THE CONVEX GEOMETRY OF BACKPROPAGATION: NEURAL NETWORK GRADIENT FLOWS CONVERGE TO EXTREME POINTS OF THE DUAL CONVEX PROGRAM

Yifei Wang
Department of Electrical Engineering
Stanford University
Stanford, CA 94305, USA
wangyf18@stanford.edu

Mert Pilanci
Department of Electrical Engineering
Stanford University
Stanford, CA 94305, USA
pilanci@stanford.edu

ABSTRACT

We study non-convex subgradient flows for training two-layer ReLU neural networks from a convex geometry and duality perspective. We characterize the implicit bias of unregularized non-convex gradient flow as convex regularization of an equivalent convex model. We then show that the limit points of non-convex subgradient flows can be identified via primal-dual correspondence in this convex optimization problem. Moreover, we derive a sufficient condition on the dual variables which ensures that the stationary points of the non-convex objective are the KKT points of the convex objective, thus proving convergence of non-convex gradient flows to the global optimum. For a class of regular training data distributions such as orthogonal separable data, we show that this sufficient condition holds. Therefore, non-convex gradient flows converge to optimal solutions of a convex optimization problem. We present numerical results verifying the predictions of our theory for non-convex subgradient descent.

1 INTRODUCTION

Neural networks (NNs) exhibit remarkable empirical performance in various machine learning tasks. However, a full characterization of the optimization and generalization properties of NNs is far from complete. Non-linear operations inherent to the structure of NNs, over-parameterization and the associated highly nonconvex training problem makes their theoretical analysis quite challenging.

In over-parameterized models such as NNs, one natural question arises: Which particular solution does gradient descent/gradient flow find in unregularized NN training problems? Suppose that $\mathbf{X} \in \mathbb{R}^{N\times d}$ is the training data matrix and $\mathbf{y} \in \{1, -1\}^N$ is the label vector. For linear classification problems such as logistic regression, it is known that gradient descent (GD) exhibits implicit regularization properties, see, e.g., (Soudry et al., 2018; Gunasekar et al., 2018). To be precise, under certain assumptions, GD converges to the following solution which maximizes the margin:

$$\arg\min_{\mathbf{w}\in\mathbb{R}^d} \frac{1}{2}\|\mathbf{w}\|_2^2, \text{ s.t. } y_n\mathbf{w}^T\mathbf{x}_n \geq 1, n \in [N]. \tag{1}$$

Here we denote $[N] = \{1, \ldots, N\}$. Recently, there are several results on the implicit regularization of the (stochastic) gradient descent method for NNs. In (Lyu & Li, 2019), for the multi-layer homogeneous network with exponential or cross-entropy loss, with separable training data, it is shown that the gradient flow (GF) and GD finds a stationary point of the following non-convex max-margin problem:

$$\arg\min_{\boldsymbol{\theta}} \frac{1}{2}\|\boldsymbol{\theta}\|_2^2, \text{ s.t. } y_n f(\boldsymbol{\theta}; \mathbf{x}_n) \geq 1, n \in [N], \tag{2}$$

where $f(\boldsymbol{\theta}; \mathbf{x})$ represents the output of the neural network with parameter $\boldsymbol{\theta}$ given input $\mathbf{x}$. In (Phuong & Lampert, 2021), by further assuming the orthogonal separability of the training data, it is shown that all neurons converge to one of the two max-margin classifiers. One corresponds to the data with positive labels, while the other corresponds to the data with negative labels.

However, as the max-margin problem of the neural network (2) is a non-convex optimization problem, the existing results only guarantee that it is a stationary point of (2), which can be a local minimizer or even a saddle point. In other words, the global optimality is not guaranteed.

In a different line of work (Pilanci & Ergen, 2020; Ergen & Pilanci, 2020; 2021b), exact convex optimization formulations of two and three-layer ReLU NNs are developed, which have global optimality guarantees in polynomial-time when the data has a polynomial number of hyperplane arrangements, e.g., in any fixed dimension or with convolutional networks of fixed filter size. The convex optimization framework was extended to vector output networks (Sahiner et al., 2021b), quantized networks (Bartan & Pilanci, 2021b), autoencoders (Sahiner et al., 2021c; Gupta et al., 2021), networks with polynomial activation functions (Bartan & Pilanci, 2021a), networks with batch normalization (Ergen et al., 2021), univariate deep ReLU networks, deep linear networks (Ergen & Pilanci, 2021c) and Generative Adversarial Networks (Sahiner et al., 2021a).

In this work, we first derive an equivalent convex program corresponding to the maximal margin problem(2). We then consider non-convex subgradient flow for unregularized logistic loss. We show that the limit points of non-convex subgradient flow can be identified via primal-dual correspondence in the convex optimization problem. We then present a sufficient condition on the dual variable to ensure that all stationary points of the non-convex max-margin problem are KKT points of the convex max-margin problem. For certain regular datasets including orthogonal separable data, we show that this sufficient condition on the dual variable holds, thus implies the convergence of gradient flow on the unregularized problem to the global optimum of the non-convex maximalo margin problem (2). Consequently, this enables us to fully characterize the implicit regularization of unregularized gradient flow or gradient descent as convex regularization applied to a convex model.

1.1 Related Work

There are several works studying the property of two-layer ReLU networks trained by gradient descent/gradient flow dynamics. The following papers study the gradient descent like dynamics in training two-layer ReLU networks for regression problems. Ma et al. (2020) show that for two-layer ReLU networks, only a group of a few activated neurons dominate the dynamics of gradient descent. In (Mei et al., 2018), the limiting dynamics of stochastic gradient descent (SGD) is captured by the distributional dynamics from a mean- field perspective and they utlize this to prove a general convergence result for noisy SGD. Li et al. (2020) focus on the case where the weights of the second layer are non-negative and they show that the over-parameterized neural network can learn the ground-truth network in polynomial time with polynomial samples. In (Zhou et al., 2021), it is shown that mildly over-parameterized student network can learn the teacher network and all student neurons converge to one of the teacher neurons.

Beyond (Lyu & Li, 2019) and (Phuong & Lampert, 2021), the following papers study the classification problems. In (Chizat & Bach, 2018), under certain assumptions on the training problem, with over-parameterized model, the gradient flow can converge to the global optimum of the training problem. For linear separable data, utilizing the hinge loss for classification, Wang et al. (2019) introduce a perturbed stochastic gradient method and show that it can attain the global optimum of the training problem. Similarly, for linear separable data, Yang et al. (2021) introduce a modified loss based on the hinge loss to enable (stochastic) gradient descent find the global minimum of the training problem, which is also globally optimal for the training problem with the hinge loss.

1.2 Problem setting

We focus on two-layer neural networks with ReLU activation, i.e., $f(\theta; X) = (XW_1)_+ w_2$, where $W_1 \in \mathbb{R}^{d \times m}$, $w_2 \in \mathbb{R}^m$ and $\theta = (W_1; w_2)$ represents the parameter. Due to the ReLU activation, this neural network is homogeneous, i.e., for any scalar $c \geq 0$, we have $f(c\theta; X) = c^2 f(\theta; X)$. The training problem is given by

$$\min_{\theta} \sum_{n=1}^{N} \ell(y_n f(\theta; x_n)); \tag{3}$$

where $\ell(q) : \mathbb{R} \to \mathbb{R}_+$ is the loss function. We focus on the logistic, i.e, cross-entropy loss, i.e., $\ell(q) = \log(1 + \exp(-q))$.

Then, we brie y review gradient descent and gradient ow. The gradient descent takes the update rule

$$(t+1) = (t) \quad (t)g(t);$$

where $g(t) \in @L((t))$ and $@$ represents the Clarke's subdifferential.

The gradient ow can be viewed as the gradient descent with in nitesimal step size. The trajectory of the parameter during training is an arc $: [0;+1) \to$, where $= f = (W_1; w_2) | W_1 \in R^{d \ m}; W_2 \in R^m g$. More precisely, the gradient ow is given by the differential inclusion

$$\frac{d}{dt} (t) \in @L((t)); \quad \text{for } t \quad 0, \text{ almost everywhere.}$$

2 Main Results

In this section, we present our main results and defer the detailed analysis to the following sections. Consider the more general multi-class version of the problem with classes. Suppose that $y \in [K]^N$ is the label vector. Let $Y = (y_{n;k})_{n \in [N]; k \in [K]} \in R^{N \ K}$ be the encoded label matrix such that

$$y_{n;k} = \begin{cases} 1; & \text{if } y_n = k; \\ 1; & \text{otherwise} \end{cases}$$

Similarly, we consider the following two-layer vector-output neural networks with ReLU activation:

$$F(; X) = \begin{bmatrix} f_1(_1; X) \\ \vdots \\ f_K(_K; X) \end{bmatrix} = \begin{bmatrix} (XW_1^{(1)})_+ w_2^{(1)} \\ \vdots \\ (XW_1^{(K)})_+ w_2^{(K)} \end{bmatrix};$$

where we write $= (_1; \ldots; _K)$. For $k = 1; \ldots; K$, we have $_k = (W_1^{(k)}; w_2^{(k)})$ where $W_1^{(k)} \in R^{N \ m}$ and $w_2^{(k)} \in R^m$. One can view each of the K outputs of $F(; X)$ as the output of a two-layer scalar-output neural network. Consider the following training problem:

$$\min \sum_{k=1}^{K} \sum_{n=1}^{N} `(y_{n;k} f_k(_k; x_n)): \tag{4}$$

According to (Lyu & Li, 2019), the gradient ow and the gradient descent nds a stationary point of the following non-convex max-margin problem:

$$\arg\min \sum_{k=1}^{K} \frac{1}{2} k _k k_2^2; \text{ s.t. } y_{n;k} f(_k; x_n) \quad 1; n \in [N]; k \in [K]: \tag{5}$$

Denote the set of all possible hyperplane arrangement as

$$P = f \text{diag}(I(Xw \quad 0)) | w \in R^d g; \tag{6}$$

and let $p = |P|$. We can also write $P = f D_1; \ldots; D_p g$. From (Cover, 1965), we have an upper bound $p \quad 2r \left(\frac{e(N \quad 1)}{r}\right)^r$ where $r = \text{rank}(X)$. We rst reformulate (5) as convex optimization.

Proposition 1 The non-convex problem (5) is equivalent to the following convex program

$$\begin{aligned} \min & \sum_{k=1}^{K} \sum_{j=1}^{p} (k u_{j;k} k_2 + k u'_{j;k} k_2); \\ \text{s.t. } & \text{diag}(y_k) \sum_{j=1}^{p} D_j X (u_{j;k} \quad u'_{j;k}) \quad 1; \\ & (2D_j \quad I) X u_{j;k} \quad 0; (2D_j \quad I) X u'_{j;k} \quad 0; j \in [p]; k \in [K]: \end{aligned} \tag{7}$$

where y_k is the k-th column of Y. The dual problem of (7) is given by

$$\begin{aligned} & \max \ \text{tr}(^T Y); \\ & \text{s.t. } \text{diag}(y_k) \ _k \quad 0; \max_{k w k_2 \ 1} j \ _k^T (X^T w)_+ j \quad 1; k \in [K]: \end{aligned} \tag{8}$$

where $_k$ is the k-th column of $$.

We present the detailed derivation of the convex formulation (7) and its dual problem (8) in the appendix. Given $u \in R^d$, we define $D(u) = \mathrm{diag}(I(Xu > 0))$. For two vectors $u, v \in R^d$, we define the cosine angle between u and v by $\cos\angle(u, v) = \frac{u^T v}{\|u\|_2 \|v\|_2}$.

2.1 Our contributions

The following theorem illustrate that for neurons satisfying $\mathrm{sign}(y_k^T (X w_{1,i}^{(k)})_+) = \mathrm{sign}(w_{2,i}^{(k)})$ at initialization, $w_{1,i}^{(k)}$ align to the direction of $X^T D(w_{1,i}^{(k)}) y_k$ at a certain time T, depending on $\mathrm{sign}(w_{2,i_{k,+}}^{(k)})$ at initialization. In Section 2.3, we show that these are dual extreme points of (7).

Theorem 1 Consider the K-class classification training problem (4) for any dataset. Suppose that the neural network is scaled at initialization such that $\|w_{1,i}^{(k)}\|_2 = |w_{2,i}^{(k)}|$ for $i \in [m]$ and $k \in [K]$. Assume that at initialization, for $k \in [K]$, there exists neurons $(w_{1,i_k}^{(k)}, w_{2,i_k}^{(k)})$ such that

$$\mathrm{sign}(y_k^T (X w_{1,i_k}^{(k)})_+) = \mathrm{sign}(w_{2,i_k}^{(k)}) = s, \tag{9}$$

where $s \in \{+1, -1\}$. Consider the subgradient flow applied to the non-convex problem (4). Let $\in (0, 1)$. Suppose that the initialization is sufficiently close to the origin. Then, for $k \in [K]$, there exist $T = T(\ , k)$ such that

$$\cos\angle\left(w_{1,i_k}^{(k)}(T), sX^T D(w_{1,i_k}^{(k)}(T)) y_k\right) \quad 1 \quad .$$

Next, we impose conditions on the dataset to prove a stronger global convergence results on the flow. We say that the dataset (X, y) is orthogonal separable among multiple classes if for all $n, n' \in [N]$,

$$x_n^T x_{n'} > 0, \text{ if } y_n = y_{n'},$$
$$x_n^T x_{n'} \quad 0, \text{ if } y_n \neq y_{n'}.$$

For orthogonal separable dataset among multiple classes, the subgradient flow for the non-convex problem (4) can find the global optimum of (5) up to a scaling constant.

Theorem 2 Suppose that $(X, y) \in R^{N \times d} \times [K]^N$ is orthogonal separable among multiple classes. Consider the non-convex subgradient flow applied to the non-convex problem (4). Suppose that the initialization is sufficiently close to the origin and scaled as in Theorem 1. Then, the non-convex subgradient flow converges to the global optimum of the convex program (7) and hence the non-convex objective (5) up to scaling.

Therefore, the above result characterizes the implicit regularization of unregularized gradient flow as convex regularization, i.e., group ℓ_1 norm, in the convex formulation (7). It is remarkable that group sparsity is enforced by small initialization magnitude with no explicit form of regularization.

2.2 Convex Geometry of Neural Gradient Flow

Suppose that $\in R^N$. Here we provide an interesting geometric interpretation behind the formula

$$\cos\angle(u, X^T D(u) \ ) > 1 \quad .$$

which describes a dual extreme point to which hidden neurons approach to as predicted by Theorem 1. We now explain the geometric intuition behind this result. Consider an ellipsoid $\{Xu : \|u\|_2 \quad 1\}$. A positive extreme point of this ellipsoid along the direction is defined by $\arg\max_{u : \|u\|_2 \ 1} \ ^T Xu$, which is given by the formula $\frac{X^T}{\|X^T\|_2}$. Next, we consider the rectified ellipsoid set $Q := \{(Xu)_+ : \|u\|_2 \quad 1\}$ introduced in (Ergen & Pilanci, 2021a) and shown in Figure 1. The constraint $\max_{u : \|u\|_2 \ 1} | \ ^T (Xu)_+| \quad 1$ on is equivalent to $\in Q$. Here Q is the absolute polar set of Q, which appears as a constraint in the convex program (8) and is defined as the following convex set

$$Q = \{ : \max_{z \in Q} | \ ^T z| \quad 1\}. \tag{10}$$

An extreme point of this non-convex body along the direction is given by the solution of the problem

$$\max_{u:\|u\|_2 \le 1} \|(Xu)_+\|_2^2 = \max_{D_j \in \mathcal{P}} \max_{u:\|u\|_2 \le 1, (2D_j - I)Xu \ge 0} u^T X^T D_j X u. \tag{11}$$

Here, $(\cdot\,; u)$ are primal-dual pairs as they appear in the convex dual program (8). First, note that a stationary point of gradient flow on the objective (11) is given by the identity $u \in \partial \|(Xu)_+\|_2^2$ where c is a constant. In particular, by picking the zero as the subgradient of $\|(x_n^T u)_+\|_2^2$ when $x_n^T u = 0$,

$$u = \frac{X^T D(u)}{\|X^T D(u)\|_2} = \frac{\sum_{n=1}^N a_n x_n \mathbb{I}(u^T x_n > 0)}{\|\sum_{n=1}^N a_n x_n \mathbb{I}(u^T x_n > 0)\|_2}. \tag{12}$$

Note that the formula $\cos\angle(u, X^T D(u)) > 1$ appearing in Theorem 1 shows that gradient flow reaches the extreme points of projected ellipsoids $\{Xu : \|u\|_2 \le 1\}$ in the direction of $u = y_k$, where $D_j \in \mathcal{P}$ corresponds to a valid hyperplane arrangement. This interesting phenomenon is depicted in Figures 3 and 4. The one-dimensional spikes in Figures 1 and 3 are projected ellipsoids. Detailed setup for Figure 1 to 4 and additional experiments can be found in Appendix F.

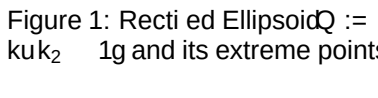

Figure 1: Rectified Ellipsoid $Q := \{(Xu)_+ : \|u\|_2 \le 1\}$ and its extreme points (spikes).

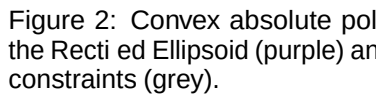

Figure 2: Convex absolute polar set Q° of the Rectified Ellipsoid (purple) and other dual constraints (grey).

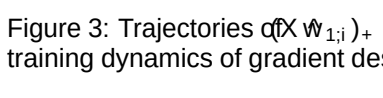

Figure 3: Trajectories of $(Xw_{1,i})_+$ along the training dynamics of gradient descent.

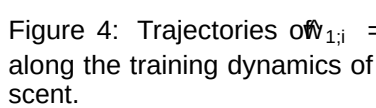

Figure 4: Trajectories of $w_{1,i} = \frac{w_{1,i}}{\|w_{1,i}\|_2}$ along the training dynamics of gradient descent.

Figure 5: Two-layer ReLU network gradient descent dynamics on an orthogonal separable dataset. $w_{1,i} = \frac{w_{1,i}}{\|w_{1,i}\|_2}$ is the normalized vector of the i-th hidden neuron in the first layer.

3 CONVEX MAX-MARGIN PROBLEM

In this section, we consider the equivalent convex model of the max-margin problem and its optimality conditions. We primarily focus on the binary classification problem for simplicity, which are later extended to the multi-class case. We can reformulate the nonconvex max-margin problem (2) as

$$\min \frac{1}{2}(\|W_1\|_F^2 + \|w_2\|_2^2), \text{ s.t. } Y(XW_1)_+ w_2 \geq 1, \tag{13}$$

where $Y = \mathrm{diag}(y)$. This is a nonconvex optimization problem due to the ReLU activation and the two-layer structure of neural network. Analogous to the convex formulation introduced in (Pilanci & Ergen, 2020) for regularized training problem of neural network, we can provide a convex optimization formulation of (13) and derive the dual problem.

Proposition 2 The problem (13) is equivalent to

$$\begin{aligned} P_{cvx} = \min & \sum_{j=1}^{p} (\|u_j\|_2 + \|u_j'\|_2), \\ \text{s.t. } & Y \sum_{j=1}^{p} D_j X(u_j' - u_j) \geq 1, \\ & (2D_j - I)Xu_j \geq 0, (2D_j - I)Xu_j' \geq 0, \forall j \in [p]. \end{aligned} \tag{14}$$

The dual problem of (14) is given by

$$d^* = \max y^T v \quad \text{s.t. } \|Y v\|_\infty \leq 0, \max_{u:\|u\|_2 \leq 1} |v^T (Xu)_+| \leq 1. \tag{15}$$

The following proposition gives a characterization of the KKT point of the non-convex max-margin problem (2). The definition of B-subdifferential can be found in Appendix A.

Proposition 3 Let (W_1, w_2, η) be a KKT point of the non-convex max-margin problem (2) (in terms of B-subdifferential). Suppose that $w_{2,i} \neq 0$ for certain $i \in [m]$. Then, there exists a diagonal matrix $\hat{D}_i \in \mathbb{R}^{N \times N}$ satisfying

$$\begin{aligned} (\hat{D}_i)_n &= 1, \text{ for } x_n^T w_{1,i} > 0, \\ (\hat{D}_i)_n &\in \{0, 1\}, \text{ for } x_n^T w_{1,i} = 0, \\ (\hat{D}_i)_n &= 0, \text{ for } x_n^T w_{1,i} < 0. \end{aligned}$$

such that

$$\frac{w_{1,i}}{w_{2,i}} = \frac{X^T \hat{D}_i w_{1,i}}{\|X^T \hat{D}_i w_{1,i}\|_2}, \|X^T \hat{D}_i w_{1,i}\|_2 = 1.$$

Based on the characterization of the KKT point of the non-convex max-margin problem (2), we provide an equivalent condition to ensure that it is also the KKT point of the convex max-margin problem (14).

Theorem 3 The KKT point of the non-convex max-margin problem (13) (in terms of B-subdifferential) corresponds to a KKT point of the convex max-margin problem (14) if v is dual feasible, i.e.,

$$\max_{u:\|u\|_2 \leq 1} |v^T (Xu)_+| \leq 1. \tag{16}$$

This condition is equivalent to for all $D_j \in \mathcal{D}$, the dual variable v satisfies that

$$\max_{\|u\|_2 \leq 1; (2D_j - I)Xu \geq 0} |v^T D_j X u| \leq 1. \tag{17}$$

3.1 DUAL FEASIBILITY OF THE DUAL VARIABLE

A natural question arises: is it possible to examine whether v is feasible in the dual problem? We say the dataset (X, y) is orthogonal separable if for all $n, n' \in [N]$,

$$\begin{aligned} x_n^T x_{n'} &> 0, \text{ if } y_n = y_{n'}, \\ x_n^T x_{n'} &\leq 0, \text{ if } y_n \neq y_{n'}. \end{aligned}$$

For orthogonal separable data, as long as the induced diagonal matrices in Proposition 3 cover the positive part and the negative part of the labels, the KKT point of the non-convex max-margin problem (2) is the KKT point of the convex max-margin problem (14).

Proposition 4 Suppose that (X, y) is orthogonal separable. Suppose that the KKT point of the non-convex problem include two neurons $(w_{1;i_+}, w_{2;i_+})$ and $(w_{1;i_-}, w_{2;i_-})$ such that the corresponding diagonal matrices $\hat{D}_{i_+}$ and $\hat{D}_{i_-}$ defined in Proposition 3 satisfy that

$$\hat{D}_{i_+} = \operatorname{diag}(\mathbb{I}(y = 1)), \quad \hat{D}_{i_-} = \operatorname{diag}(\mathbb{I}(y = -1)).$$

Then, the dual variable is dual feasible, i.e., satisfying (16).

The spike-free matrices discussed in (Ergen & Pilanci, 2021a) also makes examining the dual feasibility of easier. The definition of spike-free matrices can be found in Appendix A

Proposition 5 Suppose that X is spike-free. Suppose that the KKT point of the non-convex problem include two neurons $(w_{1;i_+}, w_{2;i_+})$ and $(w_{1;i_-}, w_{2;i_-})$ such that the corresponding diagonal matrices $\hat{D}_{i_+}$ and $\hat{D}_{i_-}$ defined in Proposition 3 satisfy that

$$\hat{D}_{i_+} = \operatorname{diag}(\mathbb{I}(y = 1)), \quad \hat{D}_{i_-} = \operatorname{diag}(\mathbb{I}(y = -1)).$$

Then, the dual variable is dual feasible, i.e., satisfying (16).

Remark 1 For the spike-free data, the constraint on the dual problem is equivalent to

$$\max_{Xu \ 0; \|u\|_2 \ 1} |v^T Xu| \ 1; \quad \text{or equivalently}$$

$$\max_{Xu \ 0; \|u\|_2 \ 1} v^T Y_+ Xu \ 1; \quad \min_{Xu \ 0} v^T Y_- Xu \ 1.$$

4 Sub-gradient flow dynamics of logistic loss

In this section, we consider the following sub-gradient flow of the logistic loss (3)

$$\begin{aligned} \frac{\partial}{\partial t} w_{1;i}(t) &= w_{2;i}(t) \Big(\sum_{n:(w_{1;i}(t))^T x_n > 0} \tilde{\ell}_n(t) x_n(t) \Big), \\ \frac{\partial}{\partial t} w_{2;i}(t) &= \sum_{n=1}^{N} \tilde{\ell}_n(t) ((w_{1;i}(t))^T x_n(t))_+ . \end{aligned} \tag{18}$$

where the n-th entry of $e(t) \in \mathbb{R}^N$ is defined

$$\tilde{\ell}_n = -y_n \ell'(q_n), \quad q_n = y_n (x_n^T W_1)_+ w_2. \tag{19}$$

For simplicity, we omit the term (t). For instance, we write $w_{1;i} = w_{1;i}(t)$. To be specific, when $w_{1;i}^T x_n = 0$, we select 0 as the subgradient of $w_{2;i}(w_{1;i}^T x_n)_+$ with respect to $w_{1;i}$. Denote $_i = \operatorname{sign}(Xu_i)$. For $\in \{-1, 1, 0\}^N$, we define

$$g(\ ; e) = \sum_{n: \ _n > 0} \tilde{\ell}_n x_n. \tag{20}$$

For simplicity, we also write

$$g(u; e) := g(\operatorname{sign}(Xu); \ ) = \sum_{n: w_{1;i}^T x_n > 0} \tilde{\ell}_n x_n. \tag{21}$$

Then, we can rewrite sub-gradient flow of the logistic loss (3) as follows:

$$\frac{\partial}{\partial t} w_{i;1} = w_{2;i} g(u; e), \quad \frac{\partial}{\partial t} w_{i;2} = w_{i;1}^T g(u; e). \tag{22}$$

Assume that the neural network is scaled at initialization, i.e., $\|w_{1;i}(0)\|_2^2 = w_{2;i}^2(0)$ for $i \in [m]$. Then, the neural network is scaled for 0.

Lemma 1 Suppose that $\|w_{1,i}(0)\|_2 = |w_{2,i}(0)| > 0$ for $i \in [m]$. Then, for any $t > 0$, we have $\|w_{1,i}(t)\|_2 = |w_{2,i}(t)| > 0$.

According to Lemma 1, for all $t \geq 0$, $\mathrm{sign}(w_{2,i}(t)) = \mathrm{sign}(w_{2,i}(0))$. Therefore, we can simply write $s_i = s_i(t) = \mathrm{sign}(w_{2,i}(t))$. As the neural network is scaled for $t \geq 0$, it is interesting to study the dynamics of $w_{1,i}$ in the polar coordinate. We write $w_{1,i}(t) = e^{r_i(t)} u_i(t)$, where $\|u_i(t)\|_2 = 1$. The gradient flow in terms of polar coordinate writes

$$\frac{\partial}{\partial t} r_i = s_i u_i^T g(u_i; \theta); \quad \frac{\partial}{\partial t} u_i = s_i \left(I_d - u_i u_i^T\right) g(u_i; \theta) \qquad (23)$$

Let $x_{\max} = \max_{i \in [n]} \|x_i\|_2$. Define $g_{\min}$ to be

$$g_{\min} = \min_{\theta \in Q} \|g(\cdot; \theta/4)\|_2, \text{ s.t. } g(\cdot; \theta/4) \neq 0, \quad \text{where we denote} \qquad (24)$$

$$Q = \{\theta \in \{-1, 0, 1\}^N, \theta = \mathrm{sign}(Xw), w \in \mathbb{R}^d\}. \qquad (25)$$

As the set $Q \subseteq \{-1, 1, 0\}^N$ is finite, we note that $g_{\min} > 0$. We note that when $\max_{n \in [N]} |q_n| \to 0$, we have $\theta \to \frac{\gamma}{4}$. The following lemma shows that with initializations sufficiently close to 0, $\|g(u(t); \theta(t)) - g(u(t); \gamma/4)\|_2$ and $\left\|\frac{d}{dt} g(u(t); \theta(t))\right\|_2$ can be very small.

Lemma 2 Suppose that $T > 0$ and $\varepsilon > 0$. Suppose that $(u(t); r(t))$ follows the gradient flow (23) with $s = 1$ and the initialization $u(0) = u_0$ and $r(0) = r_0$. Suppose that α_0 is sufficiently small. Then, the following two statements hold.

- For all $t \leq T$, we have $\|g(u(t); \theta(t)) - g(u(t); \gamma/4)\|_2 \leq \frac{g_{\min}}{8}$.
- For $t \leq T$ such that $\mathrm{sign}(Xu(s))$ is constant in a small neighborhood of t, we have $\left\|\frac{d}{dt} g(u(t); \theta(t))\right\|_2 \leq \frac{g_{\min}^2}{16}$.

Based on the above lemma on the property of $g(u(t); \theta(t))$, we introduce the following lemma to upper-bound the time such that $\cos\angle(u(t); g(u(t); \theta(t)))$ approaches 1 or $\mathrm{sign}(Xu(t))$ changes.

Lemma 3 Let $\varepsilon \in (0, 1)$. Suppose that u_0 satisfies that $\|u_0\|_2 = 1$ and $\theta(0)^T (Xu_0)_+ > 0$. Suppose that $(u(t); r(t))$ follows the gradient flow (23) with $s = 1$ and the initialization $u(0) = u_0$ and $r(0) = r_0$. Let $v(t) = \frac{g(u(t); \theta(t))}{\|g(u(t); \theta(t))\|_2}$. We write $v_0 = v(0)$, $\theta_0 = \theta(0)$ and $g_0 = \|g(\theta_0; \gamma/4)\|_2$. Denote

$$T = \frac{1}{2g_0 \sqrt{1 - \varepsilon/8}} \log\left(\frac{\sqrt{1 - \varepsilon/8} + 1}{1 - \varepsilon/8 + \varepsilon}\right) \log\left(\frac{\sqrt{1 - \varepsilon/8} + v_0^T u_0}{\sqrt{1 - \varepsilon/8} - v_0^T u_0}\right). \qquad (26)$$

For $c \in (0, 1 - \varepsilon]$, define

$$T^{\mathrm{shift}}(c) = \frac{1}{2g_0 \sqrt{1 - \varepsilon/8}} \log\left(\frac{\sqrt{1 - \varepsilon/8} + c}{\sqrt{1 - \varepsilon/8} - c}\right) \log\left(\frac{\sqrt{1 - \varepsilon/8} + v_0^T u_0}{\sqrt{1 - \varepsilon/8} - v_0^T u_0}\right) \qquad (27)$$

Suppose that α_0 is sufficiently small such that the statements in Lemma 2 holds for T. Then, at least one of the following event happens

- There exists a time T such that we have $\mathrm{sign}(Xu(t)) = \mathrm{sign}(Xu_0)$ for $t \in [0, T)$ and $\mathrm{sign}(Xu(t)) \neq \mathrm{sign}(Xu_0)$. Let $u_1 = u(T)$ and $v_1 = \lim_{t \to T - 0} v(t)$. If $u_1^T v_1 \leq 1 - \varepsilon$, then the time T satisfies that $T \geq T^{\mathrm{shift}}(v_1^T u_1)$. Otherwise, there exists a time T^0 satisfying $T^0 \leq T$, such that we have $\mathrm{sign}(Xu(t)) = \mathrm{sign}(Xu_0)$ for $t \in [0, T^0]$ and $u(T^0)^T v(T^0) \geq 1 - \varepsilon$.
- There exists a time $T \leq T$, such that we have $\mathrm{sign}(Xu(t)) = \mathrm{sign}(Xu_0)$ for $t \in [0, T]$ and $u(T)^T v(T) \geq 1 - \varepsilon$.

Corollary 1 Suppose that there exists a time T such that we have $\text{sign}(Xu(t)) = \text{sign}(Xu_0)$ for $t \in [0, T)$ and $\text{sign}(Xu(t)) \neq \text{sign}(Xu_0)$. If $T > T^{\text{shift}}(v_1^T u_1) = \frac{1}{g_0\sqrt{1-\delta^2}} \log\frac{\sqrt{1-\delta^2}+v_1^T u_1}{1-\delta^2 v_1^T u_1} - \log\frac{\sqrt{1-\delta^2}+v_0^T u_0}{1-\delta^2 v_0^T u_0}$, then, we have $u_1^T v_1 > 1-\delta$.

Proposition 6 Consider the sub-gradient flow (23) with $s = 1$ and the initialization $u(0) = u_0$ and $r(0) = r_0$. Here at initilization the neuron u_0 satisfies that $\|u_0\|_2 = 1$ and $y^T(Xu_0)_+ > 0$. Let $v(t) = \frac{g(u(t);\theta(t))}{\|g(u(t);\theta(t))\|_2}$. For any $\delta > 0$, for sufficiently small r_0, there exists a time $T = O(\log(\delta^{-1}))$ such that $u(T)^T v(T) \geq 1-\delta$ and $\cos\angle(u(T); g(u(T); y)) \geq 1-\delta$.

Remark 2 The statement of proposition is similar to Lemma 4 in (Maennel et al., 2018). However, their proof contains a problem because they did not consider the change of $\text{sign}(Xw)$ along the gradient flow. Our proof in Appendix D.4 corrects this error.

We next study the properties of orthogonal separable datasets. Denote $B = \{w \in \mathbb{R}^d : \|w\|_2 \leq 1\}$. The following lemma give a sufficient condition on w to satisfy the condition in Proposition 4.

Lemma 4 Assume that (X, y) is orthogonal separable. Suppose that $w \in B$ is a local maximizer of $y^T(Xw)_+$ in B and $(Xw)_+ \neq 0$. Then, $\langle w, x_n\rangle > 0$ for $n \in [N]$ such that $y_n = 1$. Suppose that $w \in B$ is a local minimizer of $y^T(Xw)_+$ in B and $(Xw)_+ \neq 0$. Then, $\langle w, x_n\rangle > 0$ for $n \in [N]$ such that $y_n = -1$.

We show an equivalent condition of $u \in B$ being the local maximizer/minimizer of $y^T(Xu)_+$ in B.

Proposition 7 Assume that (X, y) is orthogonal separable. Then, $u \in B$ is a local maximizer of $y^T(Xu)_+$ in B is equivalent to $\cos\angle(u; g(u; y)) = 1$. Similarly, $u \in B$ is a local minimizer of $y^T(Xu)_+$ in B is equivalent to $\cos\angle(u; g(u; y)) = -1$.

Based on Proposition 4 and 7, we present the main theorem.

Theorem 4 Suppose that the dataset is orthogonal separable and $\theta(t)$ follows the gradient flow. Suppose that the neural network is scaled at initialization, $\|w_{1;i}(0)\|_2 = |w_{2;i}(0)|$ for all $i \in [m]$. For almost all initializations which are sufficiently close to zero, the limiting point of $\frac{\theta(t)}{\|\theta(t)\|_2}$ is $\frac{\theta^*}{\|\theta^*\|_2}$, where θ^* is a global minimizer of the max-margin problem (2).

We present a sketch of the proof. According to Proposition 6, for initialization sufficiently close to zero, there exist two neurons and time $T_+, T_- > 0$ such that $\cos\angle(w_{1;i_+}(T_+); g(w_{1;i_+}(T_+); y)) \geq 1-\delta$ and $\cos\angle(w_{1;i_-}(T_-); g(w_{1;i_-}(T_-); y)) \geq (1-\delta)$. This implies that $w_{1;i_+}(T_+)$ and $w_{1;i_-}(T_-)$ are sufficiently close to certain stationary points of gradient flow maximizing/minimizing $y^T(Xu_+)$ over B, i.e., $\{u \in B | \cos\angle(u; g(u; y)) = 1\}$. As the dataset is orthogonal separable, from Proposition 7 and Lemma 4, the induced masking matrices $\hat{D}_{i_+}(T_+)$ and $\hat{D}_{i_-}(T_-)$ by $w_{1;i_+}(T_+)/w_{1;i_-}(T_-)$ in Proposition 3 satisfy that $\hat{D}_{i_+}(T_+) = \text{diag}(I(y = 1))$ and $\hat{D}_{i_-}(T_-) = \text{diag}(I(y = -1))$. According to Lemma 3 in (Phuong & Lampert, 2021), for $\max\{T_+, T_-\}$, we also have $\hat{D}_{i_+}(t) = \text{diag}(I(y = 1))$ and $\hat{D}_{i_-}(t) = \text{diag}(I(y = -1))$. According to Theorem 3 and Proposition 4, the KKT point of the non-convex problem (2) that gradient flow converges to corresponds to the KKT point of the convex problem (14).

5 Conclusion

We provide a convex formulation of the non-convex max-margin problem for two-layer ReLU neural networks and uncover a primal-dual extreme point relation between non-convex subgradient flow. Under the assumptions on the training data, we show that flows converge to KKT points of the convex max-margin problem, hence a global optimum of the non-convex objective.

6 Acknowledgements

This work was partially supported by the National Science Foundation under grants ECCS-2037304, DMS-2134248, and US Army Research Office.

REFERENCES

Burak Bartan and Mert Pilanci. Neural spectrahedra and semide nite lifts: Global convex optimization of polynomial activation neural networks in fully polynomial-time. arXiv preprint arXiv:2101.02429, 2021a.

Burak Bartan and Mert Pilanci. Training quantized neural networks to global optimality via semide - nite programming. International Conference on Machine Learning (ICML), 2021, 2021b.

Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. Advances in Neural Information Processing Systems, 31: 3036–3046, 2018.

Thomas M Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. IEEE transactions on electronic computers, (3):326–334, 1965.

Tolga Ergen and Mert Pilanci. Implicit convex regularizers of cnn architectures: Convex optimization of two-and three-layer networks in polynomial time. International Conference on Learning Representations (ICLR), 2021, 2020.

Tolga Ergen and Mert Pilanci. Convex geometry and duality of over-parameterized neural networks. Journal of Machine Learning Research, 22(212):1–63, 2021a.

Tolga Ergen and Mert Pilanci. Global optimality beyond two layers: Training deep relu networks via convex programs. In International Conference on Machine Learning, pp. 2993–3003. PMLR, 2021b.

Tolga Ergen and Mert Pilanci. Revealing the structure of deep neural networks via convex duality. In International Conference on Machine Learning, pp. 3004–3014. PMLR, 2021c.

Tolga Ergen, Arda Sahiner, Batu Ozturkler, John Pauly, Morteza Mardani, and Mert Pilanci. Demystifying batch normalization in relu networks: Equivalent convex optimization models and implicit regularization. arXiv preprint arXiv:2103.01499, 2021.

Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In International Conference on Machine Learning, pp. 1832–1841. PMLR, 2018.

Vikul Gupta, Burak Bartan, Tolga Ergen, and Mert Pilanci. Exact and relaxed convex formulations for shallow neural autoregressive models. In International Conference on Acoustics, Speech, and Signal Processing, 2021.

Yuanzhi Li, Tengyu Ma, and Hongyang R Zhang. Learning over-parametrized two-layer neural networks beyond ntk. In Conference on Learning Theory, pp. 2613–2682. PMLR, 2020.

Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. arXiv preprint arXiv:1906.05890, 2019.

Chao Ma, Lei Wu, and Weinan E. The quenching-activation behavior of the gradient descent dynamics for two-layer neural network models. arXiv preprint arXiv:2006.14450, 2020.

Hartmut Maennel, Olivier Bousquet, and Sylvain Gelly. Gradient descent quantizes relu network features. arXiv preprint arXiv:1803.08367, 2018.

Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean eld view of the landscape of two-layer neural networks. Proceedings of the National Academy of Sciences, 115(33):E7665–E7671, 2018.

Mary Phuong and Christoph H Lampert. The inductive bias of relu networks on orthogonally separable data. 2021.

Mert Pilanci and Tolga Ergen. Neural networks are convex regularizers: Exact polynomial-time convex optimization formulations for two-layer networks. pp. 7695–7705, 2020.

Arda Sahiner, Tolga Ergen, Batu Ozturkler, Burak Bartan, John Pauly, Morteza Mardani, and Mert Pilanci. Hidden convexity of wasserstein gans: Interpretable generative models with closed-form solutions.arXiv preprint arXiv:2107.05680, 2021a.

Arda Sahiner, Tolga Ergen, John Pauly, and Mert Pilanci. Vector-output relu neural network problems are copositive programs: Convex analysis of two layer networks and polynomial-time algorithms. International Conference on Learning Representations (ICLR), 2021b.

Arda Sahiner, Morteza Mardani, Batu Ozturkler, Mert Pilanci, and John Pauly. Convex regularization behind neural reconstruction. International Conference on Learning Representations (ICLR), 2021c.

Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. The Journal of Machine Learning Research, 19(1): 2822–2878, 2018.

Gang Wang, Georgios B Giannakis, and Jie Chen. Learning relu networks on linearly separable data: Algorithm, optimality, and generalization. IEEE Transactions on Signal Processing, 67(9): 2357–2370, 2019.

Qiuling Yang, Alireza Sadeghi, Gang Wang, and Jian Sun. Learning two-layer relu networks is nearly as easy as learning linear classi ers on separable data. IEEE Transactions on Signal Processing, 69:4416–4427, 2021.

Mo Zhou, Rong Ge, and Chi Jin. A local convergence theory for mildly over-parameterized two-layer neural network. arXiv preprint arXiv:2102.02410, 2021.

A DEFINITIONS AND NOTIONS

We introduce several useful definitions and notions which will be utilized in the proof.

A.1 DEFINITIONS

Definition 1 Let $O \subseteq R^n$ be an open set and let $F : O \to R$ be locally Lipschitz continuous at $x \in O$. Let D_F be the differentiable points of F in O. The **B**-subdifferential of F at x is defined by

$$\partial^B F(x) := \lim_{k \to \infty} \{ F'(x^k) | x^k \in D_F, x_k \to x \}. \tag{28}$$

The set $\partial F(x) = co(\partial^B F(x))$ is called Clarke's subdifferential, where co denotes the convex hull.

Definition 2 A matrix A is spike-free if and only if the following conditions hold: for all $\|u\|_2 = 1$, there exists $\|z\|_2 = 1$ such that

$$(Au)_+ = Az. \tag{29}$$

This is equivalent to say that

$$\max_{u: \|u\|_2 = 1, (I - XX^\dagger)(Xu)_+ = 0} \|X^\dagger (Xu)_+\|_2 = 1. \tag{30}$$

A.2 NOTIONS

We use the following letters for indexing.

- The index n is for the n-th data sample x_n.
- We use the index i to represent the i-th neuron-pair $(w_{1,i}, w_{2,i})$.
- The index j is for the j-th masking matrix $D_i \in \mathcal{P}$.

B PROOFS IN SECTION 3

B.1 PROOF FOR PROPOSITION 2

Consider the following loss function $\ell : R^N \times R^N \to R \cup \{+\infty\}$

$$\ell(z, y) = \begin{cases} 0, & y_n z_n \geq 1, \forall n \in [N], \\ +\infty, & \text{otherwise} \end{cases} \tag{31}$$

For a given $y \in \{+1, -1\}^N$, $\ell(z, y)$ is a convex loss function of z. The non-convex max-margin is equivalent to

$$\min \ell((XW_1)_+ w_2, y) + \frac{1}{2} \|W_1\|_F^2 + \|w_2\|_2^2. \tag{32}$$

According to Appendix A.13 in (Pilanci & Ergen, 2020), the problem (32) is equivalent to

$$\begin{aligned} \min \ell\Big(\sum_{i=1}^p D_i X (u_i^+ - u_i^-), y \Big) + \frac{1}{2} \|W_1\|_F^2 + \|w_2\|_2^2, \\ \text{s.t.} (2D_i - I) X u_i \geq 0, (2D_i - I) X u_i^- \geq 0, \forall i \in [p]. \end{aligned} \tag{33}$$

This is equivalent to (14). For fixed $y \in \{+1, -1\}^N$, the Fenchel conjugate function $\ell^*(\hat{z}, y)$ with respect to z can be computed by

$$\begin{aligned} \ell^*(\hat{\ }, y) &= \max_{z \in R^N} z^T \hat{\ } - \ell(z, y) \\ &= \max_{z \in R^N} z^T \hat{\ }, \text{ s.t. } \operatorname{diag}(y) z \geq 1, \\ &= \begin{cases} y^T \hat{\ }, & \operatorname{diag}(y) \hat{\ } \geq 0 \\ +\infty, & \text{otherwise} \end{cases} \end{aligned} \tag{34}$$

According to Theorem 6 in (Pilanci & Ergen, 2020), the dual problem of (14) writes

$$\max \quad \tilde{r}(\ ; y);\ \text{s.t.} \max_{u:\|u\|_2 \ 1} |\ ^T(Xu)_+| \quad 1; \tag{35}$$

which is equivalent to

$$\max \quad y^T \ ;\ \text{s.t.}\ \text{diag}(y) \quad 0;\ \max_{u:\|u\|_2 \ 1} |\ ^T(Xu)_+| \quad 1: \tag{36}$$

By taking $=\ \hat{}$, we derive (15). This completes the proof.

B.2 Proof for Proposition 3

For the non-convex max-margin problem (13), consider the Lagrange function

$$L(W_1; w_2; \ ) = \frac{1}{2}(\|W_1\|_F^2 + \|w_2\|_2^2) \quad (Y \ )^T(Y(XW_1)_+ w_2 \quad 1)$$

where $Y \quad 0$. The KKT point of the non-convex max-margin problem (13) (in terms of B-subdifferential) satisfies

$$\begin{aligned} 0 &2 \ @^B_{W_1} L(W_1; w_2; \ ); \\ w_2 & (XW_1)_+^T \ = 0; \\ _n(y_n(x_n^T W_1)_+ w_2 & 1) = 0: \end{aligned} \tag{37}$$

The KKT condition on the i-th column of W_1 is equivalent to

$$w_{1;i} = \sum_{n=1}^N w_{2;i} \ _n x_n g_{n;i}; \tag{38}$$

where $g_{n;i} \ 2 \ @^B(z)_+|_{z=x_n^T w_{1;i}}$. In other words, we have

$$g_{n;i} \begin{cases} = I(x_n^T w_{1;i} \quad 0); & \text{if } x_n^T w_{1;i} \neq 0; \\ 2 \{0; 1\}; & \text{if } x_n^T w_{1;i} = 0: \end{cases} \tag{39}$$

Let $\hat{D}_i = \text{diag}([g_{1;i}; \cdots; g_{N;i}])$. Then, we can write that

$$\begin{aligned} w_{1;i} &= \sum_{n=1}^N \ _n g_{n;i} x_n w_{2;i} \\ &= w_{2;i} X^T \hat{D}_i \ : \end{aligned} \tag{40}$$

From the definition of $g_{n;i}$, we have

$$g_{n;i} \ x_n^T w_{1;i} \ = 0: \tag{41}$$

Therefore, we can compute that

$$\begin{aligned} w_{2;i} &= (Xw_{1;i})_+^T \\ &= \sum_{n=1}^N I(x_n^T w_{1;i} \quad 0) x_n^T w_{1;i} \ _n \\ &= \sum_{n=1}^N g_{n;i} \ x_n^T w_{1;i} \ _n \\ &= w_{1;i}^T X^T \hat{D}_i \ : \end{aligned} \tag{42}$$

In summary, we have

$$w_{1;i} \ = \ w_{2;i} X^T \hat{D}_i \ ; \quad w_{2;i} \ = \ w_{1;i}^T X^T \hat{D}_i \ : \tag{43}$$

Suppose that $w_{2;i} \neq 0$. This implies that

$$\frac{w_{1;i}}{w_{2;i}} \ = \ X^T \hat{D}_i \ ; \quad \|X^T \hat{D}_i \ \|_2 = 1: \tag{44}$$

This completes the proof.

B.3 PROOF FOR THEOREM 3

PROOF We can write the Lagrange function for the convex max-margin problem (14) as

$$\begin{aligned}
&L(\{u_j\}_{j=1}^p; \{u_j^0\}_{j=1}^p; \; ; \{z_j\}_{j=1}^p; \{z_j^0\}_{j=1}^p) \\
&= \sum_{j=1}^p (\|u_j\|_2 + \|u_j^0\|_2) + \lambda^T \mathrm{diag}(y) \Big(1 \quad \mathrm{diag}(y) \sum_{j=1}^p D_j X (u_j \quad u_j^0) \Big) \\
&\quad \sum_{j=1}^p (z_j^T (2D_j \quad I) X u_j + (z_j^0)^T (2D_j \quad I) X u_j^0) \\
&= \lambda^T y + \sum_{j=1}^p (\|u_j\|_2 + \|u_j^0\|_2) + \sum_{j=1}^p (u_j^0)^T (X^T D_j \quad X^T (2D_j \quad I) z_j^0) \\
&\quad + \sum_{j=1}^p (u_j)^T (\quad X^T D_j \quad X^T (2D_j \quad I) z_j):
\end{aligned} \tag{45}$$

where $z_j, z_j^0 \in \mathbb{R}^N$ satisfies that $z_j \geq 0, z_j^0 \geq 0$ for $j \in [p]$ and $\lambda \in \mathbb{R}^N$ satisfies that $\mathrm{diag}(y) \geq 0$. The KKT point shall satisfy the following KKT conditions:

$$\begin{aligned}
& X^T D_j + X^T (2D_j \quad I) z_j^0 \in \partial_{u_j^0} \|u_j^0\|_2, \\
& X^T D_j + X^T (2D_j \quad I) z_j \in \partial_{u_j} \|u_j\|_2, \\
& \Big(\sum_{j=1}^p (D_j)_n x_n^T (u_j \quad u_j^0) \quad y_n \Big) = 0, \\
& z_{j,n} (2(D_j)_{n,n} \quad 1) x_n^T u_j = 0, \\
& z_{j,n}^0 (2(D_j)_{n,n} \quad 1) x_n^T u_j^0 = 0.
\end{aligned} \tag{46}$$

Let $(W_1, w_2, \;)$ be the KKT point of the non-convex problem (2) and λ satisfies (17). Let $\hat{D}_i$ be the diagonal matrix defined in Proposition 3 with respect to $w_{1,i}$ and denote $\mathcal{P} = \{\hat{D}_i | i \in [m]\}$. Without the loss of generality, we may assume $\{\hat{D}_i\}_{i=1}^m$ are different. (Otherwise, we can merge two neurons w_{1,i_1} and w_{1,i_2} with $D_{i_1} = D_{i_2}$ together.)

Suppose that $D_j \in \mathcal{P}$, i.e., $D_j = \hat{D}_i$ for certain $i \in [m]$. By letting $u_j^0 = w_{1,i} w_{2,i}$, $z_j^0 = 0$, $u_j = \quad w_{1,i} w_{2,i}$ and $z_j = 0$, the following identities hold.

$$X^T D_j + X^T (2D_i \quad I) z_j^0 = X^T \hat{D}_i = \frac{w_{1,i}}{w_{2,i}} = \frac{u_i^0}{\|u_i^0\|}. \tag{47}$$

$$X^T D_j + X^T (2D_i \quad I) z_j = \quad X^T \hat{D}_i = \frac{w_{1,i}}{w_{2,i}} = \frac{u_i}{\|u_i\|}. \tag{48}$$

Therefore, for index j satisfying $D_i \in \mathcal{P}$, the first two KKT conditions in (46) hold.

For $D_j \neq \mathcal{P}$, we can let $u_j = u_j^0 = 0$. As λ satisfies (17), we have

$$\max_{\|u\|_2 \leq 1, (2D_j \quad I) X u \geq 0} \quad u^T D_j X u \quad 1. \tag{49}$$

According to Lemma 4 in (Pilanci & Ergen, 2020), this implies that there exists z_j, z_j^0 such that

$$\| \quad X^T D_j + Z^T (2D_j \quad I) z_j^0 \| \leq 1, \| X^T D_j + Z^T (2D_j \quad I) z_j \| \leq 1. \tag{50}$$

Therefore, the first two KKT conditions in (46) hold.

From our choice of u_j, z_j, u_j^0, z_j, the last two KKT conditions in (46) hold. We also note that

$$\sum_{j=1}^p D_j X (u_j^0 \quad u_j) = \sum_{i=1}^m (X w_{1,i})_+ w_{2,i}. \tag{51}$$

As $(W_1, w_2, \;)$ is the KKT point of the non-convex problem, the third KKT condition (46) holds. This completes the proof.

C PROOFS IN SECTION 3.1

In this section, we present several proofs for propositions in Section 3.1.

C.1 PROOF FOR PROPOSITION 4

We start with two lemmas.

Lemma 5 Suppose that $u_0 = X^T \hat{D}_0$ and $\|u_0\|_2 \leq 1$. For any masking matrix $D_j \in \mathcal{P}$ such that $(D_j - \hat{D}_0) \mathbb{I}(\beta > 0) = 0$, we have

$$\max_{(2D_j - I)Xu \geq 0; \|u\|_2 \leq 1} u^T D_j X u \leq 1. \tag{52}$$

PROOF According to Lemma 4 in (Pilanci & Ergen, 2020), the constraint (52) is equivalent to that there exist $z_j \in \mathbb{R}^N$ such that $z_j \geq 0$ and

$$\|X^T D_j u + X^T (2D_j - I) z_j\| \leq 1. \tag{53}$$

Consider the index $n \in [N]$ such that $(D_j - \hat{D}_0)_{nn} \neq 0$. As $(D_j - \hat{D}_0) \mathbb{I}(\beta > 0) = 0$, we have $\beta_n \leq 0$. We let $(z_j)_n = -\beta_n$. If $(\hat{D}_0)_{nn} = 0$, then we have $(D_j)_{nn} = 1$ and

$$(D_j - \hat{D}_0)_{nn} \beta_n x_n = -\beta_n x_n = -x_n^T (2(D_j)_{nn} - 1)(z_j)_n. \tag{54}$$

If $(\hat{D}_0)_{nn} = 1$, then we have $(D_j)_{nn} = 0$ and

$$(D_j - \hat{D}_0)_{nn} \beta_n x_n = \beta_n x_n = -x_n^T (2(D_j)_{nn} - 1)(z_j)_n. \tag{55}$$

For other index $n \in [N]$, we simply let $(z_j)_n = 0$. Then, we have

$$(D_j - \hat{D}_0)_{nn} \beta_n x_n = 0 = -x_n^T (2(D_j)_{nn} - 1)(z_j)_n. \tag{56}$$

Based on our choice of z_j, we have $z_j \geq 0$ and for $n \in [N]$

$$(D_j - \hat{D}_0)_{nn} \beta_n x_n = -x_n^T (2(D_j)_{nn} - 1)(z_j)_n. \tag{57}$$

This implies that

$$X^T (D_j - \hat{D}_0) \beta = -X^T (2D_j - I) z_j. \tag{58}$$

Hence, we have

$$X^T D_j u + X^T (2D_j - I) z_j = X^T \hat{D}_0 u = u_0. \tag{59}$$

Therefore, $\|X^T D_j u + X^T (2D_j - I) z_j\|_2 \leq 1$.

Lemma 6 Suppose that the data is orthogonal separable and $\beta \neq 0$. Suppose that $u_0 = X^T \hat{D}_0 u$ and $\|u_0\|_2 \leq 1$. For any masking matrix D_j such that $\hat{D}_0 D_j = 0$, we have $\|X^T D_j u\|_2 \leq \|u_0\|_2 \leq 1$. Therefore, (52) holds.

PROOF We note that $u_0 = X^T (\hat{D}_0 - D_j) u + X D_j u$. Denote $a = X^T (\hat{D}_0 - D_j) u$ and $b = X^T D_j u$. We note that

$$a^T b = \Big(\sum_{n : (\hat{D}_0)_{nn} = 1 ; (D_j)_{n,n} = 0} \beta_n x_n \Big)^T \Big(\sum_{n' : (\hat{D}_0)_{n'n'} = 0 ; (D_j)_{n'n'} = 0} \beta_{n'} x_{n'} \Big). \tag{60}$$

As $\mathrm{diag}(y) \geq 0$, β_n has the same signature with y_n. Therefore, from the orthogonal separability of the data, we have

$$\beta_n \beta_{n'} x_n^T x_{n'} \geq 0. \tag{61}$$

This immediately implies that $a^T b \geq 0$. Therefore,

$$1 \geq \|u_0\|_2^2 = \|a + b\|_2^2 = \|a\|_2^2 + \|b\|_2^2 + 2 a^T b \geq \|a\|_2^2. \tag{62}$$

This completes the proof.

Based on Lemma 5 and Lemma 6, we present the proof for Proposition 3. Let $i_+ = \frac{w_{1,i_+}}{w_{2,i_+}}$. From the proof of Proposition 3, we note that $\|u_+\|_2 = 1$. For any masking matrix $D_j \in \mathcal{P}$, let $D' = \hat{D}_{i_+} D_j$. As $\hat{D}_{i_+} \geq D'$, according to Lemma 6, we have

$$\|X^T D' u\|_2 \leq \|X^T \hat{D}_{i_+} u\|_2 = \|u_+\|_2 \leq 1. \tag{63}$$

As $Y \geq 0$ and $\hat{D}_{i_+} = \mathrm{diag}(\mathbb{I}(y = 1))$, we have $(D_j - D') \mathbb{I}(\beta > 0) = D_j (I - \hat{D}_{i_+}) \mathbb{I}(\beta > 0) = 0$. From Lemma 5, we note that u satisfies (52). Similarly, we can show that $-u$ also satisfies (52). This completes the proof.

C.2 PROOF FOR PROPOSITION 5

PROOF Note that $Y_- \neq 0$. Let $Y_+ = \mathrm{diag}(\mathbb{I}(y = 1))$ and $Y_- = \mathrm{diag}(\mathbb{I}(y = -1))$. We claim that

$$\max_{\|u\|_2 \leq 1} \|(Xu)_+\|_2^2 = \max_{\|u\|_2 \leq 1} u^T Y_+ X^T (Xu)_+ . \tag{64}$$

Firstly, we note that

$$\|(Xu)_+\|_2^2 = \sum_{n=1}^N \mathbb{I}(x_n^T u)_+ \sum_{n=1}^N (x_n^T u)_+ \geq \sum_{n=1}^N \mathbb{I}(x_n^T u)_+ (x_n^T u)_+ = u^T Y_+ X^T (Xu)_+ . \tag{65}$$

This implies that $\max_{\|u\|_2 \leq 1} \|(Xu)_+\|_2^2 \geq \max_{\|u\|_2 \leq 1} u^T Y_+ X^T (Xu)_+$.

On the other hand, suppose that $u^* \in \arg\max_{\|u\|_2 \leq 1} u^T Y_+ X^T (Xu)_+$. As X is spike-free, there exists z such that $\|z\|_2 \leq 1$ and $Xz = (Xu^*)_+$. Therefore, we have

$$u^{*T} Y_+ X^T (Xu^*)_+ = u^{*T} Y_+ X^T Xz = u^{*T} X^T Xz = \|(Xz)_+\|_2^2 . \tag{66}$$

This implies that $\max_{\|u\|_2 \leq 1} \|(Xu)_+\|_2^2 \leq \max_{\|u\|_2 \leq 1} u^T Y_+ X^T (Xu)_+$.

For any $D_j \in \mathcal{P}$ with $D_j \neq Y_+$. We note that

$$u^T (Xu)_+ \geq u^T D_j (Xu)_+ \geq u^T Y_+ (Xu)_+ . \tag{67}$$

Combining with (65), this implies that $\max_{\|u\|_2 \leq 1} \|(Xu)_+\|_2^2 = \max_{\|u\|_2 \leq 1} u^T D (Xu)_+$.

Let us go back to the original problem. Let $\hat{D}_{i+} = \frac{w_{1,i+}}{w_{2,i+}}$. We note that $(Xw_+) = \hat{D}_{i+} Xw_+ = \hat{D}_{i+} XX^T \hat{D}_{i+}$. Therefore, we have

$$\|(Xw_+)\|_2^2 = u^T \hat{D}_{i+} XX^T \hat{D}_{i+} u = \|X^T \hat{D}_{i+} u\|_2^2 = \|u_+\|_2^2 = 1 . \tag{68}$$

Thus, for any $\|u\|_2 \leq 1$, suppose that $(Xu)_+ = Xz$, where $\|z\|_2 \leq 1$. Then, we have

$$u^T \hat{D}_{i+} (Xu)_+ = u^T \hat{D}_{i+} Xz \leq \|z\|_2 \leq 1 . \tag{69}$$

Therefore, $\max_{\|u\|_2 \leq 1} \|(Xu)_+\|_2^2 = \max_{\|u\|_2 \leq 1} u^T D_+ (Xu)_+ \leq 1$. Similarly, we have

$$\min_{\|u\|_2 \leq 1} \|(Xu)_+\|_2^2 = \min_{\|u\|_2 \leq 1} u^T D (Xu)_+ \geq 1 .$$

This completes the proof.

D PROOFS IN SECTION 4

D.1 PROOF FOR LEMMA 1

PROOF According to the sub-gradient flow (22), we can compute that

$$\frac{\partial}{\partial t} \|w_{1,i}\|_2^2 - w_{2,i}^2 = 2 w_{1,i}^T w_{2,i} g(u, e) - 2 w_{2,i} w_{1,i}^T g(u, e) = 0 . \tag{70}$$

Let $T_0 = \sup\{T : \|w_{1,i}(t)\|_2 = |w_{2,i}(t)|, \forall i \in [n], t \in [0, T)\}$. For $t \in [0, T_0)$, as the neural network is scaled, it is sufficient study the dynamics of $w_{1,i}$ in the polar coordinate. Let us write $w_{1,i}(t) = e^{r_i(t)} u_i(t)$, where $\|u_i(t)\|_2 = 1$. Then, in terms of polar coordinate, the projected gradient flow follows

$$\begin{aligned} \frac{\partial}{\partial t} r_i &= \mathrm{sign}(w_{2,i}) u_i^T g(u_i, e), \\ \frac{\partial}{\partial t} u_i &= \mathrm{sign}(w_{2,i}) \big(g(u_i, e) - u_i^T g(u_i, e) u_i \big) . \end{aligned} \tag{71}$$

Without the loss of generality, we may assume that $w_{2,i}(0) \neq 0$ for $i \in [m]$. Denote

$$x_{\max} = \max_{i \in [n]} \|x_i\|_2 . \tag{72}$$

From the definition of e, we have $\|e\|_1 \leq 1/4$. Therefore, we have

$$\frac{\partial}{\partial t} r_i \geq \|g(u_i, e)\|_2 \geq \Big\| \sum_{j: x_j^T u > 0} \tilde{\gamma}_j x_j \Big\|_2 \geq \frac{n x_{\max}}{2}. \tag{73}$$

Therefore, for finite $t > 0$, we have

$$r_i(t) \geq r_i(0) + \frac{n x_{\max} t}{4}, \tag{74}$$

which implies that $|w_{2,i}(t)| > 0$. This implies that $T_0 = +\infty$.

D.2 Proof of Lemma 2

Proof As we have $\|w_{1,i}\|_2 = |w_{2,i}|$, for $n \in [N]$, we can compute that

$$\begin{aligned} |q_n| &= |(x_n^T W_1)_+ w_2| \\ &\leq \sum_{i=1}^m |(x_n^T w_{1,i})_+ w_{2,i}| \\ &= \sum_{i=1}^m \|w_{1,i}\|_2 |(x_n^T w_{1,i})_+| \\ &\leq \sum_{i=1}^m \|w_{1,i}\|_2 |x_n^T w_{1,i}| \\ &\leq \|x_n\|_2 \sum_{i=1}^m \|w_{1,i}\|_2^2. \end{aligned} \tag{75}$$

Note that $\tilde{\gamma}_n = -y_n \ell'(q_n)$ and $\frac{y_n}{4} = -y_n \ell'(0)$. As ℓ' is $\frac{1}{4}$-Lipschitz continuous, we have

$$\tilde{\gamma}_n - y_n/4 \leq \frac{1}{4}|q_n| \leq \frac{\|x_n\|_2}{4} \sum_{i=1}^m \|w_{1,i}\|_2^2. \tag{76}$$

For any $\hat{\theta} \in Q$, as $\tilde{\gamma}_n \in [0, 1/4]$ for $n \in [N]$, we have

$$\begin{aligned} &\|g(\hat{\theta}, e) - g(\hat{\theta}, y/4)\|_2 \\ &\leq \Big\| \sum_{n:(\hat{\theta})_n > 0} (\tilde{\gamma}_n - y_n/4) x_n \Big\|_2 \\ &\leq \sum_{k:(\hat{\theta})_k > 0} |\tilde{\gamma}_k| \|x_k\|_2 \\ &\leq \sum_{k=1}^n \frac{\|x_k\|_2^2}{4} \sum_{j=1}^m \|w_{1,j}\|_2^2 \\ &= c_1 \sum_{i=1}^m \|w_{1,i}\|_2^2, \end{aligned} \tag{77}$$

where $c_1 = \frac{1}{4}\|X\|_F^2 > 0$ is a constant. Therefore, we can bound $\|g(\hat{\theta}, e(t))\|$ by

$$\|g(\hat{\theta}, e(t))\|_2 \leq \|g(\hat{\theta}, y/4)\|_2 + c_1 \sum_{i=1}^m \|w_{1,i}(t)\|_2^2 \leq d_{\max} + c_1 \sum_{i=1}^m e^{2 r_i(t)}, \tag{78}$$

where we let

$$g_{\max} = \max_{\hat{\theta} \in Q} \|g(\hat{\theta}, y/4)\|_2. \tag{79}$$

Let $r(t) = \max_{i \in [m]} r_i(t)$. We note that

$$\frac{@}{@t} r(t) \quad d_{max} + c_1 n e^{2r(t)} \quad c_2(1 + e^{2r(t)}); \tag{80}$$

where $c_2 = \max\{nc_1; d_{max}\} > 0$ is a constant. If we start with $r(0) \quad 0$, then, $r(t)$ cannot grow much faster than $c_2 t$. Let $\tilde{r}(t)$ satisfy the following ODE:

$$\frac{@}{@t} \tilde{r}(t) = c_2(1 + e^{2\tilde{r}(t)}): \tag{81}$$

The solution is given by

$$\tilde{r}_a(t) = c_2(t \quad a) \quad \frac{1}{2}\log(1 \quad e^{2c_2(t \quad a)}); \tag{82}$$

where $a > 0$ is a parameter depending on the initialization. For any initial $r(0)$, we have a unique a satisfying $\tilde{r}_a(0) = r(0)$. Therefore, we have $e(t) \quad \tilde{r}_a(t)$ and

$$kg(\ ; e(t)) \quad g(\ ; y=4)k_2 \quad c_1 n e^{2\tilde{r}_a(t)}: \tag{83}$$

According to the bound (83), by choosing a sufficiently small l_0, (which leads to a sufficiently small a), such that

$$e^{2\tilde{r}_a(T)} \quad \min \quad \frac{d_{min}}{16c_1}; \frac{d_{min}^2}{4n^2 x_{max}^3} \quad : \tag{84}$$

Therefore, for $t \quad T$, we have

$$kg(\hat{}; e(t)) \quad g(\hat{}; y=4)k_2 \quad c_1 n e^{2\tilde{r}_a(t)} \quad c_1 n e^{2\tilde{r}_a(T)} \quad \frac{d_{min}}{8}: \tag{85}$$

Hence, we have

$$kg(\ (t); e(t)) \quad g(\ (t); y=4)k_2 \quad c_1 n e^{2\tilde{r}_a(t)} \quad c_1 n e^{2\tilde{r}_a(T)} \quad \frac{d_{min}}{8}: \tag{86}$$

We can compute that

$$\frac{d}{dt}\tilde{}_i = \quad y_i l^{(2)}(q_i)\frac{d}{dt}q_i: \tag{87}$$

As $l^{(2)}(q) \ 2 \ (0; 1=4]$, we can compute that

$$\begin{aligned} \frac{d}{dt}q_i \quad & \sum_{j=1}^{m} jw_{2;j}\, jkx_i k_2 \quad \frac{d}{dt}w_{1;j} \quad {}^2 \\ & + \sum_{j=1}^{m} kw_{1;j} k_2 kx_i k_2 \quad \frac{d}{dt}w_{2;j} \\ & \frac{n}{4}\sum_{j=1}^{m} kw_{1;j} k_2^2 x_{max}^2 + \frac{n}{4}\sum_{j=1}^{m} w_{2;j}^2 x_{max}^2 \\ & \frac{nx_{max}^2}{2} e^{2r(t)}: \end{aligned} \tag{88}$$

Therefore, we have

$$\frac{d}{dt}\tilde{}_i = j\grave{}^{00}(q_i)j \ \frac{d}{dt}q_i \quad \frac{1}{4}\frac{d}{dt}q_i \quad \frac{nx_{max}^2}{8} e^{2r(t)}: \tag{89}$$

Suppose that $\text{sign}(Xu(s)) = \ (t)$ holds for s in a small neighborhood of t. Then, we have

$$\begin{aligned} \frac{d}{dt} g(u(t); e(t)) \quad {}^2 = \quad & \frac{d}{dt} g(\ (t); e(t)) \quad {}^2 \\ & \sum_{i=1}^{n} kx_i k_2 \ \frac{d}{dt}\tilde{}_i \quad \frac{n^2 x_{max}^3}{8} e^{2r(t)} \\ & \frac{n^2 x_{max}^3}{8} e^{2\tilde{r}_a(T)} \quad \frac{d_{min}^2}{16}: \end{aligned} \tag{90}$$

This completes the proof.

D.3 PROOF OF LEMMA 3

PROOF Let $T_0 = \sup\{T | \text{sign}(Xu(t)) = \text{sign}(Xu_0), \forall t \in [0, T)\}$. We analyze the dynamics of $u(t)$ in the interval $[0, \min\{T_0, T\}]$. For $t \leq \min\{T_0, T\}$, as the statements in Lemma 2 hold, we can compute that

$$\begin{aligned}
&\frac{d}{dt} v(t)^T u(t) \\
&= \frac{d}{dt} \left(\frac{g(\theta_0, e(t))}{\|g(\theta_0, e(t))\|_2} \right)^T u(t) \\
&= \left(\frac{g(\theta_0, e(t))}{\|g(\theta_0, e(t))\|_2} \right)^T \frac{d}{dt} u(t) + u(t)^T \frac{1}{\|g(\theta_0, e(t))\|_2} \frac{d}{dt} g(\theta_0, e(t)) \\
&\quad - u(t)^T g(\theta_0, e(t)) \frac{g(\theta_0, e(t))^T \frac{d}{dt} g(\theta_0, e(t))}{\|g(\theta_0, e(t))\|_2^3} \qquad (91) \\
&g(\theta_0, e(t))^T \frac{d}{dt} u_t \leq \frac{2}{g_{\min}} \left\| \frac{d}{dt} g(\theta_0, e(t)) \right\|_2^2 \\
&\|g(\theta_0, e(t))\|_2 \leq 1 - (v(t)^T u(t))^2 \leq \frac{g_{\min}}{8} \\
&g_0 (1 - \epsilon/8) \leq 1 - (v(t)^T u(t))^2 \leq \frac{g_{\min}}{8} \\
&g_0 \leq 1 - \epsilon/4 \leq (v(t)^T u(t))^2 .
\end{aligned}$$

Here we utilize that $g_0 \geq g_{\min}$, where $g_{\min}$ is defined in (24). Let $z(t)$ satisfies the ODE

$$\frac{dz(t)}{dt} = (1 - \epsilon/4 - z(t)^2) g_0, \qquad (92)$$

with initialization $z(0) = v_0^T u_0$. Then, we note that

$$z(t) = \frac{\sqrt{1 - \epsilon/8} - 2\sqrt{1 - \epsilon/8}}{1 + c_3 \exp(2 g_0 t / (\sqrt{1 - \epsilon/8}))}, \qquad (93)$$

where $c_3 = \left(\frac{\sqrt{1 - \epsilon/8}}{v_0^T u_0} - 1 \right)^{-1}$. We can compute that

$$z(T_3) = 1 - \epsilon/4. \qquad (94)$$

According to the comparison theorem, for $t \leq \min\{T_0, T_3\}$, we have

$$v(t)^T u(t) \geq z(t). \qquad (95)$$

We first consider the case where $T_0 = \infty$. As $T_0 = \infty$, we have

$$v(T_3)^T u(T_3) \geq z(T_3) = 1 - \epsilon/4. \qquad (96)$$

Therefore, the second event holds for $T \leq T_3$.

Otherwise, we have $T_0 < \infty$. Recall that $u_1 = u(T_0)$ and $v_1 = \lim_{t \to T_0} v(t)$. Let $T_1 = \sup\{T | v(t)^T u(t) < v_1^T u_1, \forall t \in [0, T)\}$ and $T_2 = \sup\{T | v(t)^T u(t) < 1 - \epsilon/4, \forall t \in [0, T)\}$. If $T_2 \leq T_0$, for $t \in [0, T_2]$, we have

$$\frac{d}{dt} v(t)^T u(t) \geq 1 - \epsilon/4 - (1 - \epsilon)^2 g_0 > 0. \qquad (97)$$

Therefore, $v(t)^T u(t)$ monotonically increases in $[0, T_2]$. As $v(t)^T u(t) \geq z(t)$ for $t \in [0, T_0]$, we have that $z(T_2) \leq v(T_2)^T u(T_2) = 1 - \epsilon/4 = z(T_3)$. Hence, we have $T_2 \geq T_3$. Therefore, the second condition of the first event holds at $T = T_2$.

Then, we consider the case where $T_2 > T_0$. For $t \leq T_0$, we have $v(t)^T u(t) \leq 1 - \epsilon/4$. This implies that $v_1^T u_1 \leq 1 - \epsilon/4$. Apparently, we have $T_1 \geq T_0$. If $T_1 < T_0$, as $T_0 \geq T_2$, for $t \in [0, T_0]$, the inequality

(97) holds. This implies that $\lim_{t \to T_0 - 0} v(t)^T u(T_0) > v(T_1)^T u(T_1) = \lim_{t \to T_0 - 0} v(t)^T u(T_0)$, which leads to a contradiction. Therefore, we have $T_0 = T_1$. We note that

$$z(T^{\text{shift}}(u_1^T v_1)) = u_1^T v_1. \tag{98}$$

As $u(t)^T g(u(t); e(t)) \geq z(t)$ for $t \in [0, T_0]$, we have that $z(T_1) \geq u_1^T v_1 \geq z(T_4)$. Hence, we have $T_0 = T_1 \geq T^{\text{shift}}(u_1^T v_1)$. This completes the proof.

D.4 Proof of Proposition 6

We first introduce a lemma.

Lemma 7 Let $a, b \in \mathbb{R}^d$ and $0 < \epsilon < c$. Suppose that $\|a - b\|_2 \leq \epsilon$ and $\|a\|_2 \geq c$. Then, we have

$$\left\| \frac{a}{\|a\|_2} - \frac{b}{\|b\|_2} \right\|_2 \leq \frac{2\epsilon}{c}. \tag{99}$$

Proof As $\epsilon < c$, we have $\|b\|_2 > \|a\|_2 - \|a - b\|_2 \geq c - \epsilon > 0$. We first note that

$$\left| \|a\|_2^{-1} - \|b\|_2^{-1} \right| = \frac{|\|a\|_2 - \|b\|_2|}{\|a\|_2 \|b\|_2} \leq \frac{\epsilon}{c\|b\|_2}. \tag{100}$$

Therefore, we can compute that

$$\begin{aligned} &\left\| \frac{a}{\|a\|_2} - \frac{b}{\|b\|_2} \right\|_2 \\ &\leq \left\| \frac{a}{\|a\|_2} - \frac{b}{\|a\|_2} \right\|_2 + \left\| \frac{1}{\|a\|_2} - \frac{1}{\|b\|_2} \right\| \|b\|_2 \\ &\leq \frac{\epsilon}{c} + \frac{\epsilon}{c} = \frac{2\epsilon}{c}. \end{aligned} \tag{101}$$

This completes the proof.

Then we present the proof of Proposition 6.

Proof As $y^T (X u_0)_+ > 0$, with sufficiently small initialization and sufficiently small $\epsilon > 0$, we also have $e(0)^T (X u_0)_+ - y^T (X u_0)_+ / 4 \leq \|X\|_2 \|e(0) - y/4\|_2 > 0$. We prove that there exists a time T such that $u(T)^T v(T) \geq 1 - \frac{3}{4}$ by contradiction. Denote $v_0 = v(0)$. For all possible values of $\|g(u; y/4)\|_2$, we can arrange them from the smallest to the largest by $g_{(1)} \leq g_{(2)} \leq \cdots \leq g_{(p)}$. Let $T_i = \frac{1}{2\sqrt{1 - 8 g_{(i)}}} \log \frac{\sqrt{1 - 8 g_{(i)}} + 1 + 2\epsilon}{\sqrt{1 - 8 g_{(i)}} - 1 + 2\epsilon} \leq \frac{1}{\sqrt{1 - 8 g_{(i)}}} \log \frac{\sqrt{1 - 8 g_{(i)}} + g_{(i)}^{-1} v_0^T u_0}{\sqrt{1 - 8 g_{(i)}} - g_{(i)}^{-1} v_0^T u_0}$ and $T = \sum_{i=1}^{p} T_i$. Suppose that r_0 is sufficiently small such that statements in Lemma 2 holds. For According to Lemma 3, we can find $0 = t_0 < t_1 < \cdots$ such that for $i = 1, \ldots,$ $\text{sign}(Xu(t))$ is constant on $[t_{i-1}, t_i)$ and $\text{sign}(Xu(t_{i-1})) \neq \text{sign}(Xu(t_i))$. We write $u_i = u(t_i)$, $g_i = \|g(u(t_i); y/4)\|_2$,

$$g_i = \lim_{t \nearrow t_i} g(u(t); e(t)); \quad g_i = g(u(t_i); e(t)); \tag{102}$$

$v_i = \frac{g_i}{\|g_i\|_2}$ and $v_i = \frac{g_i}{\|g_i\|_2}$. We note that $g_i = g_{i+1}^+$. According to Lemma 3, we have

$$\begin{aligned} t_i - t_{i-1} &\leq \frac{1}{2\sqrt{1 - 8 g_{i-1}}} \log \frac{\sqrt{1 - 8 g_{i-1}} + (v_i)^T u_i}{\sqrt{1 - 8 g_{i-1}} - (v_i)^T u_i} \log \frac{\sqrt{1 - 8 g_{i-1}} + v_{i-1}^T u_{i-1}}{\sqrt{1 - 8 g_{i-1}} - v_{i-1}^T u_{i-1}} \\ &\leq \frac{1}{2\sqrt{1 - 8 g_{\min}}} \log \frac{\sqrt{1 - 8 g_{i-1}} + (v_i)^T u_i}{\sqrt{1 - 8 g_{i-1}} - (v_i)^T u_i} \log \frac{\sqrt{1 - 8 g_{i-1}} + v_{i-1}^T u_{i-1}}{\sqrt{1 - 8 g_{i-1}} - v_{i-1}^T u_{i-1}}. \end{aligned} \tag{103}$$

Here we utilize that $g_{i-1} \geq g_{\min}$, where $g_{\min}$ is defined in (24). This implies that

$$\frac{\sqrt{1 - 8 g_{i-1}} + (v_i)^T u_i}{\sqrt{1 - 8 g_{i-1}} - (v_i)^T u_i} \leq e^{2\sqrt{1 - 8 g_{\min}} (t_i - t_{i-1})} \frac{\sqrt{1 - 8 g_{i-1}} + v_{i-1}^T u_{i-1}}{\sqrt{1 - 8 g_{i-1}} - v_{i-1}^T u_{i-1}}. \tag{104}$$

We can show that for t satisfying $t \quad \frac{1}{2\sqrt{1 \quad =8 g_{min}}} \log \frac{\sqrt{1 \quad =8}+1 \quad =2}{\sqrt{1 \quad =8} \quad 1+ \quad =2} \quad \log \frac{1+ g_{min}^{1} v_0^T u_0}{1 \quad g_{min}^{1} v_0^T u_0}$ and $t \quad T$, we have $kg(u(t); \quad)k_2 > g_{min}$. According to Lemma 3, as $g_i \quad g_{min}$, we have

$$\frac{\sqrt{1 \quad =8} + g_{min}^{1} (g_i \quad)^T u_i}{\sqrt{1 \quad =8} \quad g_{min}^{1} (g_i \quad)^T u_i} \quad e^{2\sqrt{1 \quad =8} g_{min} (t_i \quad t_i \quad 1)} \frac{\sqrt{1 \quad =8} + g_{min}^{1} g_i^T \quad 1 u_i \quad 1}{\sqrt{1 \quad =8} \quad g_{min}^{1} g_i^T \quad 1 u_i \quad 1}. \tag{105}$$

This implies that

$$\frac{\sqrt{1 \quad =8} + g_{min}^{1} (g_i \quad)^T u_i}{\sqrt{1 \quad =8} \quad g_{min}^{1} (g_i \quad)^T u_i} \quad e^{2\sqrt{1 \quad =8} g_{min} \, t_i} \frac{\sqrt{1 \quad =8} + g_{min}^{1} v_0^T u_0}{\sqrt{1 \quad =8} \quad g_{min}^{1} v_0^T u_0}; \tag{106}$$

or equivalently, for any > 0, we have

$$\frac{\sqrt{1 \quad =8} + g_{min}^{1} (g(u(t); e(t))^T u(t)}{\sqrt{1 \quad =8} \quad g_{min}^{1} g(u(t); e(t))^T u(t)} \quad e^{2\sqrt{1 \quad =8} g_{min} \, t} \frac{\sqrt{1 \quad =8} + g_{min}^{1} v_0^T u_0}{\sqrt{1 \quad =8} \quad g_{min}^{1} v_0^T u_0}. \tag{107}$$

Here we utilize that $g(u(t); \quad (t))^T u(t)$ is continuous w.r.t. t. Therefore, for $t \quad \frac{1}{2g_{min}} \log \frac{2 \quad}{\quad} \log \frac{1+ g_{min}^{1} v_0^T u_0}{1 \quad g_{min}^{1} v_0^T u_0}$, we have

$$\frac{1 + g_{min}^{1} (g(u(t); e(t)))^T u(t)}{1 \quad g_{min}^{1} g(u(t); e(t))^T u(t)} \quad \frac{\sqrt{1 \quad =8} + 1 \quad =2}{\sqrt{1 \quad =8} \quad 1 + \quad =2}. \tag{108}$$

This implies that

$$g_{min}^{1} (g(u(t); e(t)))^T u(t) \quad 1 \quad =2. \tag{109}$$

If $kg(u(t); \quad)k_2 = g_{min}$, as the statements in Lemma 2 hold, we can compute that

$$kg(u(t); \quad) \quad g(u(t); e(t))k_2 \quad \frac{g_{min}}{4} = \frac{}{4} kg(u(t); \quad)k_2; \tag{110}$$

which implies that

$$kg(u(t); e(t))k_2 \quad (1 + \quad =4) kg(u(t); \quad)k_2. \tag{111}$$

Therefore, we have

$$\begin{aligned} v(t)^T u(t) &= \frac{(g(u(t); e(t)))^T u(t)}{kg(u(t); e(t))k_2} \\ &\quad \frac{1}{1 + \quad =4} \frac{(g(u(t); e(t)))^T u(t)}{g_{min}} \\ &\quad \frac{1 \quad =2}{1 + \quad =4} \quad 1 \quad \frac{3}{4}. \end{aligned} \tag{112}$$

This leads to a contradiction.

Analogously, we can show that for $\sum_{j=1}^{i} T_i$, we have $kg(u(t); =4)k_2 > g_{(i)}$. Thus, by taking $t \quad \sum_{i=1}^{p} T_i$, we have $kg(u(t); =4)k_2 > g_{(p)} = g_{max}$. However, from the definition of g_{max}, we have $kg(u(t); =4)k_2 \quad g_{max}$. This leads to a contradiction. Therefore, there exists a time $T = \sum_{i=1}^{p} T_i = O(\log \quad ^{1})$ such that $v(T)^T u(T) \quad 1 \quad \frac{3}{4}$.

We note that $kg(u(T); =4)k_2 \quad g_{min}$. As the statements in Lemma (2) hold, we have

$$kg(u(T); =4) \quad g(u(T); e(T))k_2 \quad \frac{g_{min}}{8} \tag{113}$$

According to Lemma 7, we have

$$\frac{g(u(T); =4)}{kg(u(T); =4)k_2} \quad \frac{g(u(T); e(T))}{kg(u(T); e(T))k_2} \quad \frac{2kg(u(T); =4) \quad g(u(T); e(T))k_2}{g_{min}} \quad \frac{}{4}. \tag{114}$$

This implies that

$$u(T)^T \frac{g(u(T); y=4)}{\|g(u(T); y=4)\|_2}$$

$$u(T)^T v(T) \quad \frac{g(u(T); \quad)}{\|g(u(T); \quad)\|_2} \quad \frac{g(u(T); e(T))}{\|g(u(T); e(T))\|_2} \quad _2 \tag{115}$$

$$1 \quad :$$

Hence, we have

$$\cos\angle(u(T); g(u(T); y)) = u(T)^T \frac{g(u(T); y)}{\|g(u(T); y)\|_2} = u(T)^T \frac{g(u(T); y=4)}{\|g(u(T); y=4)\|_2} \quad 1 \quad :$$

This completes the proof.

D.5 Proof of Lemma 4

Proof This is proved in Lemma 2 in (Phuong & Lampert, 2021). Here we provide an alternative proof. It is sufficient to prove for the case of local maximizer. Suppose that w is a local maximizer of $y^T(Xw)_+$ in B. We first consider the case where $y^T(Xw)_+ > 0$.

If there exists $n \in [N]$ such that $\langle w; x_n\rangle \quad 0$ and $y_n = 1$. Consider $v = x_n/\|x_n\|_2$ and let $w = \frac{w+v}{\|w+v\|_2}$, where > 0. For index $n' \in [N]$ such that $y_{n'} = 1$, as the dataset is orthogonal separable, we have $x_{n'}^T x_n > 0$ and

$$x_{n'}^T(w + v) = x_{n'}^T w + \frac{}{\|x_n\|_2} x_{n'}^T x_n > x_{n'}^T w: \tag{116}$$

This implies that $(x_{n'}^T w)_+ \quad (x_{n'}^T w)_+$. For $y_{n'} = \quad 1$, as the data is orthogonal separable, we note that $x_{n'}^T x_n \quad 0$ and

$$x_{n'}^T(w + v) = x_{n'}^T w + \frac{}{\|x_n\|_2} x_{n'}^T x_n \quad x_{n'}^T w: \tag{117}$$

This implies that $(x_j^T w)_+ \quad (x_j^T w)_+$. In summary, we have

$$y^T(X(w + v))_+ = \sum_{n=1}^{N} y_n (x_j^T(w + v))_+ \quad \sum_{n=1}^{N} y_n (x_j^T w)_+ = y^T(Xw)_+ > 0 \tag{118}$$

If $\langle w; x_n\rangle < 0$, then $w^T v < 0$. This implies that with sufficiently small $$, we have $\|w + v\|_2 < \|w\|_2 = 1$. Therefore,

$$y^T(Xw \quad))_+ = \frac{1}{\|w + v\|_2} y^T(X(w + v))_+ > y^T(X(w + v))_+ \quad y^T(Xw)_+; \tag{119}$$

which leads to a contradiction. If $\langle w; x_n\rangle = 0$, we note that

$$(x_n^T(w + v))_+ = \quad > (x_n^T w)_+: \tag{120}$$

This implies that

$$y^T(X(w + v))_+ \quad y^T(Xw)_+ + \quad : \tag{121}$$

We also note that $\|w + v\|_2 = \sqrt{1 + ^2} = 1 + O(^2)$. Therefore, with sufficiently small $$, we have

$$y^T(Xw \quad)_+ \quad \frac{y^T(Xw)_+ +}{\sqrt{1 + ^2}} > y^T(Xw)_+: \tag{122}$$

We then consider the case where $y^T(Xw)_+ < 0$. Apparently, we can make $y^T(Xw)_+$ larger by replacing w by $(1 \quad)w$, where $\in (0; 1)$, which leads to a contradiction.

Finally, we consider the case where $y^T(Xw)_+ = 0$. This implies that

$$\sum_{n: y_n = 1} (x_j^T w)_+ = \sum_{n: y_n = \quad 1} (x_j^T w)_+: \tag{123}$$

As $(Xw)_+ \neq 0$, this implies that there exists at least for one index $n \in [N]$ such that $y_n = 1$ and $x_n^T w > 0$. Let $v = x_n/\|x_n\|_2$. We note that $\frac{1}{\|w+v\|_2} y^T(X(w + v))_+ > 0$ for > 0. This leads to a contradiction.

D.6 PROOF OF PROPOSITION 7

It is sufficient to consider the case of the local maximizer. Denote $Q = \{\epsilon \in \{-1, 0, 1\}^N \mid \mathrm{diag}(\epsilon) \in P\}$. For $\epsilon, \epsilon' \in Q$, we say $\epsilon \leq \epsilon'$ if for all index $n \in [N]$ with $\epsilon_n \neq 0$, $\epsilon'_n = \epsilon_n$. We say $\epsilon \in Q$ is open if $\epsilon_n \neq 0$ for $n \in [N]$. Define

$$S_\epsilon = \{u \mid \mathrm{sign}(Xu) = \epsilon\}. \tag{124}$$

We start with the two lemmas.

Lemma 8 Let $\epsilon \in R^N$. Suppose that u_0 satisfies that $u_0 = \frac{g(u_0; \epsilon)}{\|g(u_0; \epsilon)\|_2}$. Let $\epsilon = \mathrm{sign}(u_0)$. Then, $v \in B_2$ is a local maximizer of $\epsilon^T (Xu)_+$ in B_2 if for any open $\epsilon' \leq \epsilon$ satisfying $\epsilon' \neq 0$, we have $\|g(\epsilon; y)\|_2 = \|g(\epsilon', y)\|_2$.

PROOF Suppose that ϵ is open. Then S_ϵ is an open set. In a small neighborhood around u_0, $u = \frac{g(u_0; \epsilon)}{\|g(u_0; \epsilon)\|_2} = \frac{g(u; \epsilon)}{\|g(u; \epsilon)\|_2}$, $\epsilon^T (Xu)_+ = u^T g(\epsilon; y)$ is a linear function of u. The Riemannian gradient of $u^T g(\epsilon; y)$ at v is zero. This implies that u locally maximizes $\epsilon^T (Xu)_+$.

Suppose that there exists at least one zero in ϵ. Consider any $v \in B$ satisfying $u_0^T v = 0$. Let $s > 0$ be a small constant such that for any $s \in (0, s]$, $u_0 + sv \in S_{\epsilon'}$ where $\epsilon' \neq 0$. Let $u_s = \frac{u + sv}{\sqrt{1 + s^2}}$.

Suppose that $\|g(\epsilon''; y)\|_2 < \|g(\epsilon; y)\|_2$ for all open $\epsilon'' \leq \epsilon$ satisfying $\epsilon'' \neq \epsilon'$. For any $\epsilon' \neq \epsilon$ with $\epsilon' \neq 0$, we construct ϵ'' by $\epsilon''_i = \epsilon_i$ for $n \in [N]$ such that $\epsilon'_n = 0$ and $\epsilon''_n = \epsilon'_n$ for $n \in [N]$ such that $\epsilon'_n \neq 0$. We note that $\|g(\epsilon''; y)\|_2 < \|g(\epsilon'; y)\|_2$. Thus, $\|g(\epsilon'; y)\|_2 - \|g(\epsilon''; y)\|_2 > \|g(\epsilon'; y)\|_2 - \|g(\epsilon; y)\|_2$. As $|\epsilon^T (Xu_s)_+| = |g(\epsilon'; y)^T u_s| > \|g(\epsilon'; y)\|_2 - \|g(\epsilon; y)\|_2$, we have $|\epsilon^T (Xu_s)_+| > \|g(\epsilon'; y)\|_2 - \|g(\epsilon; y)\|_2$. Therefore, u is a local maximizer of $\epsilon^T (Xu)_+$.

Lemma 9 Suppose that the dataset is orthogonal separable. Let $\epsilon \in R^N$ satisfy that $\mathrm{diag}(\epsilon) \in P$. Suppose that u_0 satisfies that $u_0 = \frac{g(u_0; \epsilon)}{\|g(u_0; \epsilon)\|_2}$. Then, for any $\epsilon' \neq 0$ satisfying $\epsilon' \leq \epsilon$, we have $\|g(\epsilon'; y)\|_2 = \|g(\epsilon; y)\|_2$.

PROOF If there exists $n \in [N]$ such that $\epsilon_n = 1$ and $y_n = -1$, as the data is orthogonal separable, we note that

$$x_n^T g(\epsilon; y) = x_n^T \Big(\sum_{n': \epsilon_{n'} > 0} \epsilon_{n'} x_{n'} \Big) = y_n (y_n x_n)^T \Big(\sum_{n': \epsilon_{n'} > 0} (\epsilon_{n'} y_{n'}) y_{n'} x_{n'} \Big) \geq 0, \tag{125}$$

which contradicts with $\mathrm{sign}(x_n^T g(\epsilon; y)) = \mathrm{sign}(x_n^T u_0) = \epsilon_n = 1$.

Suppose that there exists $n \in [N]$ such that ϵ_n and $y_n = 1$. Then, as the dataset is orthogonal separable, then, for index $n_1 \in [N]$ such that $\epsilon_{n_1} = 0$, we note that $y_{n_1} \neq 1$. Otherwise,

$$x_{n_1}^T g(\epsilon; y) = x_{n_1}^T \Big(\sum_{n_2: \epsilon_{n_2} > 0} \epsilon_{n_2} x_{n_2} \Big) = x_{n_1}^T \Big(\sum_{n_2: \epsilon_{n_2} > 0} (\epsilon_{n_2} y_{n_2}) y_{n_2} x_{n_2} \Big) > 0; \tag{126}$$

which contradicts with $\mathrm{sign}(x_{n_1}^T g(\epsilon; y)) = \mathrm{sign}(x_{n_1}^T u_0) = \epsilon_{n_1} = 0$. This also implies that the index set $\{n \in [N] \mid \epsilon_n > 0\}$ include all data with $y_n = 1$.

If there exists $\epsilon' \leq \epsilon$ such that $\epsilon' \neq \epsilon$ and $\|g(\epsilon'; y)\|_2 > \|g(\epsilon; y)\|_2$. Then, there exists at least one index $n \in [N]$ such that $\epsilon_n \neq 0$ and $\epsilon'_n = 1$. However, from the previous derivation, we note that $y_n = -1$ and

$$x_n^T g(\epsilon'; y) = x_n^T \Big(\sum_{j: \epsilon'_{n_1} > 0} \epsilon_{n_1} x_{n_1} \Big) = x_n^T \Big(\sum_{n_1: \epsilon'_{n_1} > 0} (\epsilon_{n_1} y_{n_1}) y_{n_1} x_{n_1} \Big) < 0; \tag{127}$$

which contradicts with $\epsilon'_n = 1$.

By combining Lemma 8 and 9, we complete the proof.

D.7 PROOF OF THEOREM 4

PROOF For almost all initialization, we can find two neurons such that $\text{sign}(w_{2;i_+}) = \text{sign}(y^T(Xw_{1;i_+})_+) = 1$ and $\text{sign}(w_{2;i_-}) = \text{sign}(y^T(Xw_{1;i_-})_+) = -1$ at initialization. By choosing a sufficiently small > 0 in Proposition 6, there exist two neurons $w_{1;i_+}; w_{1;i_-}$ and times $T_+; T_- > 0$ such that $\cos\angle(w_{1;i_+}(T_+); g(w_{1;i_+}(T_+); y)) > 1 -$ and $\cos\angle(w_{1;i_+}(T_+); g(w_{1;i_+}(T_+); y)) < -(1 -)$. This implies that $w_{1;i_+}(T_+)$ and $w_{1;i_-}(T_+)$ are sufficiently close to certain stationary points of gradient flow maximizing/minimizing $y^T g(Xu_+)$ over B, i.e., $\{u \in B : \cos\angle(u; g(u; y)) = 1\}$. As the dataset is orthogonal separable, according to Lemma 4 and Proposition 7, the corresponding diagonal matrices $\hat{D}_{i_+}(T_+)$ and $\hat{D}_{i_-}(T_-)$ satisfy that $\hat{D}_{i_+}(T_+) = \text{diag}(I(y = 1))$ and $\hat{D}_{i_-}(T_-) = \text{diag}(I(y = -1))$. According to Lemma 3 in (Phuong & Lampert, 2021), we have $\hat{D}_{i_+}(t) = \text{diag}(I(y = 1))$ and $\hat{D}_{i_-}(t) = \text{diag}(I(y = -1))$ hold for $t = \max\{T_+; T_-\}$.

With $t \to \infty$, according to Proposition 4, the dual variable in the KKT point of the non-convex max-margin problem (13) is dual feasible, i.e., satisfies (16). Suppose that π is a limiting point of $\frac{(t)}{\|(t)\|_2}_{t>0}$ and θ is the corresponding dual variable. From Theorem 1, we note that the pair $(\pi; \theta)$ corresponds to the KKT point of the convex max-margin problem (14).

E PROOFS OF MAIN RESULTS ON MULTI-CLASS CLASSIFICATION

E.1 PROOF OF PROPOSITION 1

The neural network training problem (4) can be separated into K subproblems. Each of these subproblems corresponds to the neural network training problem (3) for binary classification. For each subproblem, by applying Proposition 2, we complete the proof.

E.2 PROOF OF THEOREM 1

We note that the neural network training problem (4) can be separated into K subproblems. Each of these subproblems corresponds to the neural network training problem (3) for binary classification. By applying Proposition 6 with to each subproblem with $= y_k$, we complete the proof.

E.3 PROOF OF THEOREM 2

Similarly, the corresponding non-convex max-margin problem (5) and the convex max-margin problem (7) can be separated into K subproblems. Each of these subproblems corresponds to the non-convex max-margin problem (2) and the convex max-margin problem (14) for binary classification. By applying Theorem 4 to each subproblem with $= y_k$, we complete the proof.

F NUMERICAL EXPERIMENT

F.1 DETAILS ON FIGURE 5

We provide the experiment setting in Figure 1 and 5 as follows. The dataset is given by $X = \begin{bmatrix} 1.65 & 0.47 \\ 0.47 & 1.35 \end{bmatrix} \in \mathbb{R}^{2 \times 2}$ and $y = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \in \mathbb{R}^2$. Here we have $N = 2$ and $d = 2$. We note that this dataset is orthogonal separable but not spike-free. We plot the ellipsoid set and the rectified ellipsoid set in Figure 6.

Figure 6: The ellipsoid set and the rectified ellipsoid set. Orthogonal separable dataset.

We enumerate all possible hyperplane arrangements in the set $\mathcal{P}$ and solve the convex max-margin problem (14) via CVXPY to obtain the following non-zero neurons

$$u_{1;3} = \begin{matrix} 0.58 \\ 0.16 \end{matrix}; \quad w^{0}_{1;2} = \begin{matrix} 0.23 \\ 0.66 \end{matrix}: \tag{128}$$

We note that the dual problem (15) is equivalent to

$$\begin{aligned} \max \ & {}^{T} y; \\ \text{s.t.} \ & k X^{T} D_j \quad X^{T} (2D_j \quad I) z_{j;+} k_2 \quad 1; 8j \ 2 \ [p]; \\ & k \quad X^{T} D_j \quad X^{T} (2D_j \quad I) z_{j;} \quad k_2 \quad 1; 8j \ 2 \ [p]; \\ & z_{j;+} \quad 0; z_{j;} \quad 0; 8j \ 2 \ [p]; \operatorname{diag}(y) \quad 0: \end{aligned} \tag{129}$$

The above problem is a second-order cone program (SOCP) and can be solved via standard convex optimization frameworks such as CVX and CVXPY. We solve (129) to obtain the optimal dual variable . For the geometry of the dual problem, as the dataset is orthogonal separable, the set $f : \max_{kuk_2 \ 1} j \ {}^{T}(Xu)_{+} j \quad 1g$ reduces to $f : \max_{kuk_2 \ 1} j \ {}^{T}(Xu_1)_{+} j \quad 1; \ {}^{T}(Xu_2)_{+} j \quad 1g$, where $u_1; u_2$ correspond to two vectors at the spikes of the rectified ellipsoid set. We draw the sets $f : \max_{kuk_2 \ 1} j \ {}^{T}(Xu)_{+} j \quad 1g$, $f :$, the optimal dual variable and the direction of in Figure 2.

For each $D_j \ 2 \ \mathcal{P}$, we solve for the vector u_j which maximize/minimize ${}^{T} D_j X u_j$ with the constraints $ku_j k_2 \quad 1$ and $(2D_j \quad I) X u_j \quad 0$. We plot the rectified ellipsoid set $f(Xu)_{+} j k u k_2 \quad 1g$, vectors u_j, neurons in the optimal solution to (14) scaled to unit ℓ_2-norm and the direction of in Figure 1. We note that each neuron in the optimal solution from (14) (scaled to unit ℓ_2-norm) maximize/minimize the corresponding ${}^{T} D_j X u_j$ given $(2D_j \quad I) X u_j \quad 0$.

Then, we consider a two-layer ReLU network with $= 10$ neurons and apply the gradient descent method to train on the logistic loss (3). Let $\hat{w}_{1;i} = \frac{w_{1;i}}{k w_{1;i} k_2}$ for $i \ 2 \ [m]$. We plot $\hat{w}_{1;i}$ and $(X \hat{w}_{1;i})_{+}$ at iteration $f 10^{l} j l = 0; : : : ; 4g$ along with neurons in the optimal solution to (14) scaled to unit ℓ_2-norm in Figure 5. Certain neurons do not move, while the activated neurons trained by gradient descent tend to converge to the direction of the neurons in the optimal solution to (14).

We repeat the training on the logistic loss (3) with the gradient descent method several times and we plot the trajectories in Figure 7.

Figure 7: Multiple independent random initializations of gradient descent trajectories on the same orthogonal separable dataset.

F.2 EXPERIMENT ON SPIKE-FREE DATASET

We repeat the previous numerical experiment on a non-spike-free dataset: $X = \begin{bmatrix} 1.65 & 0.47 \\ 0.47 & 1.35 \end{bmatrix}$ 2 R^2 ^{2} and $y = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ 2 R^2. Similarly, we plot the ellipsoid set and the recti ed set in Figure 8.

Figure 8: The ellipsoid set and the recti ed ellipsoid set for a non-spike-free dataset.

We enumerate all possible hyperplane arrangements in the and solve the convex max-margin problem (14) via CVXPY to obtain the following non-zero neuron

$$u_{1;4} = \begin{bmatrix} 0.43 \\ 0.59 \end{bmatrix} \tag{130}$$

We plot the recti ed ellipsoid set $f(Xu)_+$ $jkuk_2$ $1g$, vectors u_j, neurons in the optimal solution to (14) scaled to unit $_2$-norm and the direction of in Figure 9. We also plot $v_{1;i}$ and $(Xv_{1;i})_+$ at iteration $f10^j jl = 0; : : : ; 4g$ along with neurons in the optimal solution to (14) scaled to unit $_2$-norm in Figure 10.

Figure 9: Reciti ed Ellipsoidal set and corresponding extreme points for a non-spike-free dataset.