_i}\left[\left\|\sum_{i=1}^n \epsilon_i \varphi(x_i)\right\|_2^2\right]} \qquad \text{(via Jensen's inequality)} \\
&\leq \frac{1}{n} B \sqrt{\mathbb{E}_{\epsilon_i}\left[\sum_{i=1}^n \sum_{j=1}^n \epsilon_i \epsilon_j \langle \varphi(x_i), \varphi(x_i)\rangle\right]} \\
&\leq \frac{1}{n} B \sqrt{\sum_{i=1}^n \|\varphi(x_i)\|_2^2} \qquad \text{(terms where } i \neq j \text{ cancel out)}
\end{aligned}$$

□

We will now complete the bound on $\hat{\mathfrak{R}}(\mathcal{F}_B^{1,\phi})$ for general Lipschitz activations ϕ.

Lemma C.4. *Suppose that our activation ϕ is M-Lipschitz. Then*

$$\hat{\mathfrak{R}}(\mathcal{F}_B^{1,\phi}) \leq \frac{B(\sqrt{\sum_{i=1}^n \|\varphi(x_i)\|_\infty^2} + 2M\sqrt{\sum_{i=1}^n \|x_i\|_2^2})}{n}$$

Proof. We will show that

$$E_{\epsilon_i}\left[\left\|\sum_{i=1}^n \epsilon_i \varphi(x_i)\right\|_\infty\right] \leq \sqrt{\sum_{i=1}^n \|\varphi(x_i)\|_\infty^2} + 2M\sqrt{\sum_{i=1}^n \|x_i\|_2^2}$$

from which the statement of the lemma follows via Lemma C.2. Fix any $\bar{u}' \in \mathbb{S}^{d-1}$. Then we get the decomposition

$$\begin{aligned} E_{\epsilon_i}\left[\sup_{\bar{u}\in\mathbb{S}^{d-1}}\left|\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i)\right|\right] \leq E_{\epsilon_i}\left[\left|\sum_{i=1}^n \epsilon_i \phi(\bar{u}'^\top x_i)\right|\right] + \\ E_{\epsilon_i}\left[\sup_{\bar{u}\in\mathbb{S}^{d-1}}\left|\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i)\right| - \inf_{\bar{u}\in\mathbb{S}^{d-1}}\left|\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i)\right|\right] \end{aligned} \tag{C.3}$$

We can bound the first term as

$$\begin{aligned} E_{\epsilon_i}\left[\left|\sum_{i=1}^n \epsilon_i \phi(\bar{u}'^\top x_i)\right|\right] &\leq \sqrt{E_{\epsilon_i}\left[\left(\sum_{i=1}^n \epsilon_i \phi(\bar{u}'^\top x_i)\right)^2\right]} \\ &\leq \sqrt{\sum_{i=1}^n \phi(\bar{u}'^\top x_i)^2} \\ &\leq \sqrt{\sum_{i=1}^n \|\varphi(x_i)\|_\infty^2} \end{aligned} \tag{C.4}$$

We can bound the second term of equation C.3 by

$$\begin{aligned} &\mathbb{E}_{\epsilon_i}\left[\sup_{\bar{u}\in\mathbb{S}^{d-1}}\left|\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i)\right| - \inf_{\bar{u}\in\mathbb{S}^{d-1}}\left|\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i)\right|\right] \\ &\leq \mathbb{E}_{\epsilon_i}\left[\sup_{\bar{u}\in\mathbb{S}^{d-1}}\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i) - \inf_{\bar{u}\in\mathbb{S}^{d-1}}\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i)\right] \\ &\leq 2\mathbb{E}_{\epsilon_i}\left[\sup_{\bar{u}\in\mathbb{S}^{d-1}}\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i)\right] \end{aligned} \tag{C.5}$$

This simply gives an empirical Rademacher complexity of the hypothesis class $\mathcal{F} \triangleq \{x \mapsto \phi(\bar{u}^\top x) : \bar{u} \in \mathbb{S}^{d-1}\}$ scaled by n. By the Lipschitz contraction property of Rademacher complexity, using the fact that ϕ is M-Lipschitz, we can therefore bound equation C.5 by

$$2\mathbb{E}_{\epsilon_i}\left[\sup_{\bar{u}\in\mathbb{S}^{d-1}}\sum_{i=1}^n \epsilon_i \phi(\bar{u}^\top x_i)\right] \leq 2M\sqrt{\sum_{i=1}^n \|x_i\|_2^2} \tag{C.6}$$

Plugging equation C.4 and equation C.6 back into equation C.3 gives the desired bound. □

As an example, we can apply this bound to relu features:

Corollary C.5. *Suppose that ϕ is the* relu *activation. Let $\kappa \triangleq Vol(\mathbb{S}^{d-1})$. Then $\hat{\mathfrak{R}}(\mathcal{F}_B^{1,\phi}) \leq \frac{3B\|X\|_F}{n}$, and $\hat{\mathfrak{R}}(\mathcal{F}_B^{2,\phi}) \lesssim \frac{B\sqrt{\kappa}\|X\|_F}{n\sqrt{d}}$.*

Proof. In the setting of Lemma C.4, we have $\|\varphi(x_i)\|_\infty \leq \|x_i\|_2$ and $M = 1$. Applying this in Lemma C.4 gives the desired 1-norm Rademacher bound.

For the 2-norm case, we first show that $\|\varphi(x_i)\|_2^2 = \Theta\left(\frac{\kappa}{d}\|x_i\|_2^2\right)$. We can compute

$$\begin{aligned}
\|\varphi(x_i)\|_2^2 &= \text{Vol}(\mathbb{S}^{d-1})\mathbb{E}_{\bar{u}\sim\mathbb{S}^{d-1}}[\text{relu}(\bar{u}^\top x_i)^2] \\
&= \frac{\kappa}{d}\mathbb{E}_{\bar{u}\sim\mathbb{S}^{d-1}}[\text{relu}(\sqrt{d}\bar{u}^\top x_i)^2] \\
&= \frac{\kappa}{d}\frac{1}{M_2}\mathbb{E}_{u\sim\mathcal{N}(0,I_{d\times d})}[\text{relu}(u^T x_i)^2] && (M_2 \text{ is the second moment of } \mathcal{N}(0,1)) \\
&= \Theta\left(\frac{\kappa}{d}\|x_i\|_2^2\right) && \text{(C.7)}
\end{aligned}$$

where the last line uses the computation provided in Lemma A.1 by Du et al. (2017). Now we plug this into Lemma C.3 to get the desired bound. □

C.1 Proof of Network and Kernel Generalization Bounds

We use our Rademacher complexity-based generalization bounds to provide a proof of Theorem 3.1 and Lemma 3.4.

Proof of Theorem 3.1. Since ϕ is 1-homogeneous and 1-Lipschitz, it must follow that $|\phi(\bar{u}^\top x)| \leq |\bar{u}^\top x|$, and so $\|\varphi(x)\|_\infty \leq \|x\|_2$. Thus, $\sup_{x\in\mathcal{X}}\|\varphi(x)\|_\infty \leq \sup_{x\in\mathcal{X}}\|x\|_2$. Furthermore, $\hat{\mathfrak{R}}(\mathcal{F}_1^{1,\phi}) \leq \frac{3\|X\|_F}{n}$. Finally, we note that the mapping given in the proof of Theorem 3.3 between neural network and linear functional on $\varphi(\cdot)$ applies here as well: we construct α to put a mass of $2w_j\|u_j\|_2$ on $\bar{u}_j$ for each $(w_j, u_j) \in \Theta$. Then $\|\alpha\|_1 \leq \|\Theta\|_2^2$, and $\langle\alpha, \varphi(x)\rangle = 2f(\Theta; x)$. Thus, following the notation of Theorem C.1, $\frac{\gamma_\alpha}{2} \geq \gamma \triangleq \min_i y_i f(\bar{\Theta}; x_i)$. Now we can apply equation C.2 of Theorem C.1 to get with probability $1-\delta$,

$$\begin{aligned}
L(\Theta) &= L(\alpha) \\
&\leq \frac{4\hat{\mathfrak{R}}(\mathcal{F}_1^{1,\phi})}{\gamma_\alpha} + \sqrt{\frac{\log\log_2\frac{4\sup_{x\in\mathcal{X}}\|x\|_2}{\gamma_\alpha}}{n}} + \sqrt{\frac{\log(1/\delta)}{2n}} \\
&\leq \frac{6\|X\|_F}{n\gamma} + \sqrt{\frac{\log\log_2\frac{4C}{\gamma}}{n}} + \sqrt{\frac{\log(1/\delta)}{2n}} && \text{(plugging in our bound for } \hat{\mathfrak{R}}(\mathcal{F}_1^{1,\phi}))
\end{aligned}$$

which gives us equation 3.2. To conclude equation 3.3, we apply the above on $\Theta_{\lambda,m}$ and use Theorem A.3. □

We will now prove Lemma 3.4.

Proof of Lemma 3.4. From equation C.7, we first obtain $\sup_{x\in\mathcal{X}}\|\varphi(x)\|_2 \lesssim C\sqrt{\frac{\kappa}{d}}$. Denote the optimizer for equation 3.8 by α_{ℓ_2}. Since α_{ℓ_2} has normalized margin $\sqrt{\kappa}\gamma_{\ell_2}$, we apply Theorem C.1 to get with probability $1-\delta$,

$$\begin{aligned}
L_{\ell_2\text{-svm}} = L(\alpha_{\ell_2}) &\leq \frac{4\hat{\mathfrak{R}}(\mathcal{F}_1^{2,\phi})}{\sqrt{\kappa}\gamma_{\ell_2}} + \sqrt{\frac{\log\log_2\frac{4\sup_{x\in\mathcal{X}}\|\varphi(x)\|_2}{\gamma_{\ell_2}}}{n}} + \sqrt{\frac{\log(1/\delta)}{2n}} \\
&\lesssim \frac{\|X\|_F}{n\gamma_{\ell_2}\sqrt{d}} + \sqrt{\frac{\log\max\left\{\log_2\frac{C\sqrt{\kappa/d}}{\gamma_{\ell_2}}, 2\right\}}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} \\
&\qquad\qquad \text{(applying Corollary C.5)}
\end{aligned}$$

□

D MISSING PROOFS FOR COMPARISON TO KERNEL METHODS

D.1 CLASSIFICATION

In this section we will complete a proof of Theorem 3.5. Recall the construction of the distribution $\mathcal{D}$ provided in Section 3.4. We first provide a classifier of this data with small ℓ_1 norm.

Lemma D.1. *In the setting of Theorem 3.5, we have that*

$$\gamma_{\ell_1} \geq \frac{\sqrt{2}}{4}.$$

Proof. Consider the network $f(x) = \frac{1}{4}\Big((x^\top(e_1+e_2)/\sqrt{2})_+ + (x^\top(-e_1-e_2)/\sqrt{2})_+ - (x^\top(-e_1+e_2)/\sqrt{2})_+ - (x^\top(e_1-e_2)/\sqrt{2})_+\Big)$. The attained margin $\gamma = \frac{\sqrt{2}}{4}$, so $\gamma_{\ell_1} \geq \frac{\sqrt{2}}{4}$. □

Now we will upper bound the margin attainable by the ℓ_2 SVM.

Proof of Lemma 3.6. By the definition of γ_{ℓ_2}, we have that for any α with $\sqrt{\kappa}\|\alpha\|_2 \leq 1$, we have

$$\gamma_{\ell_2} \leq \max_{\sqrt{\kappa}\|\alpha\|_2 \leq 1} \frac{1}{n}\sum_{i=1}^{n} \langle \alpha, y_i \varphi(x_i)\rangle$$

Setting $\alpha = \frac{1}{\sqrt{\kappa}}\frac{1}{n}\sum_{i=1}^{n}\varphi(x_i)y_i / \|\frac{1}{n}\sum_{i=1}^{n}\varphi(x_i)y_i\|_2$ completes the proof. (Attentive readers may realize that this is equivalent to setting the dual variable of the convex program 3.8 to all 1's function.) □

Proof of Lemma 3.7. Let $W_i = \varphi(x_i)y_i$. We will bound several quantities regarding W_i's. In the rest of the proof, we will condition on the event E that $\forall i, \|x_i\|_2^2 \lesssim d\log n$. Note that E is a high probability event and conditioned on E, x_i's are still independent. We omit the condition on E in the rest of the proof for simplicity.

We first show that assuming the following three inequalities that the conclusion of the Lemma follows.

1. $\forall i, \|W_i\|_2^2 \lesssim \kappa \log n$.
2. $\sigma^2 \triangleq \mathrm{Var}[\sum_i W_i] \triangleq \sum_{i=1}^{n}\mathbb{E}[\|W_i - \mathbb{E}\, W_i\|_2^2] \lesssim n\kappa\log n$
3. $\|\,\mathbb{E}\,[\sum W_i]\,\|_2 \lesssim \sqrt{\kappa}n/d$.

By bullets 1, 2, and Bernstein inequality, we have that with probability at least $1 - dn^{-10}$ over the randomness of the data (X, Y),

$$\left\|\sum_{i=1}^{n} W_i - \mathbb{E}\left[\sum_{i=1}^{n} W_i\right]\right\|_2 \lesssim \sqrt{\kappa}\log^{1.5} n + \sqrt{n\kappa\log^2 n} \lesssim \sqrt{n\kappa\log^2 n}$$

By bullet 3 and equation above, we complete the proof with triangle inequality:

$$\left\|\sum_{i=1}^{n} W_i\right\|_2 \leq \left\|\mathbb{E}\left[\sum_{i=1}^{n} W_i\right]\right\|_2 + \sqrt{n\kappa\log^2 n} \lesssim \sqrt{n\kappa\log^2 n} + \sqrt{\kappa}n/d$$

Therefore, it suffices to prove bullets 1, 2 and 3. Note that 2 is a direct corollary of 1 so we will only prove 1 and 3. We start with 3:

By the definition of the ℓ_2 norm in $\mathcal{L}^2(\mathbb{S}^{d-1})$ and the independence of (x_i, y_i)'s, we can rewrite

$$\left\|\mathbb{E}\left[\sum_{i=1}^{n} W_i\right]\right\|_2^2 = \kappa \cdot n^2 \mathop{\mathbb{E}}_{\bar{u}\sim\mathbb{S}^{d-1}}\left[\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}} \varphi(x)[\bar{u}]\cdot y\right]^2 \tag{D.1}$$

Let $\bar{u} = (\bar{u}_1, \ldots, \bar{u}_d)$ and $\bar{u}_{-2} = (\bar{u}_3, \ldots, \bar{u}_d) \in \mathbb{R}^{d-2}$, and define $\tau \triangleq \|\bar{u}_{-2}\|_2$. Let $x_{-2} = (x^\top e_3, \ldots, x^\top e_d)$. Note that $\varphi(x)[\bar{u}]y = y[\bar{u}_1 \cdot x^\top e_1 + \bar{u}_2 \cdot x^\top e_2 + \bar{u}_{-2}^\top x_{-2}]$ and $\bar{u}_{-2}^\top x_{-2}$ has distribution $\|\bar{u}_{-2}\|_2 \cdot \mathcal{N}(0,1) = \tau \cdot \mathcal{N}(0,1)$. Let $z = \bar{u}_{-2}^\top x_{-2}/\tau$, and therefore z has standard normal distribution. With this change of the variables, by the definition of the distribution $\mathcal{D}$, we have

$$\begin{aligned}\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}} \varphi(x)[\bar{u}] \cdot y = \frac{1}{4} \mathop{\mathbb{E}}_{z\sim N(0,1)} [(\bar{u}_1 + \bar{u}_2 + \tau z)_+] + \frac{1}{4} \mathop{\mathbb{E}}_{z\sim N(0,1)} [(-\bar{u}_1 - \bar{u}_2 + \tau z)_+] \\ - \frac{1}{4} \mathop{\mathbb{E}}_{z\sim N(0,1)} [(\bar{u}_1 - \bar{u}_2 + \tau z)_+] - \frac{1}{4} \mathop{\mathbb{E}}_{z\sim N(0,1)} [(+\bar{u}_1 - \bar{u}_2 + \tau z)_+]\end{aligned}$$

By claim D.2, and the 1-homogeneity of relu, we can simplify the above equation to

$$\begin{aligned}\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}} \varphi(x)[\bar{u}] \cdot y = \frac{1}{4}\tau \cdot \big(2c_1 + O(\min\{|\bar{u}_1 + \bar{u}_2|/\tau, |\bar{u}_1 + \bar{u}_2|^2/\tau^2\}) \\ - \frac{\tau}{4}\left(2c_1 - O(\min\{|\bar{u}_1 - \bar{u}_2|/\tau, |\bar{u}_1 - \bar{u}_2|^2/\tau^2\})\right)\big) \\ \lesssim \min\{|\bar{u}_1| + |\bar{u}_2|, (|\bar{u}_1| + |\bar{u}_2|)^2/\tau\}\end{aligned}$$

It follows that

$$\begin{aligned}&\mathop{\mathbb{E}}_{\bar{u}\sim\mathbb{S}^{d-1}} \left[\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}} \varphi(x)[\bar{u}] \cdot y\right]^2 \lesssim \mathbb{E}_{\bar{u}}\left[\min\{(|\bar{u}_1| + |\bar{u}_2|)^2, (|\bar{u}_1| + |\bar{u}_2|)^4/\|\bar{u}_{-2}\|_2^2\}\right] \\ &\lesssim \mathbb{E}_{\bar{u}}\left[(|\bar{u}_1| + |\bar{u}_2|)^2 \cdot \mathbf{1}[\|\bar{u}_{-2}\|_2 \le 1/2]\right] + \mathbb{E}_{\bar{u}}\left[(|\bar{u}_1| + |\bar{u}_2|)^4/\|\bar{u}_{-2}\|_2^2 \cdot \mathbf{1}[\|\bar{u}_{-2}\|_2 \ge 1/2]\right] \\ &\lesssim \exp(-\sqrt{d}) + \mathbb{E}_{\bar{u}}\left[(|\bar{u}_1| + |\bar{u}_2|)^4\right] \lesssim 1/d^2 \qquad \text{(D.2)}\end{aligned}$$

Combining equation D.1 and equation D.2 we complete the proof of bullet 3. Next we prove bullet 1. Note that $\varphi(x)[\bar{u}]y$ is bounded by $|\bar{u}_1| + |\bar{u}_2| + \|\bar{u}_{-2}^\top x_{-2}\|_2$. Therefore, conditioned on $\|x_i\|_2 \lesssim d \log n$

$$\begin{aligned}\|W_i\|_2^2 &\le \mathop{\mathbb{E}}_{\bar{u}\sim\mathbb{S}^{d-1}} \left[(|\bar{u}_1| + |\bar{u}_2| + \|\bar{u}_{-2}^\top x_{-2}\|_2)^2\right] \\ &\lesssim \mathop{\mathbb{E}}_{\bar{u}\sim\mathbb{S}^{d-1}} \left[|\bar{u}_1|^2\right] + \mathop{\mathbb{E}}_{\bar{u}\sim\mathbb{S}^{d-1}} \left[|\bar{u}_2|^2\right] + \mathop{\mathbb{E}}_{\bar{u}\sim\mathbb{S}^{d-1}} \left[\|\bar{u}_{-2}^\top x_{-2}\|_2^2\right] \\ &\lesssim 1/d + \|x_{-2}\|_2^2/d \lesssim \log n\end{aligned}$$

Hence we complete the proof. □

Claim D.2. *Let $Z \sim \mathcal{N}(0,1)$ and $a \in \mathbb{R}$. Then, there exists a universal constant c_1 and c_2 such that*

$$|\mathbb{E}\left[(a+Z)_+ + (-a+Z)_+\right] - 2c_1| \le c_2 \min\{|a|, a^2\}.$$

Proof. Without loss of generality we can assume $a \ge 0$. Then,

$$\begin{aligned}&\mathbb{E}\left[(a+Z)_+ + (-a+Z)_+\right] = \mathbb{E}\left[(a+Z)\mathbf{1}[Z \ge -a]\right] + \mathbb{E}\left[(Z-a)\mathbf{1}[Z \ge a]\right] \\ &= \mathbb{E}\left[a \cdot \mathbf{1}[Z \ge -a]\right] + \mathbb{E}\left[Z \cdot \mathbf{1}[Z \ge -a]\right] - \mathbb{E}\left[a \cdot \mathbf{1}[Z \ge a]\right] + \mathbb{E}\left[Z \cdot \mathbf{1}[Z \ge a]\right] \\ &= \mathbb{E}\left[a\mathbf{1}[-a \le Z \le a]\right] + 2\,\mathbb{E}\left[Z \cdot \mathbf{1}[Z \ge a]\right] \qquad \text{(by } \mathbb{E}\left[Z \cdot \mathbf{1}[-a \le Z \le a]\right] = 0\text{)} \\ &= \mathbb{E}\left[a\mathbf{1}[-a \le Z \le a]\right] + 2\,\mathbb{E}\left[Z \cdot \mathbf{1}[Z \ge 0]\right] - 2\,\mathbb{E}\left[Z \cdot \mathbf{1}[a \ge Z \ge 0]\right] \\ &= 2c_1 + O(\min\{a, a^2\})\end{aligned}$$

where the last equality uses the fact that $c_1 \triangleq \mathbb{E}\left[Z \cdot \mathbf{1}[Z \ge 0]\right]$ and $\mathbb{E}\left[a\mathbf{1}[-a \le Z \le a]\right] \le a\mathbb{E}\left[\mathbf{1}[-a \le Z \le a]\right] \lesssim a \min\{1, a\}$. □

Now we will prove Theorem 3.5.

Proof of Theorem 3.5. To circumvent the technical issue of bounded support in Theorem 3.1 and Lemma 3.4, we construct p_{data} to be a slightly modified version of $\mathcal{D}$: perform rejection sampling of $(x, y) \sim \mathcal{D}$ until we obtain a sample with $\|x\|_2^2 \lesssim d \log n$. Since this occurs with very high

probability, the high probability result of Lemma 3.7 still translates to p_{data}. Now apply Lemma 3.6 to conclude that $\gamma_{\ell_2} \lesssim \frac{\log n}{\sqrt{n}} + \frac{1}{d}$. Furthermore, Lemma D.1 allows us to conclude that $\gamma_{\ell_1} \gtrsim 1$.

With high probability over the draws of the samples, we will also get that $\|X\|_F/\sqrt{n} \lesssim \sqrt{d}$. We can therefore apply Theorem 3.1, and conclude that with probability $1-\delta$,

$$L_{\ell_1\text{-svm}} \lesssim \sqrt{\frac{d}{n}} + \sqrt{\frac{\log\log(d\log n)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}}$$

Furthermore, plugging γ_{ℓ_2} into the bound of Lemma 3.4 gives us

$$\min\left\{\frac{1}{\log n}, \frac{d}{\sqrt{n}}\right\} + \sqrt{\frac{\log\log(dn)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}}$$

□

D.2 Regression

We will first define the 1-norm and 2-norm regression problems. The regression equivalent of equation 3.6 for $\alpha \in \mathcal{L}^1(\mathbb{S}^{d-1})$ is as follows:

$$\begin{aligned} \alpha_{\ell_1} \in \arg\min_{\alpha} \|\alpha\|_1 \\ \text{subject to } \langle \alpha, \varphi(x_i)\rangle = y_i \end{aligned} \tag{D.3}$$

Next we define the regression version of equation 3.8:

$$\begin{aligned} \alpha_{\ell_2} \in \arg\min_{\alpha} \|\alpha\|_2 \\ \text{subject to } \langle \alpha, \varphi(x_i)\rangle = y_i \end{aligned} \tag{D.4}$$

where $\alpha \in \mathcal{L}^2(\mathbb{S}^{d-1})$.

We will briefly motivate our study of the regression setting by connecting the minimum 1-norm solution to neural networks. To compare, in the classification setting, optimizing the weakly regularized loss over neural networks is equivalent to solving the ℓ_1 SVM. In the regression setting, solving the weakly regularized squared error loss is equivalent is equivalent to finding the minimum 1-norm solution that fits the datapoints exactly.

Theorem D.3. *Let $f(\Theta;\cdot)$ be some two-layer neural network with $m \geq n$ hidden units parametrized by Θ, as in Section 3.1. Define the λ-regularized squared error loss*

$$L_{\lambda,m}(\Theta) \triangleq \frac{1}{n}\sum_{i=1}^{n}(f(\Theta;x_i)-y_i)^2 + \lambda\|\Theta\|_2^2$$

with $\Theta_{\lambda,m} \in \arg\min_{\Theta} L_{\lambda,m}(\Theta)$. Suppose that equation D.3 is feasible with optimal solution α_{ℓ_1}. Then as $\lambda \to 0$, $L_{\lambda,m}(\Theta_{\lambda,m}) \to 0$ and $\|\Theta_{\lambda,m}\|_2^2 \to 2\|\alpha_{\ell_1}\|_1$.

Proof. We can see that equation D.3 will have a n-sparse solution $\alpha^\star$ using the same reasoning as the proof of Lemma B.1. Furthermore, following the proof of Theorem 3.3, the function $x \mapsto \langle \alpha^\star, \varphi(x)\rangle$ is implementable by a neural network $\Theta^{\star,m}$ with $\|\Theta^{\star,m}\|_2^2 = 2\|\alpha^\star\|_1 = 2\|\alpha_{\ell_1}\|_1$. Following the same reasoning as before, we can also conclude that $\Theta^{\star,m}$ is an optimal solution for:

$$\begin{aligned} \min_{\Theta} \|\Theta\|_2^2 \\ \text{subject to } f(\Theta;x_i) = y_i \end{aligned} \tag{D.5}$$

Now we note that $\lambda\|\Theta_{\lambda,m}\|_2^2 \leq L_{\lambda,m}(\Theta_{\lambda,m}) \leq L_{\lambda,m}(\Theta^{\star,m}) = \lambda\|\Theta^{\star,m}\|_2^2$, so as $\lambda \to 0$, and also $\|\Theta_{\lambda,m}\|_2 \leq \|\Theta^{\star,m}\|_2$. Now assume for the sake of contradiction that $\exists B$ with $\|\Theta_{\lambda,m}\|_2 \leq B < \|\Theta^{\star,m}\|_2$ for arbitrarily small λ. We define

$$\begin{aligned} r^\star \triangleq \min_{\Theta} \frac{1}{n}\sum_{i=1}^{n}(f(\Theta;x_i)-y_i)^2 \\ \text{subject to } \|\Theta\|_2 \leq B \end{aligned}$$

Note that $r^\star > 0$ since $\Theta^{\star,m}$ is optimal for equation D.5. However, $L_{\lambda,m} \geq r^\star$ for arbitrarily small λ, a contradiction. Thus, $\lim_{\lambda\to 0}\|\Theta^{\star,m}\|_2^2 = \|\Theta^{\star,m}\|_2^2$. □

We proceed to provide similar generalization bounds as the classification setting. This time, our bounds depend on the norms of the solution. We borrow notation from Section C: again let $f^{\phi}(\alpha;\cdot) \triangleq x \mapsto \langle \alpha, \varphi(x)\rangle$ (for $\varphi(x)$ defined in equation 3.5), and define hypothesis classes $\mathcal{F}_B^{p,\phi}$ of linear functionals bounded by B in p-norm. As before, let $\hat{\mathfrak{R}}(\mathcal{F})$ denote the empirical Rademacher complexity of hypothesis class $\mathcal{F}$. The following is a generalization bound based on the 1-norm:

Lemma D.4. *Let $l(\cdot;y) : \mathbb{R} \to [-c,c]$ be a bounded M-Lipschitz loss function. Assume that ϕ is a 1-homogeneous and 1-Lipschitz activation. Let $(x_i,y_i)_{i=1}^n$ be drawn i.i.d from p_{data}. Then with probability at least $1-\delta$ over the dataset, every $\alpha \in \mathcal{L}^1(\mathbb{S}^{d-1})$ satisfies*

$$\mathbb{E}_{(x,y)\sim p_{\mathrm{data}}}[l(f^{\phi}(\alpha;x);y)] \leq \frac{1}{n}\sum_{i=1}^n l(f^{\phi}(\alpha;x_i);y_i) + 12M\frac{\max\{1,\|\alpha\|_1\|X\|_F\}}{n} + c\sqrt{\frac{\log(1/\delta)+\log(\max\{1,2\|\alpha\|_1\|X\|_F\})}{2n}}$$

Proof. Our starting point is Theorem 1 of Kakade et al. (2009), which states that with probability $1-\delta$, for any fixed hypothesis class $\mathcal{F}$ and $f \in \mathcal{F}$,

$$\mathbb{E}_{(x,y)\sim p_{\mathrm{data}}}[l(f(x);y)] \leq \frac{1}{n}\sum_{i=1}^n l(f(x_i);y_i) + 2M\hat{\mathfrak{R}}(\mathcal{F}) + c\sqrt{\frac{\log(1/\delta)}{n}} \tag{D.6}$$

We define $B_j \triangleq \frac{2^j}{\|X\|_F}$ for $j \geq 0$. We note that by Lemma C.4, $\hat{\mathfrak{R}}(\mathcal{F}_{B_j}^{1,\phi}) \leq 3\frac{2^j}{n}$. and apply the above on $\mathcal{F}_{B_j}^{1,\phi}$ using $\delta_j \triangleq \frac{\delta}{2^{j+1}}$. Then using a union bound, with probability $1-\sum_{j=0}^{\infty}\delta_j = 1-\delta$, for all $j \geq 0$ and $f^{\phi}(\alpha;\cdot) \in \mathcal{F}_{B_j}^{1,\phi}$

$$\begin{aligned}\mathbb{E}_{(x,y)\sim p_{\mathrm{data}}}[l(f^{\phi}(\alpha;x);y)] &\leq \frac{1}{n}\sum_{i=1}^n l(f^{\phi}(\alpha;x_i);y_i) + 2M\hat{\mathfrak{R}}(\mathcal{F}_{B_j}^{1,\phi}) + c\sqrt{\frac{\log(1/\delta_j)}{n}} \\ &\leq \frac{1}{n}\sum_{i=1}^n l(f^{\phi}(\alpha;x_i);y_i) + 6M\frac{2^j}{n} + c\sqrt{\frac{\log(1/\delta)+\log(2^{j+1})}{n}}\end{aligned}$$

Now for every α with $\|\alpha\|_1 < \frac{1}{\|X\|_F}$, we use the inequality for $\mathcal{F}_{B_0}^{1,\phi}$, and for every other α, we apply the inequality corresponding to $\mathcal{F}_{B_{j+1}}^{1,\phi}$, where $2^j \leq \|\alpha\|_1\|X\|_F \leq 2^{j+1}$. This gives the desired statement. □

We can also provide the same generalization error bound for the 2-norm and relu features:

Lemma D.5. *In the setting of Lemma D.4, choose ϕ to be the* relu *activation. Then with probability $1-\delta$,*

$$\mathbb{E}_{(x,y)\sim p_{\mathrm{data}}}[l(f^{\phi}(\alpha;x);y)] \lesssim \frac{1}{n}\sum_{i=1}^n l(f^{\phi}(\alpha;x_i);y_i) + M\frac{\sqrt{\kappa}\max\{1,\|\alpha\|_2\|X\|_F\}}{n\sqrt{d}} + c\sqrt{\frac{\log(1/\delta)+\log(\max\{1,\|\alpha\|_2\|X\|_F\})}{2n}}$$

Proof. We proceed the same way as in the proof of Lemma D.4. We define B_j as before, and this time have $\hat{\mathfrak{R}}(\mathcal{F}_{B_j}^{2,\phi}) \lesssim \frac{\sqrt{\kappa}2^j}{n\sqrt{d}}$ from Lemma C.5. Thus, again union bounding over all j, equation equation D.6 gives with probability $1-\delta$, for all $j \geq 0$ and $f^{\phi}(\alpha;\cdot) \in \mathcal{F}_{B_j}^{2,\phi}$

$$\mathbb{E}_{(x,y)\sim p_{\mathrm{data}}}[l(f^{\phi}(\alpha;x);y)] \lesssim \frac{1}{n}\sum_{i=1}^n l(f^{\phi}(\alpha;x_i);y_i) + M\frac{\sqrt{\kappa}2^j}{n\sqrt{d}} + c\sqrt{\frac{\log(1/\delta)+\log(2^{j+1})}{n}}$$

Now we assign the α to different j as before to obtain the statement in the lemma. □

Note that if l is some bounded loss such that $l(y; y) = 0$ (for example, truncated squared error), for α_{ℓ_1} and α_{ℓ_2} the loss terms over the datapoints (in the bounds of Lemmas D.4 and D.5) vanish. For loss l, define

$$L_{\ell_1\text{-reg}} \triangleq \mathbb{E}_{x,y\sim p_{\text{data}}}[l(f^{\phi}(\alpha_{\ell_1}; x); y)]$$
$$L_{\ell_2\text{-reg}} \triangleq \mathbb{E}_{x,y\sim p_{\text{data}}}[l(f^{\phi}(\alpha_{\ell_2}; x); y)]$$

Next, we will define the kernel matrix K with $K_{ij} = \langle \varphi(x_i), \varphi(x_j)\rangle$. Now we are ready to state and prove the formal theorem describing the gap between the 1-norm solution and 2-norm solution.

Theorem D.6. *Recall the definitions of α_{ℓ_1} and α_{ℓ_2} in equation D.3 and equation D.4. For any activation ϕ with the property that K is full rank for any X with no repeated datapoints, there exists a distribution p_{data} such that with probability 1,*

$$\|\alpha_{\ell_1}\|_1 \leq 1$$

On the other hand,

$$\mathbb{E}_{(x_i,y_i)_{i=1}^n \sim_{iid} p_{\text{data}}}[\|\alpha_{\ell_2}\|_2^2] = \frac{n}{\kappa}$$

For i.i.d samples from this choice of p_{data}, if l is bounded ($l(\cdot; y) : \mathbb{R} \to [-1, 1]$), 0 on correct predictions ($l(y; y) = 0$), and 1-Lipschitz, then with probability $1 - \delta$,

$$L_{\ell_1\text{-reg}} \lesssim \sqrt{\frac{d}{n}} + \sqrt{\frac{\log(1/\delta) + \log n}{n}}$$

Meanwhile, in the case that $\|\alpha_{\ell_2}\|_2^2 \geq \frac{n}{\kappa}$, the upper bound on $L_{\ell_2\text{-reg}}$ from Lemma D.5 is $\Omega(1)$ and in particular does not decrease with n.

We will first show that for any dataset X, there is a distribution over Y such that the expectation of $\|\alpha_{\ell_2}\|_2$ is large. When it is clear from context, y will denote the vector corresponding to Y.

Lemma D.7. *There is a distribution $\mathcal{A}$ over $\mathcal{L}^1(\mathbb{S}^{d-1})$ such that for any dataset X with $y_i \triangleq \langle \varphi(x_i), \beta\rangle$ for $\beta \sim \mathcal{A}$,*

$$\mathbb{E}_{\beta}[\|\alpha_{\ell_2}\|_2^2] \geq \frac{n}{\kappa}$$

and with probability 1,

$$\|\alpha_{\ell_1}\|_1 \leq 1$$

We note the order of the quantifiers in Lemma D.7: the distribution $\mathcal{A}$ must not depend on the dataset X. We first provide a simple closed-form expression for $\|\alpha_{\ell_2}\|_2^2$.

Claim D.8. *If K is full rank, then $\|\alpha_{\ell_2}\|_2^2 = y^\top K^{-1} y$.*

Proof. This follows by taking the dual of equation D.4. □

Proof of Lemma D.7. We sample $\beta \sim \mathcal{A}$ as follows: first sample $\bar{u} \sim \mathbb{S}^{d-1}$ uniformly. Then set β to have a delta mass of 1 at $\bar{u}$ and be 0 everywhere else. Define the vector $v_{\bar{u}} \triangleq [\phi(\bar{u}^\top x_1) \cdots \phi(\bar{u}^\top x_n)]$; then it follows that we set our labels y to $v_{\bar{u}}$. It is immediately clear that $\|\alpha_{\ell_1}\|_1 \leq \|\beta\|_1 \leq 1$.

To lower bound $\mathbb{E}_{\beta}[\|\alpha_{\ell_2}\|_2^2]$, from Claim D.8 we get

$$\begin{aligned}
\mathbb{E}_{\beta\sim\mathcal{A}}[\|\alpha_{\ell_2}\|_2^2] &= E_{\bar{u}\sim\mathbb{S}^{d-1}}[v_{\bar{u}}^\top K^{-1} v_{\bar{u}}] \\
&= E_{\bar{u}\sim\mathbb{S}^{d-1}}[\text{trace} K^{-1}(v_{\bar{u}} v_{\bar{u}}^\top)] \\
&= \text{trace}(K^{-1} E_{\bar{u}\sim\mathbb{S}^{d-1}}[v_{\bar{u}} v_{\bar{u}}^\top] \\
&= \frac{1}{\kappa}\text{trace}(K^{-1}K) && \text{(by definition of } K) \\
&= \frac{n}{\kappa}
\end{aligned}$$

□

Proof of Theorem D.6. We note that since the distribution $\mathcal{A}$ of Lemma D.7 does not depend on the dataset X, it must follow that

$$\mathbb{E}_{(x_i)_{i=1}^n \sim_{\text{iid}} \mathcal{N}(0, I_{d\times d})} \left[\mathbb{E}_{\beta\sim\mathcal{A}}[\|\alpha_{\ell_2}\|_2^2]\right] = \frac{n}{\kappa}$$
$$\mathbb{E}_{\beta\sim\mathcal{A}} \left[\mathbb{E}_{(x_i)_{i=1}^n \sim_{\text{iid}} \mathcal{N}(0, I_{d\times d})}[\|\alpha_{\ell_2}\|_2^2]\right] = \frac{n}{\kappa}$$

Thus, there exists $\beta^\star$ such that if we sample X i.i.d. from the standard normal and set $y_i = \langle \varphi(x_i), \beta^\star\rangle$, the expectation of $\|\alpha_{\ell_2}\|_2^2$ is at least $\frac{n}{\kappa}$. We choose p_{data} corresponding to this $\beta^\star$, with x sampled from the standard normal. Now it is clear that p_{data} will satisfy the norm conditions of Theorem D.6.

For the generalization bounds, with high probability $\|X\|_F = \Theta(\sqrt{nd})$ as x is sampled from the standard normal distribution. Thus, Lemma D.4 immediately gives the desired generalization error bounds for $L_{\ell_1\text{-reg}}$. On the other hand, if $\|\alpha_{\ell_2}\|_2 \geq \sqrt{\frac{n}{\kappa}}$, then the bound of Lemma D.5 is at least

$$\frac{\sqrt{\kappa}\|\alpha_{\ell_2}\|_2\|X\|_F}{n\sqrt{d}} \geq \Omega(1)$$

□

E Missing Proofs in Section 4

E.1 Detailed Setup

We first write our regularity assumptions on Φ, R, and V in more detail:

Assumption E.1 (Regularity conditions on Φ, R, V)**.** R is nonnegative, Lipschitz, and smooth: $\exists M_R, C_R$ such that $\|\nabla^2 R\|_{op} \leq C_R$, and $\|\nabla R\|_2 \leq M_R$.

Assumption E.2. Φ is differentiable, bounded and Lipschitz on the sphere: $\exists B_\Phi, M_\Phi$ such that $\|\Phi(\bar\theta)\| \leq B_\Phi\ \forall \bar\theta \in \mathbb{S}^d$, and $|\Phi_i(\bar\theta) - \Phi_i(\bar\theta')| \leq M_\Phi \|\bar\theta - \bar\theta'\|_2\ \forall \bar\theta, \bar\theta' \in \mathbb{S}^d$.

Assumption E.3. V is Lipschitz and upper and lower bounded on the sphere: $\exists b_V, B_V, M_V$ such that $0 < b_V \leq V(\bar\theta) \leq B_V\ \forall \bar\theta \in \mathbb{S}^d$, and $\|\nabla V(\bar\theta)\|_2 \leq M_V\ \forall \bar\theta \in \mathbb{S}^d$.

We state the version of Theorem 4.3 that collects these parameters:

Theorem E.4 (Theorem 4.3 with problem parameters)**.** *Suppose that Φ and V are 2-homogeneous and Assumptions E.1, E.2, and E.3 hold. Fix a desired error threshold $\epsilon > 0$. Suppose that from a starting distribution ρ_0, a solution to the dynamics in equation 4.2 exists. Choose*

$$\sigma \triangleq \exp(-d\log(1/\epsilon)\text{poly}(k, M_V, M_R, M_\Phi, b_V, B_V, C_R, B_\Phi, L[\rho_0] - L^\star))$$
$$t_\epsilon \triangleq \frac{d^2}{\epsilon^4}\text{poly}(\log(1/\epsilon), k, M_V, M_R, M_\Phi, b_V, B_V, C_R, B_\Phi, L[\rho_0] - L^\star)$$

Then it must hold that $\min_{0\leq t\leq t_\epsilon} L[\rho_t] - \inf_\rho L[\rho] \leq 2\epsilon$.

E.2 Proof of Theorem E.4

Throughout the proof, it will be useful to keep track of $W_t \triangleq \sqrt{\mathbb{E}_{\theta\sim\rho_t}[\|\theta\|_2^2]}$, the second moment of ρ_t. We first introduce a general lemma on integrals over vector field divergences.

Lemma E.5. *For any $h_1 : \mathbb{R}^{d+1} \to \mathbb{R}$, $h_2 : \mathbb{R}^{d+1} \to \mathbb{R}^{d+1}$ and distribution ρ with $\rho(\theta) \to 0$ as $\|\theta\| \to \infty$,*

$$\int h_1(\theta)\nabla\cdot(h_2(\theta)\rho(\theta))d\theta = -E_{\theta\sim\rho}[\langle \nabla h_1(\theta), h_2(\theta)\rangle]$$

Proof. The proof follows from integration by parts. □

We note that ρ_t will satisfy the boundedness condition of Lemma E.5 during the course of our algorithm - ρ_0 starts with this property, and Lemma E.9 proves that ρ_t will continue to have this property. We therefore freely apply Lemma E.5 in the remaining proofs. We first bound the absolute value of $L'[\rho_t]$ over the sphere by $B_L \triangleq M_R B_\Phi + B_V$.

Lemma E.6. *For any* $\bar{\theta} \in \mathbb{S}^{d-1}, t \geq 0$, $|L'[\rho_t](\bar{\theta})| \leq \triangleq B_L$.

Proof. We compute

$$|L'[\rho_t](\bar{\theta})| = \left| \left\langle \nabla R \left(\int \Phi d\rho \right), \Phi(\bar{\theta}) \right\rangle + V(\bar{\theta}) \right|$$
$$\leq \left\| \nabla R \left(\int \Phi d\rho \right) \right\|_2 \|\Phi(\bar{\theta})\|_2 + V(\bar{\theta}) \leq M_R B_\Phi + B_V$$

□

Now we analyze the decrease in $L[\rho_t]$.

Lemma E.7. *Under the perturbed Wasserstein gradient flow*

$$\frac{d}{dt} L[\rho_t] = -\sigma \mathbb{E}_{\theta \sim \rho_t}[L'[\rho_t](\theta)] + \sigma \mathbb{E}_{\bar{\theta} \sim U^d}[L'[\rho_t](\bar{\theta})] - \mathbb{E}_{\theta \sim \rho_t}[\|v[\rho_t](\theta)\|_2^2]$$

Proof. Applying the chain rule, we can compute

$$\frac{d}{dt} L[\rho_t] = \left\langle \nabla R \left(\int \Phi d\rho_t \right), \frac{d}{dt} \int \Phi d\rho_t \right\rangle + \frac{d}{dt} \int V d\rho_t$$
$$= \frac{d}{dt} \mathbb{E}_{\theta \sim \rho_t}[L'[\rho_t](\theta)]$$
$$= \int L'[\rho_t](\theta) \rho_t'(\theta) d\theta$$
$$= -\sigma \int L'[\rho_t] d\rho_t + \sigma \int L'[\rho_t] dU^d - \int L'[\rho_t](\theta) \nabla \cdot (v[\rho_t](\theta) \rho_t(\theta)) d\theta$$
$$= -\sigma \mathbb{E}_{\theta \sim \rho_t}[L'[\rho_t](\theta)] + \sigma \mathbb{E}_{\bar{\theta} \sim U^d}[L'[\rho_t](\bar{\theta})] - \mathbb{E}_{\theta \sim \rho_t}[\|v[\rho_t](\theta)\|_2^2],$$

where we use Lemma E.5 with $h_1 = L'[\rho_t]$ and $h_2 = v[\rho_t]$. □

Now we show that the decrease in objective value is approximately the average velocity of all parameters under ρ_t plus some additional noise on the scale of σ. At the end, we choose σ small enough so that the noise terms essentially do not matter.

Corollary E.8. *We can bound* $\frac{d}{dt} L[\rho_t]$ *by*

$$\frac{d}{dt} L[\rho_t] \leq \sigma B_L (W_t^2 + 1) - \mathbb{E}_{\theta \sim \rho_t}[\|v[\rho_t](\theta)\|_2^2] \tag{E.1}$$

Proof. By homogeneity, and Lemma E.6, $\mathbb{E}_{\theta \sim \rho_t}[L'[\rho_t](\theta)] = \mathbb{E}_{\theta \sim \rho_t}[L'[\rho_t](\bar{\theta}) \|\theta\|_2^2] \leq B_L W_t^2$. We also get $\mathbb{E}_{\bar{\theta} \sim U^d}[L'[\rho_t](\bar{\theta})] \leq B_L$ since U^d is only supported on $\mathbb{S}^d$. Combining these with Lemma E.7 gives the desired statement. □

Now we show that if we run the dynamics for a short time, the second moment of ρ_t will grow slowly, again at a rate that is roughly the scale of the noise σ.

Lemma E.9. *For all* $0 \leq t' \leq t$, $W_{t'}^2 \leq \frac{L[\rho_0] + \sigma t B_L}{b_V - t \sigma B_L}$.

Proof. Let $t^* \triangleq \arg\max_{t' \in [0,t]} W_{t'}^2$. Integrating both sides of equation E.1, and rearranging, we get

$$0 \leq \int_0^{t^*} \mathbb{E}_{\theta \sim \rho_s}[\|v[\rho_s](\theta)\|_2^2] ds \leq L[\rho_0] - L[\rho_t] + \sigma B_L \int_0^{t^*} (W_s^2 + 1) ds$$
$$\leq L[\rho_0] - L[\rho_{t^*}] + t^* \sigma B_L (W_{t^*}^2 + 1)$$

Now since R is nonnegative, we apply $L[\rho_{t^*}] \geq E_{\theta \sim \rho_{t^*}}[V(\theta)] \geq E_{\theta \sim \rho_{t^*}}[V(\bar{\theta}) \|\theta\|_2^2] \geq b_V W_{t^*}^2$. We now plug this in and rearrange to get $W_{t'}^2 \leq W_{t^*}^2 \leq \frac{L[\rho_0] + t^* \sigma B_L}{b_V - t^* \sigma B_L} \leq \frac{L[\rho_0] + t \sigma B_L}{b_V - t \sigma B_L}$ $\forall 0 \leq t' \leq t$. □

Now let $W_\epsilon^2 \triangleq \frac{L[\rho_0]+\sigma t_\epsilon B_L}{b_v - t_\epsilon \sigma B_L}$. By Lemma E.9, $\forall 0 \le t \le t_\epsilon$, $W_t^2 \le W_\epsilon^2$.

The next statement allows us to argue that our dynamics will never increase the objective by too much.

Lemma E.10. *For any t_1, t_2 with $0 \le t_1 \le t_2 \le t_\epsilon$, $L[\rho_{t_2}] - L[\rho_{t_1}] \le \sigma(t_2 - t_1)B_L(W_\epsilon^2 + 1)$.*

Proof. From Corollary E.8, $\forall t \in [t_1, t_2]$ we have

$$\frac{d}{dt}L[\rho_t] \le \sigma B_L(W_\epsilon^2 + 1)$$

Integrating from t_1 to t_2 gives the desired result. □

The following lemma bounds the change in expectation of a 2-homogeneous function over ρ_t. At a high level, we lower bound the decrease in our loss as a function of the change in this expectation.

Lemma E.11. *Let $h : \mathbb{R}^{d+1} \to \mathbb{R}$ that is 2-homogeneous, with $\|\nabla h(\bar\theta)\| \le M$ $\forall \bar\theta \in \mathbb{S}^d$ and $|h(\bar\theta)| \le B$ $\forall \bar\theta \in \mathbb{S}^d$. Then $\forall 0 \le t \le t_\epsilon$, we have*

$$\left|\frac{d}{dt}\int h d\rho\right| \le \sigma B(W_\epsilon^2 + 1) + MW\left(-\frac{d}{dt}L[\rho_t] + \sigma B_L(W_\epsilon^2 + 1)\right)^{1/2} \tag{E.2}$$

Proof. Let $Q(t) \triangleq \int h d\rho_t$. We can compute:

$$\begin{aligned} Q'(t) &= \int h(\theta)\frac{d\rho_t}{dt}(\theta)d\theta \\ &= \int h(\theta)(-\sigma\rho_t(\theta) - \nabla\cdot(v[\rho_t](\theta)\rho_t(\theta)))d\theta + \sigma\int h dU^d \\ &= -\sigma\int h(\bar\theta)\|\theta\|_2^2\rho_t(\theta)d\theta + \sigma\int h dU^d - \int h(\theta)\nabla\cdot(v[\rho_t](\theta)\rho_t(\theta))d\theta \end{aligned} \tag{E.3}$$

Note that the first two terms are bounded by $\sigma B(W_\epsilon^2 + 1)$ by the assumptions for the lemma. For the third term, we have from Lemma E.5:

$$\begin{aligned} \left|\int h(\theta)\nabla\cdot(v[\rho_t](\theta)\rho_t(\theta))d\theta\right| &= |E_{\theta\sim\rho_t}[\langle\nabla h(\theta), v[\rho_t](\theta)\rangle]| \\ &\le \sqrt{E_{\theta\sim\rho_t}[\|\nabla h(\theta)\|_2^2]E_{\theta\sim\rho_t}[\|v[\rho_t](\theta)\|_2^2]} && \text{(by Cauchy-Schwarz)} \\ &\le \sqrt{E_{\theta\sim\rho_t}[\|\nabla h(\bar\theta)\|_2^2\|\theta\|_2^2]E_{\theta\sim\rho_t}[\|v[\rho_t](\theta)\|_2^2]} && \text{(by homogeneity of } \nabla h) \\ &\le MW_\epsilon\sqrt{E_{\theta\sim\rho_t}[\|v[\rho_t](\theta)\|_2^2]} && \text{(since } h \text{ is Lipschitz on the sphere)} \\ &\le MW_\epsilon\left(-\frac{d}{dt}L[\rho_t] + \sigma B_L(W_\epsilon^2 + 1)\right)^{1/2} && \text{(by Corollary E.8)} \end{aligned}$$

Plugging this into equation E.3, we get that

$$|Q'(t)| \le \sigma B(W_\epsilon^2 + 1) + MW_\epsilon\left(-\frac{d}{dt}L[\rho_t] + \sigma B_L(W_\epsilon^2 + 1)\right)^{1/2}$$

□

We apply this result to bound the change in $L'[\rho_t]$ over time in terms of the change of the objective value. For clarity, we write the bound in terms of c_1 that is some polynomial in the problem constants.

Lemma E.12. *Define $Q(t) \triangleq \int \Phi d\rho_t$. For every $\bar\theta \in \mathbb{S}^d$ and $0 \le t \le t + l \le t_\epsilon$, $\exists c_1 \triangleq$ poly$(k, C_R, B_\Phi, M_\Phi, B_L)$ such that*

$$|L'[\rho_t](\bar\theta) - L'[\rho_{t+l}](\bar\theta)| \le C_R B_\Phi \int_t^{t+l} \|Q'(t)\|_1 \tag{E.4}$$

$$\le \sigma l c_1(W_\epsilon^2 + 1) + c_1 W_\epsilon\sqrt{l}(L[\rho_t] - L[\rho_{t+l}] + \sigma l c_1(W_\epsilon^2 + 1))^{1/2} \tag{E.5}$$

Proof. Recall that $L'[\rho_t](\bar{\theta}) = \langle \nabla R(\int \Phi d\rho_t), \Phi(\bar{\theta})\rangle + V(\bar{\theta})$. Differentiating with respect to t,

$$\begin{aligned}\frac{d}{dt}L'[\rho_t](\bar{\theta}) &= \left\langle \frac{d}{dt}\nabla R\left(\int \Phi d\rho_t\right), \Phi(\bar{\theta})\right\rangle \\ &= \Phi(\bar{\theta})^\top \nabla^2 R(Q(t))Q'(t) \\ &\leq C_R B_\Phi \|Q'(t)\|_2 \\ &\leq C_R B_\Phi \|Q'(t)\|_1 \end{aligned} \tag{E.6}$$

Integrating and applying the same reasoning to $-L'[\rho_t]$ gives us equation E.4. Now we apply Lemma E.11 to get

$$\begin{aligned}\|Q'(t)\|_1 &= \sum_{i=1}^k \left|\frac{d}{dt}\int \Phi_i d\rho_t\right| \\ &\leq \sum_{i=1}^k \left[\sigma B_\Phi(W_\epsilon^2+1) + M_\Phi W_\epsilon \left(-\frac{d}{dt}L[\rho_t] + \sigma B_L(W_\epsilon^2+1)\right)^{1/2}\right] \\ &\leq k\sigma B_\Phi(W_\epsilon^2+1) + kM_\Phi W_\epsilon \left(-\frac{d}{dt}L[\rho_t] + \sigma B_L(W_\epsilon^2+1)\right)^{1/2}\end{aligned}$$

We plug this into equation E.6 and then integrate both sides to obtain

$$\begin{aligned}&C_R B_\Phi \int_t^{t+l} \|Q'(t)\|_1 \\ &\leq k\sigma l C_R B_\Phi^2(W_\epsilon^2+1) + kC_R B_\Phi M_\Phi W_\epsilon \int_t^{t+l}\left(-\frac{d}{dt}L[\rho_t] + \sigma B_L(W_\epsilon^2+1)\right)^{1/2} \\ &\leq k\sigma l C_R B_\Phi^2(W_\epsilon^2+1) + kC_R B_\Phi M_\Phi W_\epsilon \sqrt{l}(L[\rho_t] - L[\rho_{t+l}] + \sigma l B_L(W_\epsilon^2+1))^{1/2}\end{aligned}$$

Using $c_1 \triangleq \max\{kC_R B_\Phi^2, kC_R B_\Phi M_\Phi, B_L\}$ gives the statement in the lemma. □

Now we also show that L' is Lipschitz on the unit ball. For clarity, we let $c_2 \triangleq \sqrt{k}M_R M_\Phi + M_V$.

Lemma E.13. *For all $\bar{\theta}, \bar{\theta}' \in \mathbb{S}^d$,*

$$|L'[\rho](\bar{\theta}) - L'[\rho](\bar{\theta}')| \leq c_2\|\bar{\theta} - \bar{\theta}'\|_2 \tag{E.7}$$

Proof. Using the definition of L' and triangle inequality,

$$\begin{aligned}|L'[\rho](\bar{\theta}) - L'[\rho](\bar{\theta}')| &\leq \left\|\nabla R\left(\int \Phi d\rho\right)\right\|_2 \|\Phi(\bar{\theta}) - \Phi(\bar{\theta}')\|_2 + |V(\bar{\theta}) - V(\bar{\theta}')| \\ &\leq (\sqrt{k}M_R M_\Phi + M_V)\|\bar{\theta} - \bar{\theta}'\|_2 \qquad \text{(by definition of } M_\Phi, M_R, M_V)\end{aligned}$$

□

Now the remainder of the proof will proceed as follows: we show that if ρ_t is far from optimality, either the expected velocity of θ under ρ_t will be large in which case the loss decreases from Corollary E.8, or there will exist $\bar{\theta}$ such that $L'[\rho_t](\bar{\theta}) \ll 0$. We will first show that in the latter case, the σU^d noise term will grow mass exponentially fast in a descent direction until we make progress. Define $K_t^{-\tau} \triangleq \{\bar{\theta} \in \mathbb{S}^d : L'[\rho_t](\bar{\theta}) \leq -\tau\}$, the $-\tau$-sublevel set of $L'[\rho_t]$, and let $m(S) \triangleq \mathbb{E}_{\theta\sim U^d}[\mathbf{1}(\theta \in S)]$ be the normalized spherical area of the set S.

Lemma E.14. *If $K_t^{-\tau}$ is nonempty, for $0 \leq \delta \leq \tau$, $\log m(K_t^{-\tau+\delta}) \geq -2d\log\frac{c_2}{\delta}$.*

Proof. Let $\bar{\theta} \in K_t^{-\tau}$. From Lemma E.13, $L'[\rho](\bar{\theta}') \leq -\tau + \delta$ for all $\bar{\theta}'$ with $\|\bar{\theta}' - \bar{\theta}\|_2 \leq \frac{\delta}{c_2}$. Thus, we have

$$m(K_t^{-\tau+\delta}) \geq \mathbb{E}_{\bar{\theta}'\sim U^d}\left[\mathbf{1}[\|\bar{\theta}' - \bar{\theta}\|_2 \leq \frac{\delta}{c_2}]\right]$$

Now the statement follows by Lemma 2.3 of (Ball et al., 1997). □

Now we show that if a descent direction exists, the added noise will find it and our function value will decrease. We start with a general lemma about the magnitude of the gradient of a 2-homogeneous function in the radial direction.

Lemma E.15. *Let $h : \mathbb{R}^{d+1} \to \mathbb{R}$ be a 2-homogeneous function. Then for any $\theta \in \mathbb{R}^{d+1}$, $\bar{\theta}^\top \nabla h(\theta) = 2\|\theta\|_2 h(\bar{\theta})$.*

Proof. We have $h(\theta + \alpha\bar{\theta}) = (\|\theta\|_2 + \alpha)^2 h(\bar{\theta})$. Differentiating both sides with respect to α and evaluating the derivative at 0, we get $\bar{\theta}^\top \nabla h(\theta) = 2\|\theta\|_2 h(\bar{\theta})$, as desired. □

We state the lemma claiming that our objective will decrease if $L'[\rho_t](\bar{\theta}) \ll 0$ for some $\bar{\theta} \in \mathbb{S}^d$.

Lemma E.16. *Choose*

$$l \geq \frac{\log(W_\epsilon^2/\sigma) + 2d\log\frac{2c_2}{\tau}}{\tau - \sigma} + 1$$

If $K_{t^*}^{-\tau}$ is nonempty for some t^* satisfying $t^* + l \leq t_\epsilon$, then after l steps, we will have

$$L[\rho_{t^*+l}] \leq L[\rho_{t^*}] - \frac{(\tau/4 - \sigma l c_1(W_\epsilon^2 + 1))^2}{l c_1^2 W_\epsilon^2} + \sigma l c_1(W_\epsilon^2 + 1) \tag{E.8}$$

We will first show that a descent direction in $L'[\rho_t]$ will remain for the next l time steps. In the notation of Lemma E.12, define $z(s) \triangleq C_R B_\Phi \int_{t^*}^{t^*+s} \|Q'(t)\|_1 dt$. Note that from Lemma E.12, for all $\bar{\theta} \in \mathbb{S}^d$ we have $|L'[\rho_{t^*+s}](\bar{\theta}) - L'[\rho_{t^*}](\bar{\theta})| \leq z(s)$. Thus, the following holds:

Claim E.17. *For all $s \leq l$, $K_{t^*+s}^{-\tau+z(s)}$ is nonempty.*

Proof. By assumption, $\exists\bar{\theta}$ with $\bar{\theta} \in K_{t^*}^{-\tau}$. Then $L'[\rho_{t^*+s}](\bar{\theta}) \leq L'[\rho_{t^*}](\bar{\theta}) + z(s) \leq -\tau + z(s)$, so $K_{t^*+s}^{-\tau+z(s)}$ is nonempty. □

Let $T_s \triangleq K_{t^*+s}^{-\tau/2+z(s)}$ for $0 \leq s \leq l$. We now argue that this set T_s does not shrink as t increases.

Claim E.18. *For all $s' > s$, $T_{s'} \supseteq T_s$.*

Proof. From equation E.6 and the definition of $z(s)$, $|L'[\rho_{t+s'}](\bar{\theta}) - L'[\rho_{t+s}](\bar{\theta})| \leq z(s') - z(s)$. It follows that for $\bar{\theta} \in T_s$

$$\begin{aligned} L'[\rho_{t+s'}](\bar{\theta}) &\leq L'[\rho_{t+s}](\bar{\theta}) + z(s') - z(s) \\ &\leq -\tau/2 + z(s) - z(s) + z(s') && \text{(by definition of } T_s) \\ &\leq -\tau/2 + z(s') \end{aligned}$$

which means that $\bar{\theta} \in T_{s'}$. □

Now we show that the weight of the particles in T_s grows very fast if $z(k)$ is small.

Claim E.19. *Suppose that $z(l) \leq \tau/4$. Let $\tilde{T}_s = \{\theta \in \mathbb{R}^{d+1} : \bar{\theta} \in T_s\}$. Define $N(s) \triangleq \int_{\tilde{T}_s} \|\theta\|^2 d\rho_{t^*+s}$ and $\beta \triangleq \exp(-2d\log\frac{2c_2}{\tau})$. Then $N'(s) \geq (\tau - \sigma)N(s) + \sigma\beta$.*

Proof. From the assumption $z(l) \leq \frac{\tau}{4}$, it holds that $T_s \subseteq K_{t^*+s}^{-\tau/4}\ \forall s \leq k$. Since T_s is defined as a sublevel set, $v[\rho_{t^*+s}](\bar{\theta})$ points inwards on the boundary of T_s for all $\bar{\theta} \in T_s$, and by 1-homogeneity of the gradient, the same must hold for all $u \in \tilde{T}_s$.

Now consider any particle $\theta \in \tilde{T}_s$. We have that θ flows to $\theta + v[\rho_{t^*+s}](\theta)ds$ at time $t^* + s + ds$. Furthermore, since the gradient points inwards from the boundary, it also follows that $u + v[\rho_{t^*+s}](\theta)ds \in \tilde{T}_s$. Now we compute

$$\begin{aligned} \int_{\tilde{T}_s} \|\theta\|_2^2 d\rho_{t^*+s+ds} &= (1 - \sigma ds)\int_{\tilde{T}_s} \|\theta + v[\rho_{t^*+s}](\theta)ds\|_2^2 d\rho_{t^*+s} + \sigma ds \int_{\tilde{T}_s} 1 dU^d \\ &\geq (1 - \sigma ds)\int_{\tilde{T}_s} (\|\theta\|_2^2 + 2\theta^\top v[\rho_{t^*+s}](\theta)ds) d\rho_{t^*+s} + \sigma m(K_{t^*+s}^{-\tau/2+z(s)})ds \end{aligned} \tag{E.9}$$

Now we apply Lemma E.15, using the 2-homogeneity of F' and the fact that $L'[\rho_{t^*+s}](\bar{\theta}) \leq -\tau/4 \; \forall \theta \in \tilde{T}_s$

$$\begin{aligned} \|\theta\|_2^2 + 2\theta^\top v[\rho_{t^*+s}](\theta)ds &= \|\theta\|_2^2 - 4\|\theta\|_2^2 L'[\rho_{t^*+s}](\bar{\theta})ds \\ &\geq \|\theta\|_2^2(1+\tau ds) \end{aligned} \tag{E.10}$$

Furthermore, since $K_{t^*+s}^{-\tau+z(s)}$ is nonempty by Claim E.17, we can apply Lemma E.14 and obtain

$$m(K_{t^*+s}^{-\tau/2+z(s)}) \geq \beta \tag{E.11}$$

Plugging equation E.10 and equation E.11 back into equation E.9, we get

$$\int_{\tilde{T}_s} \|u\|_2^2 d\rho_{t^*+s+ds} \geq (1-\sigma ds)(1+2\tau ds)N(s) + \sigma\beta ds$$

Since we also have that $\tilde{T}_{s+ds} \supseteq \tilde{T}_s$, it follows that

$$N(s+ds) = \int_{\tilde{T}_{s+ds}} \|u\|_2^2 d\rho_{t^*+s+ds} \geq (1-\sigma ds)(1+\tau ds)N(s) + \sigma\beta ds$$

and so $N'(s) \geq (\tau - \sigma)N(s) + \sigma\beta$. □

Now we are ready to prove Lemma E.16.

Proof of Lemma E.16. If $z(l) = C_R B_\Phi \int_t^{t+l} \|Q'(t)\|_1 \geq \frac{\tau}{4}$, then by rearranging the conclusion of Lemma E.12 we immediately get equation E.8.

Suppose for the sake of contradiction that $z(l) \leq \tau/4$. From Claim E.19, it follows that $N(1) \geq \sigma\beta$, and $N(l) \geq \exp((\tau-\sigma)(l-1))N(1)$. Thus, in $\frac{\log(W_\epsilon^2/\sigma)+2d\log\frac{2c_2}{\tau}}{\tau-\sigma}+1$ time, $W_{t^*+l} \geq N(l) \geq W_\epsilon^2$, a contradiction. Therefore, it must be true that $z(l) \geq \tau/4$.

□

The following lemma will be useful in showing that the objective will decrease fast when ρ_t is very suboptimal.

Lemma E.20. *For any time t with $0 \leq t \leq t_\epsilon$, we have*

$$\frac{d}{dt}L[\rho_t] \leq \sigma B_L(W_\epsilon^2+1) - \frac{\mathbb{E}_{\theta\sim\rho_t}[L'[\rho_t](\theta)]^2}{W_\epsilon^2} \tag{E.12}$$

Proof. We can first compute

$$\begin{aligned} \mathbb{E}_{\theta\sim\rho_t}[L'[\rho_t](\theta)] &= \mathbb{E}_{\theta\sim\rho_t}[L'[\rho_t](\bar{\theta})\|\theta\|_2^2] \\ &= \frac{1}{2}\mathbb{E}_{\theta\sim\rho_t}[\|\theta\|_2\bar{\theta}^\top v[\rho_t](\theta)] && \text{(via Lemma E.15)} \\ &\leq \frac{1}{2}\sqrt{\mathbb{E}_{\theta\sim\rho_t}[\|\theta\|_2^2]\mathbb{E}_{\theta\sim\rho_t}[\|v[\rho_t](\theta)\|_2^2]} && \text{(by Cauchy-Schwarz)} \\ &\leq \frac{1}{2}W_\epsilon\sqrt{\mathbb{E}_{\theta\sim\rho_t}[\|v[\rho_t](\theta)\|_2^2]} \end{aligned}$$

Rearranging gives $\mathbb{E}_{\theta\sim\rho_t}[\|v[\rho_t](\theta)\|_2^2] \geq \frac{\mathbb{E}_{\theta\sim\rho_t}[L'[\rho_t](\theta)]^2}{W_\epsilon^2}$, and plugging this into equation E.1 gives the desired result. □

Proof of Theorem E.4. Let $L^\star$ denote the infimum $\inf_\rho L[\rho]$, and let $\rho^\star$ be an ϵ-approximate global minimizer of L: $L[\rho^\star] \leq L^\star + \epsilon$. (We define $\rho^\star$ because a true minimizer of L might not exist.) Let $W^\star \triangleq \mathbb{E}_{\theta\sim\rho^\star}[\|\theta\|_2^2]$. We first note that since $b_V W^{\star 2} \leq L[\rho^\star] \leq L[\rho_0]$, $W^{\star 2} \leq L[\rho_0]/b_V \leq W_\epsilon^2$.

Now we bound the suboptimality of ρ_t: since L is convex in ρ,

$$L[\rho^\star] \geq L[\rho_t] + \mathbb{E}_{\theta\sim\rho^\star}[L'[\rho_t](\theta)] - \mathbb{E}_{\theta\sim\rho_t}[L'[\rho_t](\theta)]$$

Rearranging gives

$$L[\rho_t] - L[\rho^\star] \leq \mathbb{E}_{\theta\sim\rho_t}[L'[\rho_t](\theta)] - \mathbb{E}_{\theta\sim\rho^\star}[L'[\rho_t](\theta)]$$
$$\leq \mathbb{E}_{\theta\sim\rho_t}[L'[\rho_t](\theta)] - W^{\star 2}\min\left\{\min_{\bar\theta\in\mathbb{S}^{d-1}} L'[\rho_t](\bar\theta), 0\right\} \tag{E.13}$$

Now let $l \triangleq \frac{W_\epsilon^2}{\epsilon - 2W_\epsilon^2\sigma}\left(2\log\frac{W_\epsilon^2}{\sigma} + 2d\log\frac{4W_\epsilon^2 c_2}{\epsilon}\right)$, which satisfies Lemma E.16 with the value of τ later specified. Suppose that there is a t with $0 \leq t \leq t_\epsilon - 2l$ and $\forall t' \in [t, t+2l]$, $L[\rho_{t'}] - L^\star \geq 2\epsilon$. Then $L[\rho_{t'}] - L[\rho^\star] \geq \epsilon$. We will argue that the objective decreases when we are ϵ suboptimal:

$$L[\rho_t] - L[\rho_{t+2l}] \geq \tag{E.14}$$
$$\min\left\{\frac{(\epsilon/8W_\epsilon^2 - l\sigma c_1(W_\epsilon^2+1))^2}{c_1^2 W_\epsilon^2 l} - 3\sigma l c_1(W_\epsilon^2+1), l\frac{\epsilon^2}{4W_\epsilon^2} - 2\sigma l B_L(W_\epsilon^2+1)\right\} \tag{E.15}$$

Using equation E.13 and $W_\epsilon \geq W^\star$, we first note that

$$\epsilon \leq \mathbb{E}_{\theta\sim\rho_{t'}}[L'[\rho_{t'}](\theta)] - {W_\epsilon}^2\min\left\{\min_{\bar\theta\in\mathbb{S}^{d-1}} L'[\rho_{t'}](\bar\theta), 0\right\}\ \forall t' \in [t, t+l]$$

Thus, either $\min_{\bar\theta\in\mathbb{S}^d} L'[\rho_{t'}](\bar\theta) \leq -\frac{\epsilon}{2W^{\star 2}} \leq -\frac{\epsilon}{2W_\epsilon^2}$, or $\mathbb{E}_{\theta\sim\rho_{t'}}[L'[\rho_{t'}](\theta)] \geq \frac{\epsilon}{2}$. If $\exists t' \in [t, t+l]$ such that the former holds, then the $\tau \triangleq \frac{\epsilon}{2W_\epsilon^2}$ sub-level set $K_{t'}^{-\tau}$ is non-empty. Applying Lemma E.16 gives

$$L[\rho_{t'}] - L[\rho_{t'+l}] \geq \frac{(\epsilon/8W_\epsilon^2 - l\sigma c_1(W_\epsilon^2+1))^2}{c_1^2 W_\epsilon^2 l} - \sigma l c_1(W_\epsilon^2+1)$$

Furthermore, from Lemma E.10, $L[\rho_{t+2l}] - L[\rho_{t'+l}] \leq \sigma l c_1(W_\epsilon^2+1)$ and $L[\rho_{t'}] - L[\rho_t] \leq \sigma l B_L(W_\epsilon^2+1)$, and so combining gives

$$L[\rho_t] - L[\rho_{t+2k}] \geq \frac{(\epsilon/8W_\epsilon^2 - l\sigma c_1(W_\epsilon^2+1))^2}{c_1^2 W_\epsilon^2 l} - 3\sigma l c_1(W_\epsilon^2+1) \tag{E.16}$$

In the second case $\mathbb{E}_{\theta\sim\rho_{t'}}[L'[\rho_{t'}](\theta)] \geq \frac{\epsilon}{2}$, $\forall t' \in [t, t+l]$. Therefore, we can integrate equation E.12 from t to $t+l$ in order to get

$$L[\rho_t] - L[\rho_{t+l}] \geq l\frac{\epsilon^2}{4W_\epsilon^2} - \sigma l B_L(W_\epsilon^2+1)$$

Therefore, applying Lemma E.10 again gives

$$L[\rho_t] - L[\rho_{t+2l}] \geq l\frac{\epsilon^2}{4W_\epsilon^2} - 2\sigma l B_L(W_\epsilon^2+1) \tag{E.17}$$

Thus equation E.15 follows.

Now recall that we choose

$$\sigma \triangleq \exp(-d\log(1/\epsilon)\text{poly}(k, M_V, M_R, M_\Phi, b_V, B_V, C_R, B_\Phi, L[\rho_0] - L[\rho^\star]))$$

For the simplicity, in the remaining computation, we will use $O(\cdot)$ notation to hide polynomials in the problem parameters besides d, ϵ. We simply write $\sigma = \exp(-c_3 d\log(1/\epsilon))$. Recall our choice $t_\epsilon \triangleq O(\frac{d^2}{\epsilon^4}\log^2(1/\epsilon))$. It suffices to show that our objective would have sufficiently decreased in t_ϵ steps. We first note that with c_3 sufficiently large, $W_\epsilon^2 = O(L[\rho_0]/b_v) = O(1)$. Simplifying our expression for l, we get that $l = O(\frac{d}{\epsilon}\log\frac{1}{\epsilon})$, so long as $\sigma W_\epsilon^2 = o(\epsilon)$, which holds for sufficiently large c_3. Now let

$$\delta_1 \triangleq \frac{(\epsilon/8W_\epsilon^2 - l\sigma c_1(W_\epsilon^2+1))^2}{c_1^2 W_\epsilon^2 l} - 3\sigma l c_1(W_\epsilon^2+1)$$
$$\delta_2 \triangleq l\frac{\epsilon^2}{4W_\epsilon^2} - 2\sigma l B_L(W_\epsilon^2+1)$$

Again, for sufficiently large c_3, the terms with σ become negligible, and $\delta_1 = O(\frac{\epsilon^2}{l}) = O(\frac{\epsilon^3}{d\log(1/\epsilon)})$. Likewise, $\delta_2 = O(d\epsilon\log(1/\epsilon))$.

Thus, if by time t we have not encountered 2ϵ-optimal ρ_t, then we will decrease the objective by $O(\frac{\epsilon^3}{d\log(1/\epsilon)})$ in $O(\frac{d}{\epsilon}\log\frac{1}{\epsilon})$ time. Therefore, a total of $O(\frac{d^2}{\epsilon^4}\log^2(1/\epsilon))$ time is sufficient to obtain ϵ accuracy. □

Method	CIFAR10	CIFAR100
Weight decay annealing	5.86	26.22
Fixed weight decay	6.01	27.00

Table 1: Test error on CIFAR10 and CIFAR100 for initial $\lambda = 0.0005$.

E.3 Discrete-Time Optimization

To circumvent the technical issue of existence of a solution to the continuous-time dynamics, we also note that polynomial time convergence holds for discrete-time updates.

Theorem E.21. *Along with Assumptions E.1, E.2, E.3 additionally assume that $\nabla\Phi_i$ and ∇V are C_Φ and C_V-Lipschitz, respectively. Let ρ_t evolve according to the following discrete-time update:*

$$\rho_{t+1} \triangleq \rho_t + \eta(-\sigma\rho_t + \sigma U^d - \nabla \cdot (v[\rho_t]\rho_t))$$

There exists a choice of

$$\sigma \triangleq \exp(-d\log(1/\epsilon))\text{poly}(k, M_V, M_R, b_V, B_V, C_R, B_\Phi, C_\Phi, C_V, L[\rho_0] - L[\rho^\star]))$$
$$\eta \triangleq \text{poly}(k, M_V, M_R, b_V, B_V, C_R, B_\Phi, C_\Phi, C_V, L[\rho_0] - L[\rho^\star])$$
$$t_\epsilon \triangleq \frac{d^2}{\epsilon^4}\text{poly}(k, M_V, M_R, b_V, B_V, C_R, B_\Phi, C_\Phi, C_V, L[\rho_0] - L[\rho^\star])$$

such that $\min_{0\le t\le t_\epsilon} L[\rho_t] - L^\star \le \epsilon$.

The proof follows from a standard conversion of the continuous-time proof of Theorem E.4 to discrete time, and we omit it here for simplicity.

F Additional Experiments

We train a modified WideResNet architecture (Zagoruyko & Komodakis, 2016) on CIFAR10 and CIFAR100. Our theory does not entirely apply because the identity mapping prevents ResNet architectures from being homogeneous, but our experiments show that reducing weight decay can still help generalization error in this setting. Because batchnorm can cause the regularizer to have different effects (van Laarhoven, 2017), we remove batchnorm layers and train a 16 layer deep WideResNet. We again compare a network trained with weight decayed annealing to one trained without annealing. We used a fixed learning rate schedule that starts at 0.1 and decreases by a factor of 0.2 at epochs 60, 120, and 160. For CIFAR10, we use an initial weight decay of 0.0002 and decrease the weight decay by 0.2 at epoch 60, and then by 0.5 at epochs 90, 120, 140, 160. For CIFAR100, we initialize weight decay at 0.0005 and decrease it by 0.2 at epochs 60, 120, and 160. We tried different parameters for the initial weight decay and chose the ones that worked best for the model without annealing. We also tried using small weight decays at initialization, but these models failed to generalize well – we believe this is due to an optimization issue where the algorithm fails to find a true global minimum of the regularized loss. We believe that annealing the weight decay directs the optimization algorithm closer towards the global minima for small λ.

Table 1 shows the test error achieved by models with and without annealing. We see that the simple change of annealing weight decay can decrease the test error for this architecture.