## A Preliminaries on the Clarke subdifferential and the KKT conditions

Below we define the Clarke subdifferential, and review the definition of the KKT conditions for non-smooth optimization problems (cf. Lyu and Li [2019], Dutta et al. [2013]).
Let $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be a locally Lipschitz function. The Clarke subdifferential [Clarke et al., 2008] at $\mathrm{x} \in \mathbb{R}^{d}$ is the convex set

$$
\partial^{\circ} f(\mathbf{x}):=\operatorname{conv}\left\{\lim _{i \rightarrow \infty} \nabla f\left(\mathbf{x}_{i}\right) \mid \lim _{i \rightarrow \infty} \mathbf{x}_{i}=\mathbf{x}, f \text { is differentiable at } \mathbf{x}_{i}\right\}
$$

If $f$ is continuously differentiable at $\mathbf{x}$ then $\partial^{\circ} f(\mathbf{x})=\{\nabla f(\mathbf{x})\}$. For the Clarke subdifferential the chain rule holds as an inclusion rather than an equation. That is, for locally Lipschitz functions $z_{1}, \ldots, z_{n}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ and $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$, we have

$$
\partial^{\circ}(f \circ \mathbf{z})(\mathbf{x}) \subseteq \operatorname{conv}\left\{\sum_{i=1}^{n} \alpha_{i} \mathbf{h}_{i}: \boldsymbol{\alpha} \in \partial^{\circ} f\left(z_{1}(\mathbf{x}), \ldots, z_{n}(\mathbf{x})\right), \mathbf{h}_{i} \in \partial^{\circ} z_{i}(\mathbf{x})\right\}
$$

Consider the following optimization problem

$$
\begin{equation*}
\min f(\mathbf{x}) \quad \text { s.t. } \quad \forall n \in[N] \quad g_{n}(\mathbf{x}) \leq 0 \tag{7}
\end{equation*}
$$

where $f, g_{1}, \ldots, g_{n}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ are locally Lipschitz functions. We say that $\mathbf{x} \in \mathbb{R}^{d}$ is a feasible point of Problem (7) if $\mathbf{x}$ satisfies $g_{n}(\mathbf{x}) \leq 0$ for all $n \in[N]$. We say that a feasible point $\mathbf{x}$ is a KKT point if there exists $\lambda_{1}, \ldots, \lambda_{N} \geq 0$ such that

1. $\mathbf{0} \in \partial^{\circ} f(\mathbf{x})+\sum_{n \in[N]} \lambda_{n} \partial^{\circ} g_{n}(\mathbf{x})$;
2. For all $n \in[N]$ we have $\lambda_{n} g_{n}(\mathbf{x})=0$.

## B Details on Example 1

We use here the notation $\Phi(\boldsymbol{\theta} ; x)=\mathcal{N}_{\boldsymbol{\theta}}(x)$, and denote by $\sigma^{\prime}(\cdot)$ a sub-gradient of $\sigma$, namely, $\sigma(z)=\mathbb{1}[z>0]$ if $z \neq 0$ and $\sigma^{\prime}(0) \in[0,1]$ (the exact value in this case is not important here). For every $j \in\{1,2\}$ we have

$$
\begin{aligned}
\nabla_{w_{j}} \mathcal{L}(\boldsymbol{\theta}) & =\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right) \cdot y_{i} \nabla_{w_{j}} \Phi\left(\boldsymbol{\theta} ; x_{i}\right) \\
& =\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}\left(v_{1} \sigma\left(w_{1} x_{i}+b_{1}\right)+v_{2} \sigma\left(w_{2} x_{i}+b_{2}\right)\right) \cdot v_{j} \sigma^{\prime}\left(w_{j} x_{i}+b_{j}\right) x_{i}
\end{aligned}
$$

Likewise,

$$
\begin{aligned}
\nabla_{v_{j}} \mathcal{L}(\boldsymbol{\theta}) & =\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right) \cdot y_{i} \nabla_{v_{j}} \Phi\left(\boldsymbol{\theta} ; x_{i}\right) \\
& =\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}\left(v_{1} \sigma\left(w_{1} x_{i}+b_{1}\right)+v_{2} \sigma\left(w_{2} x_{i}+b_{2}\right)\right) \cdot \sigma\left(w_{j} x_{i}+b_{j}\right)
\end{aligned}
$$

and

$$
\begin{aligned}
\nabla_{b_{j}} \mathcal{L}(\boldsymbol{\theta}) & =\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right) \cdot y_{i} \nabla_{b_{j}} \Phi\left(\boldsymbol{\theta} ; x_{i}\right) \\
& =\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}\left(v_{1} \sigma\left(w_{1} x_{i}+b_{1}\right)+v_{2} \sigma\left(w_{2} x_{i}+b_{2}\right)\right) \cdot v_{j} \sigma^{\prime}\left(w_{j} x_{i}+b_{j}\right)
\end{aligned}
$$

Note that if $w_{2}=b_{2}=v_{2}=0$ then we have $\nabla_{w_{2}} \mathcal{L}(\boldsymbol{\theta})=\nabla_{v_{2}} \mathcal{L}(\boldsymbol{\theta})=\nabla_{b_{2}} \mathcal{L}(\boldsymbol{\theta})=0$. Since these parameters are initialized at zero, then they remain zero throughout the training. Moreover, Suppose that $w_{1}=0$ and $b_{1}=v_{1}=\alpha$ for some $\alpha>0$, and that $w_{2}=b_{2}=v_{2}=0$, then we have

$$
-\frac{d w_{1}}{d t}=\nabla_{w_{1}} \mathcal{L}(\boldsymbol{\theta})=\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}(\alpha \sigma(\alpha)) \cdot \alpha \sigma^{\prime}(\alpha) x_{i}=\frac{1}{2}\left(\ell^{\prime}\left(\alpha^{2}\right) \cdot \alpha \cdot 4+\ell^{\prime}\left(\alpha^{2}\right) \cdot \alpha \cdot(-4)\right)=0
$$

$$
\begin{aligned}
-\frac{d v_{1}}{d t} & =\nabla_{v_{1}} \mathcal{L}(\boldsymbol{\theta})=\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}(\alpha \sigma(\alpha)) \cdot \sigma(\alpha)=\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}\left(\alpha^{2}\right) \cdot \alpha \\
-\frac{d b_{1}}{d t} & =\nabla_{b_{1}} \mathcal{L}(\boldsymbol{\theta})=\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}(\alpha \sigma(\alpha)) \cdot \alpha \sigma^{\prime}(\alpha)=\frac{1}{2} \sum_{i=1}^{2} \ell^{\prime}\left(\alpha^{2}\right) \cdot \alpha
\end{aligned}
$$

Hence, for every $t$ we have $w_{1}(t)=0$ and $b_{1}(t)=v_{1}(t)=\alpha(t)$ where $\alpha(t)>0$ is monotonically increasing.

As a result, the KKT point $\boldsymbol{\theta}^{*}$ is such that $w_{2}^{*}=b_{2}^{*}=v_{2}^{*}=w_{1}^{*}=0$, and $b_{1}^{*}=v_{1}^{*}=\alpha^{*}$ for some $\alpha^{*}>0$. Since $\boldsymbol{\theta}^{*}$ satisfies the KKT conditions of Problem (4), then we have

$$
\alpha^{*}=b_{1}^{*}=\sum_{i=1}^{2} \lambda_{i} y_{i} \nabla_{b_{1}} \Phi\left(\boldsymbol{\theta}^{*} ; x_{i}\right)
$$

where $\lambda_{i} \geq 0$ and $\lambda_{i}=0$ if $y_{i} \Phi\left(\boldsymbol{\theta}^{*} ; x_{i}\right) \neq 1$. Hence, there is $i$ such that $y_{i} \Phi\left(\boldsymbol{\theta}^{*} ; x_{i}\right)=1$. Thus, we have $1=y_{i} \Phi\left(\boldsymbol{\theta}^{*} ; x_{i}\right)=\left(\alpha^{*}\right)^{2}$ which implies $\alpha^{*}=1$. Therefore $\boldsymbol{\theta}^{*}=\boldsymbol{\theta}(0)$.

## C Proof of Theorem 3.1

Before we prove the theorem, we first state a few definitions that are specific for this appendix. Let $S=\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{n} \subseteq[-R, R] \times\{-1,1\}$ be a dataset such that $x_{1}<\ldots<x_{n}$ and let $I=\left\{i \in[n-1]: y_{i} \neq y_{i+1}\right\}$ where we denote the elements of $I$ using $i_{1}<\ldots<i_{r}$, and $i_{0}=-R, i_{r+1}=R$, where $r=|I|$. For all $j \in[r+1]$, define $I_{j}=\left(i_{j-1}, i_{j}+1\right)$ which is the $j$-th interval where the instances in the data do not change their classification. ${ }^{5}$ Given some function $\mathcal{L}$ and real number $\alpha$, we let $L_{\alpha}^{+}(\mathcal{L}):=\{\boldsymbol{\theta}: \mathcal{L}(\boldsymbol{\theta}) \geq \alpha\}$ denote the $\alpha$-superlevel set of $\mathcal{L}$.
Next, we state the following definitions, which establish sufficient conditions for our objective function to be well-behaved in the sense of having a strict direction of descent in a certain neighborhood.
Definition C. 1 (Separability). Under Assumption 2.2, we say that $\boldsymbol{\theta}$ is separable from $S$ with positive constants $\gamma, m<M, q<Q$ iffor all $j \in[r+1]$, there exist three neurons with weights and biases denoted by $w_{i}\left(I_{j}\right)$ and $b_{i}\left(I_{j}\right)$ for $i \in[3]$, and breakpoints $\beta_{1}\left(I_{j}\right)<\beta_{2}\left(I_{j}\right)<\beta_{3}\left(I_{j}\right) \in I_{j}$, which satisfy the following items:

1. $m \leq\left|w_{i}\left(I_{j}\right)\right|$ and $\left|b_{i}\left(I_{j}\right)\right| \leq M$ for all $i \in[3]$.
2. There exist four neurons, two with breakpoints in each of the intervals $(-\infty, 0),(0, \infty)$, that are distinct from the neurons in the previous item, are active on all the data instances and whose weights $w_{i}^{\prime}, b_{i}^{\prime}$ satisfy $m \leq\left|w_{i}^{\prime}\right|$ and $\left|b_{i}^{\prime}\right| \leq M$ for all $i \in[4]$. Moreover, their breakpoints satisfy $\left|\beta_{i}\right| \leq Q$ for all $i \in[4]$.
3. $q \leq \beta_{2}\left(I_{j}\right)-\beta_{1}\left(I_{j}\right), \beta_{3}\left(I_{j}\right)-\beta_{2}\left(I_{j}\right) \leq Q$ for all $j \in[r+1]$.
4. $\left|\beta_{1}\left(I_{j+1}\right)-\beta_{3}\left(I_{j}\right)\right| \leq Q$ for all $j \in[r]$.
5. $x_{i_{j}+1}-x_{i_{j}} \geq \gamma$ for all $j \in[r]$.

If the triplet of neurons satisfying the above items in an interval is not distinct, we assume w.l.o.g. that $\beta_{1}(\cdot)$ and $\beta_{3}(\cdot)$ return the left-most and right-most breakpoints satisfying the above, respectively.

The following definition is used to describe a neighborhood around the initialization point in which our separability assumption above holds.
Definition C. $2(\Delta$-hidden Neighborhood). Given a network $\mathcal{N}(\cdot)$, weights $\boldsymbol{\theta}$ and a constant $\Delta \geq 0$, we define the $\Delta$-hidden Neighborhood of $\mathcal{N}$ at $\boldsymbol{\theta}$ as the set

$$
U_{\Delta}(\boldsymbol{\theta}):=\left\{\boldsymbol{\theta}^{\prime}=\left[\mathbf{w}^{\prime}, \mathbf{b}^{\prime}, \mathbf{v}^{\prime}\right]:\left\|\left(w_{j}, b_{j}\right)-\left(w_{j}^{\prime}, b_{j}^{\prime}\right)\right\|_{2} \leq \Delta \quad \forall j \in[k], \quad \mathbf{v}^{\prime} \in \mathbb{R}^{k}\right\}
$$

[^0]That is, the neighborhood of balls of radius $\Delta$ centered at each hidden neuron of $\boldsymbol{\theta}$ and where the output neuron weights are arbitrary.

Following the above definitions, the following auxiliary lemmas will be used in the proof of the theorem. The technical lemma below establishes that a certain binary matrix is invertible and provides a bound on the spectral norm of its inverse.
Lemma C.1. Suppose that $A \in\{0,1\}^{d \times d}$, such that the first row of $A$ is all-ones, and each subsequent row $i$ is either $(1, \ldots, 1,0, \ldots, 0)$ with $i-1$ leading ones or $(0, \ldots, 0,1, \ldots, 1)$ with $i-1$ leading zeros. Then $A$ is invertible and we have $\left\|A^{-1}\right\|_{s p} \leq d$.

Proof. The invertability of $A$ follows from the fact that the first row of $A$ is an all-ones vector, since we can use elementary row operations to change all subsequent rows to start with a ' 0 ' and end with a ' 1 ' if needed, resulting in an upper triangular matrix with all-ones on its main diagonal which is thus invertible. To bound $\left\|A^{-1}\right\|_{\text {sp }}$, let $I_{d} \in\{0,1\}^{d \times d}$ denote the identity matrix. We will use Gaussian elimination to compute the entries of $A^{-1}$. We first subtract the first row from all the other rows that do not have leading zeros and then multiply by the constant -1 . Performing the same operation on $I_{d}$ results in a matrix $B$ whose rows are either standard unit vectors or the vector $(1, \ldots, 1,0,1, \ldots, 1)$. The resulting matrix after performing these operations on $A$ is and upper triangular matrix with ones in all of its diagonal and above the diagonal entries. Since it is readily seen that the inverse of such a matrix is a matrix with all zero entries except for the main diagonal which is all-ones and the first diagonal above it which comprises of all -1 's. Denote this matrix using $B^{\prime}$, we have that the inverse of $A$ is given by $B \cdot B^{\prime}$. The entries of $A^{-1}$ therefore must consist of dot products of a standard unit vector and vectors $(0, \ldots, 0,1,-1,0, \ldots, 0)$, or the vector $(1, \ldots, 1,0,1, \ldots, 1)$ and vectors $(0, \ldots, 0,1,-1,0, \ldots, 0)$. In both cases the dot product is an element of $\{-1,0,1\}$, and therefore we can bound $\left\|A^{-1}\right\|_{\text {sp }}$ by the Frobenius norm of $A^{-1}$ which is at most $d$.

The following key lemma establishes that when $\boldsymbol{\theta}$ is separable from $S$ then there exists a direction in weight space which strictly decreases our objective value.
Lemma C.2. Under Assumption 2.2, suppose that $\boldsymbol{\theta}$ is separable from $S$ with constants $\gamma, m, M, q, Q$, and that $\mathcal{L}(\boldsymbol{\theta}) \geq \frac{1}{2 n}$. Then

$$
\frac{1}{2}\|\nabla \mathcal{L}(\boldsymbol{\theta})\|_{2}^{2} \geq \frac{\gamma^{2} q^{2} m^{6}}{259200 n^{4} Q^{2} M^{4}}
$$

Proof. Since $\mathcal{L}(\boldsymbol{\theta})>\frac{1}{2 n}$, there must exist some $i \in[n]$ such that $\ell\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right)>\frac{1}{2 n}$. We now consider two possible cases, depending on the location of $x_{i}$ with respect to the breakpoints whose existence is guaranteed by Def. C.1, where will show the existence of a direction $\mathbf{u}$ which guarantees that $\mathcal{L}(\boldsymbol{\theta})$ is strictly decreasing.

- Suppose that $x_{i} \in I_{\ell}$ for some $\ell \in[r+1]$, such that $x_{i} \in\left(\beta_{1}\left(I_{\ell}\right), \beta_{3}\left(I_{\ell}\right)\right)$. Assume without loss of generality that $x_{i} \in\left(\beta_{2}\left(I_{\ell}\right), \beta_{3}\left(I_{\ell}\right)\right)$ (the proof is symmetric otherwise), and for ease of notation denote $\beta_{2}:=\beta_{1}\left(I_{\ell}\right), \beta_{3}:=\beta_{2}\left(I_{\ell}\right), \beta_{4}:=\beta_{3}\left(I_{\ell}\right)$. Then by Item 2 in the separability assumption, there exists a breakpoints $\beta_{1}<\beta_{2}$ such that all data instances in $\left(\beta_{2}, \beta_{4}\right)$ have the same classification and the neuron with breakpoint at $\beta_{1}$ is active on all the data instances. Moreover, Item 2 also guarantees the existence of two breakpoints that are distinct from the previous ones, which we denote by $\beta_{5}, \beta_{6}$, where $\beta_{5}<\beta_{1}<0$ and $\beta_{6}>0$ are active on all the data points. We will now show the existence of a depth-2 ReLU network which consists of six hidden neurons with weights $\mathbf{w}=\left(w_{1}, \ldots, w_{6}\right)$ and biases $\mathbf{b}=\left(b_{1}, \ldots, b_{6}\right)$ (corresponding to the breakpoints $\beta_{1}, \ldots, \beta_{6}$ defined above) which computes the piece-wise linear function $f:\left[\beta_{1}, \beta_{6}\right) \rightarrow \mathbb{R}$ given by

$$
f(x):=\left\{\begin{array}{ll}
0 & x \in\left[\beta_{1}, \beta_{2}\right] \\
\frac{1}{\beta_{3}-\beta_{2}} x-\frac{\beta_{2}}{\beta_{3}-\beta_{2}} & x \in\left(\beta_{2}, \beta_{3}\right) \\
\frac{1}{\beta_{3}-\beta_{4}} x-\frac{\beta_{4}}{\beta_{3}-\beta_{4}} & x \in\left[\beta_{3}, \beta_{4}\right] \\
0 & x \in\left(\beta_{4}, \beta_{6}\right)
\end{array} .\right.
$$

The intuition behind the approximation is that we can use the first four neurons to approximate the slopes of the function $f$, and the last two remaining neurons to simulate a bias term


Figure 1: The plots of $f(x)$ (green) and $\Phi(\boldsymbol{\theta} ; x)+h \cdot f(x)$ (black) for various values of $h$. Moving $\boldsymbol{\theta}$ in the direction of $\mathbf{u}$ which computes $f(\cdot)$ strictly decreases the loss over the positively-labeled instances in the interval $(0,2)$ without affecting the rest of the dataset. Best viewed in color.
which would shift the function approximated by the network to overlap $f$ in the relevant domain of approximation (see Figure 1 for an illustration).
More formally, define $W:=\operatorname{diag}\left(w_{1}, \ldots, w_{4}\right) \in \mathbb{R}^{4 \times 4}$, define the masking matrix $A \in$ $\mathbb{R}^{4 \times 4}$ with entries $a_{j, j^{\prime}}=\mathbb{1}\left\{j\right.$-th neuron is active on the interval starting with $\left.\beta_{j^{\prime}}\right\}$ and let

$$
\mathbf{d}=\left(0, \frac{1}{\beta_{3}-\beta_{2}}, \frac{1}{\beta_{3}-\beta_{4}}, 0\right)
$$

Thus, to match the slopes computed by the a depth-2 ReLU network with weights $\mathbf{w}, \mathbf{b}$ to those of $f$, we first want the output neuron's weights $\mathbf{v}=\left[\mathbf{v}_{1}, \mathbf{v}_{2}\right] \in \mathbb{R}^{4} \times \mathbb{R}^{2}$ to satisfy the equality $A^{\top} W \mathbf{v}_{1}=\mathbf{d}$. To this end, we have by Lemma C. 1 that $A$ is invertible and $\left\|A^{-\top}\right\|_{\text {sp }} \leq 4$. It then follows that $\mathbf{v}_{1}=W^{-1} A^{-\top} \mathbf{d}$, which entails

$$
\begin{equation*}
\left\|\mathbf{v}_{1}\right\|_{2} \leq\left\|W^{-1}\right\|_{\mathrm{sp}}\left\|A^{-\top}\right\|_{\mathrm{sp}}\|\mathbf{d}\|_{2} \cdot \leq \frac{8 \sqrt{2}}{q m} \tag{8}
\end{equation*}
$$

where we used the separability assumption, implying that $\|\mathbf{d}\|_{2} \leq \sqrt{2} q^{-1}$ due to Item 3 which guarantees that $\beta_{3}-\beta_{2}, \beta_{4}-\beta_{3} \geq q$, and the lower bound assumption $\left|w_{i}\right| \geq m$ for all $i \in[4]$ which holds by Items 1 and 2. Next, we use the two neurons with breakpoints at $\beta_{5}, \beta_{6}$ to shift the network by a constant so that it overlaps with $f$ on the interval $[-R, R]$. To perform this shift, we first compute the magnitude by which we wish to shift which is given by the expression

$$
b_{0}:=-\sum_{j=1}^{4} v_{j} b_{j} \mathbb{1}\left\{w_{j} x+b_{j}>0 \quad \forall x \in\left(\beta_{j^{\prime}}, \beta_{j^{\prime}+1}\right)\right\} .
$$

Letting

$$
P:=\left(\begin{array}{ll}
b_{5} & b_{6} \\
w_{5} & w_{6}
\end{array}\right),
$$

we have that the neurons with breakpoints at $\beta_{5}, \beta_{6}$ compute a function which equals $b_{0}$ on the interval $\left[\beta_{5}, \beta_{6}\right]$ when the equality $P \cdot \mathbf{v}_{2}=\left(b_{0}, 0\right)^{\top}$ is satisfied. We will now compute $P^{-1}$ and show that it is well-defined. The inverse of a $2 \times 2$ matrix is given by

$$
P^{-1}=\frac{1}{b_{5} w_{6}-b_{6} w_{5}}\left(\begin{array}{cc}
w_{6} & -b_{6} \\
-w_{5} & b_{5}
\end{array}\right)=\frac{1}{\beta_{6}-\beta_{5}}\left(\begin{array}{cc}
\frac{1}{w_{5}} & -\frac{b_{6}}{w_{5} w_{6}} \\
-\frac{1}{w_{6}} & \frac{b_{5}}{w_{5} w_{6}} .
\end{array}\right)
$$

Using the above, we can upper bound the spectral norm of $P^{-1}$ by upper bounding $1 /\left(\beta_{6}-\right.$ $\beta_{5}$ ) with $\frac{1}{2 R} \leq \frac{1}{2}$ since $\beta_{5}, \beta_{6}$ are outside the interval $[-R, R]$ at opposite sides and $R \geq 1$ by Assumption 2.2, and by upper bounding the spectral norm of the matrix with its Frobenius norm by using Item 2 , to obtain

$$
\left\|P^{-1}\right\|_{\mathrm{sp}} \leq \frac{M}{m^{2}}
$$

Similarly, we derive an upper bound on $\left|b_{0}\right|$ using Cauchy-Schwartz and Items 1 and 2 to obtain

$$
\left|b_{0}\right| \leq\left\|\mathbf{v}_{1}\right\|_{2}\|\mathbf{b}\|_{2} \leq 2 M\left\|\mathbf{v}_{1}\right\|_{2} .
$$

With the above, we can bound the norm of $\mathbf{v}_{2}$ as follows

$$
\begin{equation*}
\left\|\mathbf{v}_{2}\right\|_{2}=\left\|P^{-1} \cdot\left(b_{0}, 0\right)^{\top}\right\|_{\mathrm{sp}} \leq \frac{2 M^{2}}{m^{2}}\left\|\mathbf{v}_{1}\right\|_{2} \tag{9}
\end{equation*}
$$

We now define $\mathbf{u}$ as the all-zero vector, except for the six output neuron entries corresponding to the neurons with breakpoints $\beta_{1}, \ldots, \beta_{6}$, where the coordinates of $\mathbf{u}$ take the values $y_{i} v_{1}, \ldots, y_{i} v_{6}$. Note that this entails

$$
\begin{equation*}
\|\mathbf{u}\|_{2}=\sqrt{\left\|\mathbf{v}_{1}\right\|_{2}^{2}+\left\|\mathbf{v}_{2}\right\|_{2}^{2}} \leq\left\|\mathbf{v}_{1}\right\|_{2} \sqrt{1+\frac{4 M^{4}}{m^{4}}} \leq \frac{8}{q m} \sqrt{2+\frac{8 M^{4}}{m^{4}}} \leq 8 \sqrt{10} \frac{M^{2}}{q m^{3}} \tag{10}
\end{equation*}
$$

where we have used Eq. (8) and the fact that $1<M / m$. Next, we have for all $j \in[n]$ and $h>0$ that

$$
y_{j} \Phi\left(\boldsymbol{\theta}+\frac{h}{\|\mathbf{u}\|} \mathbf{u} ; x_{j}\right)=y_{j}\left(\Phi\left(\boldsymbol{\theta} ; x_{j}\right)+\frac{h}{\|\mathbf{u}\|} y_{i} f\left(x_{j}\right)\right) .
$$

Observe that all data points satisfy $x_{j} \in\left[\beta_{5}, \beta_{6}\right]$, and are therefore unaffected by the value $\Phi(\cdot, x)$ attains for $x$ 's outside of this interval. Additionally, $f(x)=0$ for all $x \in$ $\left[\beta_{1}, \beta_{2}\right] \cup\left(\beta_{4}, \infty\right)$, which also keeps $\Phi(\cdot, x)$ unaffected by moving in the direction of $\mathbf{u}$. Moreover, since the sign of data instances in $x_{j} \in\left(\beta_{2}, \beta_{4}\right]$ is always $y_{i}$, we have that

$$
y_{j} \Phi\left(\boldsymbol{\theta}+\frac{h}{\|\mathbf{u}\|} \mathbf{u} ; x_{j}\right)=y_{j} \Phi\left(\boldsymbol{\theta} ; x_{j}\right)+\frac{h}{\|\mathbf{u}\|} f\left(x_{j}\right) \geq y_{j} \Phi\left(\boldsymbol{\theta} ; x_{j}\right)
$$

Lastly, for $x_{i}$ we have that $x_{i} \in\left[\beta_{3}, \beta_{4}\right]$, and that $x_{i}$ is at distance at least $\gamma$ from the boundary by Item 5 in our separability assumption. By Item 3, this implies that $f\left(x_{i}\right)$ is at least $\frac{\gamma}{\beta_{4}-\beta_{3}} \geq \frac{\gamma}{Q}$, which with the above equation and our bound from Eq. (10) yields

$$
\begin{equation*}
y_{i} \Phi\left(\boldsymbol{\theta}+\frac{h}{\|\mathbf{u}\|} \mathbf{u} ; x_{i}\right) \geq y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)+h \frac{\gamma}{Q\|\mathbf{u}\|} \geq y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)+h \frac{\gamma q m^{3}}{8 \sqrt{10} Q M^{2}} \tag{11}
\end{equation*}
$$

- Suppose that $x_{i} \in I_{\ell}$ for some $\ell \in\{1, \ldots, r+1\}$, such that $x_{i} \notin\left(\beta_{1}\left(I_{\ell}\right), \beta_{3}\left(I_{\ell}\right)\right)$. Assume without loss of generality that $x_{i} \leq \beta_{4}:=\beta_{1}\left(I_{\ell}\right)$ (the proof is symmetric otherwise). Then by our separability assumption, there exist $\beta_{2}:=\beta_{1}\left(I_{\ell-1}\right), \beta_{3}:=\beta_{3}\left(I_{\ell-1}\right), \beta_{5}:=\beta_{3}\left(I_{\ell}\right)$ and $\beta_{1}$ whose neuron is active on all the data points where $\beta_{1} \leq \beta_{2}$. Note that by the definition of $\beta_{i}(\cdot)$, it must hold that $\beta_{4}$ is the smallest element in the interval $I_{\ell}$ which satisfies Items 1 and 3 (since otherwise we would have that $x_{i} \in\left(\beta_{3}, \beta_{5}\right) \subseteq I_{\ell}$, which is handled in the previous case), and therefore $\beta_{3} \notin I_{\ell}$, implying that $x_{i} \in\left(\beta_{3}, \beta_{4}\right)$. Moreover, we note that we may assume that $\ell>1$, since otherwise we can take the smallest two breakpoints that are active on all the data along with $\beta_{4}$ which reduces us to the previous case. We thus denote the largest data instance in $I_{\ell-1}$ as $x_{i-1}$ which implies $x_{i-1}<x_{i}$. Lastly, Item 2 also guarantees the existence of two breakpoints that are distinct from the previous ones, which we denote by $\beta_{6}, \beta_{7}$, where $\beta_{6}<\beta_{1}<0$ and $\beta_{7}>0$ are active on all the data points.
Following a similar approach as in the previous case, we will now show the existence of a depth-2 ReLU network which consists of seven hidden neurons with weights $\mathbf{w}=$ $\left(w_{1}, \ldots, w_{7}\right)$ and biases $\mathbf{b}=\left(b_{1}, \ldots, b_{7}\right)$ and computes the piece-wise linear function


Figure 2: The plots of $f(x)$ (green) and $\Phi(\boldsymbol{\theta} ; x)+h \cdot f(x)$ (black) for various values of $h$. Moving $\boldsymbol{\theta}$ in the direction of $\mathbf{u}$ which computes $f(\cdot)$ strictly decreases the loss over the positively-labeled instances in the interval $(0,2)$ without degrading the loss over the rest of the dataset. Unlike the previous simpler case, since $\boldsymbol{\theta}$ has no breakpoints between the negative instance at $x=0$ and the positive instance at $x=0.6$, we use a function $f(x)$ which pivots around $x=0$ to prevent the prediction over the negatively-labeled instances in the interval $[-1,0]$ from increasing. Best viewed in color.

$$
\begin{aligned}
& f:\left[\beta_{1}, \beta_{7}\right) \rightarrow \mathbb{R} \text { given by } \\
& \qquad f(x):= \begin{cases}0 & x \in\left[\beta_{1}, \beta_{2}\right] \\
\frac{1}{\beta_{2}-\beta_{3}} x-\frac{\beta_{2}}{\beta_{2}-\beta_{3}} & x \in\left(\beta_{2}, \beta_{3}\right) \\
\frac{1}{x_{i-1}-\beta_{3}} x-\frac{x_{i-1}}{x_{i-1}-\beta_{3}} & x \in\left[\beta_{3}, \beta_{4}\right], \\
\frac{\beta_{4}-x_{i-1}}{x_{i-1}-\beta_{3}}\left(\frac{1}{\beta_{4}-\beta_{5}} x-\frac{\beta_{5}}{\beta_{4}-\beta_{5}}\right) & x \in\left(\beta_{4}, \beta_{5}\right) \\
0 & x \in\left[\beta_{5}, \beta_{7}\right)\end{cases}
\end{aligned}
$$

where the first five neurons are used to compute the slopes of $f$ and the remaining last two neurons are used to simulate a bias term to shift the network to accord with $f$ in its domain (see Figure 2 for an illustration).
Define $W:=\operatorname{diag}\left(w_{1}, \ldots, w_{5}\right) \in \mathbb{R}^{5 \times 5}$, define the masking matrix $A \in \mathbb{R}^{5 \times 5}$ with entries $a_{j, j^{\prime}}=\mathbb{1}\left\{j\right.$-th neuron is active on the interval starting with $\left.\beta_{j^{\prime}}\right\}$ and let

$$
\mathbf{d}=\left(0, \frac{1}{\beta_{2}-\beta_{3}}, \frac{1}{x_{i-1}-\beta_{3}}, \frac{\beta_{4}-x_{i-1}}{x_{i-1}-\beta_{3}} \cdot \frac{1}{\beta_{4}-\beta_{5}}, 0\right)
$$

Thus, to match the slopes computed by the a depth-2 ReLU network with weights $\mathbf{w}, \mathbf{b}$ to those of $f$, we first want the output neuron's weights $\mathbf{v}=\left[\mathbf{v}_{1}, \mathbf{v}_{2}\right] \in \mathbb{R}^{5} \times \mathbb{R}^{2}$ to satisfy the equality $A^{\top} W \mathbf{v}_{1}=\mathbf{d}$. To this end, we have by Lemma C. 1 that $A$ is invertible and $\left\|A^{-\top}\right\|_{\text {sp }} \leq 5$. It then follows that $\mathbf{v}_{1}=W^{-1} A^{-\top} \mathbf{d}$, which entails

$$
\left\|\mathbf{v}_{1}\right\|_{2} \leq\left\|W^{-1}\right\|_{\mathrm{sp}}\left\|A^{-\top}\right\|_{\mathrm{sp}}\|\mathbf{d}\|_{2} \leq 5\left\|W^{-1}\right\|_{\mathrm{sp}}\|\mathbf{d}\|_{2}
$$

To upper bound the above, we first bound $\|\mathbf{d}\|_{2}$. By our separability assumption we have

$$
\begin{aligned}
\|\mathbf{d}\|_{2} & \leq \sqrt{\frac{1}{q^{2}}+\frac{1}{\left(x_{i-1}-\beta_{3}\right)^{2}}+\frac{4 Q^{2}}{q^{2}\left(x_{i-1}-\beta_{3}\right)^{2}}} \\
& =\sqrt{\frac{q^{2}+\left(x_{i-1}-\beta_{3}\right)^{2}+4 Q^{2}}{q^{2}\left(x_{i-1}-\beta_{3}\right)^{2}} \leq \frac{1}{q\left(x_{i-1}-\beta_{3}\right)} \sqrt{8 Q^{2}+q^{2}}}
\end{aligned}
$$

where we used Item 3 to upper bound the denominators, and Item 4 which entails $\beta_{4}-$ $x_{i-1}, x_{i-1}-\beta_{3} \leq \beta_{4}-\beta_{3} \leq Q$. By Item 1 , we have $\left|w_{i}\right| \geq m$ for all $i \in[5]$, and thus we obtain

$$
\begin{equation*}
\left\|\mathbf{v}_{1}\right\|_{2} \leq \frac{5}{m q\left(x_{i-1}-\beta_{3}\right)} \sqrt{40 Q^{2}+5 q^{2}} \leq \frac{15 \sqrt{5} Q}{m q\left(x_{i-1}-\beta_{3}\right)} \tag{12}
\end{equation*}
$$

Now, similarly to the previous case, we wish to shift the function computed by the depth-2 ReLU network by a constant. This is done in the exact same manner as in the previous case, where we shift it by a magnitude given by

$$
b_{0}=-\sum_{j=1}^{5} v_{j} b_{j} \mathbb{1}\left\{w_{j} x+b_{j}>0 \quad \forall x \in\left(\beta_{j^{\prime}}, \beta_{j^{\prime}+1}\right)\right\}
$$

Bounding $\left|b_{0}\right|$ using its above definition, Cauchy-Schwartz and Item 1 in our separability assumption, we obtain

$$
\left|b_{0}\right| \leq\left\|\mathbf{v}_{1}\right\|_{2}\|\mathbf{b}\|_{2} \leq \sqrt{5} M\left\|\mathbf{v}_{1}\right\|_{2}
$$

From the above and Eq. (9), we can shift the network by the desired magnitude using a vector $\mathbf{v}_{2}=\left(v_{6}, v_{7}\right)$ satisfying

$$
\left\|\mathbf{v}_{2}\right\|_{2} \leq \frac{\sqrt{5} M^{2}}{m^{2}}\left\|\mathbf{v}_{1}\right\|_{2}
$$

We now define $\mathbf{u}$ as the all-zero vector, except for the output neuron entries corresponding to the neurons with breakpoints $\beta_{1}, \ldots, \beta_{7}$, where the coordinates of $\mathbf{u}$ take the values $y_{i} v_{1}, \ldots, y_{i} v_{7}$. Note that this entails

$$
\begin{equation*}
\|\mathbf{u}\|_{2}=\sqrt{\left\|\mathbf{v}_{1}\right\|_{2}^{2}+\left\|\mathbf{v}_{2}\right\|_{2}^{2}} \leq\left\|\mathbf{v}_{1}\right\|_{2} \sqrt{1+\frac{5 M^{4}}{m^{4}}} \leq \frac{15 \sqrt{30} Q M^{2}}{m^{3} q\left(x_{i-1}-\beta_{3}\right)} \leq \frac{90 Q M^{2}}{m^{3} q\left(x_{i-1}-\beta_{3}\right)} \tag{13}
\end{equation*}
$$

where we used Eq. (12) and the fact that $1<M / m$. We therefore have for all $j \in[n]$ and $h>0$ that

$$
y_{j} \Phi\left(\boldsymbol{\theta}+\frac{h}{\|\mathbf{u}\|} \mathbf{u} ; x_{j}\right)=y_{j}\left(\Phi\left(\boldsymbol{\theta} ; x_{j}\right)+\frac{h}{\|\mathbf{u}\|} y_{i} f\left(x_{j}\right)\right) .
$$

Observe that all data points satisfy $x_{j} \in\left[\beta_{6}, \beta_{7}\right]$, and are therefore unaffected by the value $\Phi(\cdot, x)$ attains for $x$ 's outside of this interval. Additionally, $f(x)=0$ for all $x \in$ $\left[\beta_{1}, \beta_{2}\right] \cup\left[\beta_{5}, \infty\right)$, which also keeps $\Phi(\cdot, x)$ unaffected by moving in the direction of $\mathbf{u}$. In the interval $\left(\beta_{2}, x_{i-1}\right)$, the sign of $\Phi(\cdot, x)$ is $-y_{i}$, and in the interval $\left(x_{i-1}, \beta_{5}\right)$ its sign changes to $y_{i}$. For this reason, similarly to the previous case, we have that

$$
\begin{equation*}
y_{j} \Phi\left(\boldsymbol{\theta}+\frac{h}{\|\mathbf{u}\|} \mathbf{u} ; x_{j}\right)=y_{j} \Phi\left(\boldsymbol{\theta} ; x_{j}\right)+\frac{h}{\|\mathbf{u}\|} f\left(x_{j}\right) \geq y_{j} \Phi\left(\boldsymbol{\theta} ; x_{j}\right) \tag{14}
\end{equation*}
$$

for all $x_{j} \in\left(\beta_{2}, \beta_{5}\right)$. Lastly, for $x_{i}$ we have that $x_{i} \in\left[\beta_{3}, \beta_{4}\right]$, and that $x_{i}$ is at distance at least $\gamma$ from $x_{i-1}$ by Item 5 in our separability assumption. This implies that

$$
f\left(x_{i}\right)=\frac{x_{i}}{x_{i-1}-\beta_{3}}-\frac{x_{i-1}}{x_{i-1}-\beta_{3}} \geq \frac{\gamma}{x_{i-1}-\beta_{3}}
$$

and therefore $h f\left(x_{i}\right) \geq \frac{\gamma}{x_{i-1}-\beta_{3}} h$, which with Eqs. $(13,14)$ implies that

$$
\begin{equation*}
y_{i} \Phi\left(\boldsymbol{\theta}+\frac{h}{\|\mathbf{u}\|} \mathbf{u} ; x_{i}\right) \geq y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)+h \frac{\gamma}{\|\mathbf{u}\|\left(x_{i-1}-\beta_{3}\right)} \geq y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)+h \frac{\gamma q m^{3}}{90 Q M^{2}} \tag{15}
\end{equation*}
$$

which is a weaker lower bound than the one derived in Eq. (11), and thus always holds if Eq. (11) is satisfied.

We now turn to lower bound the norm of the gradient by analyzing the directional derivative of $\mathcal{L}(\cdot)$ in the direction of the vector $\mathbf{u}$ defined by the above two cases. To this end, denote $\overline{\mathbf{u}}=\frac{\mathbf{u}}{\|\mathbf{u}\|}$ and compute

$$
\begin{align*}
\|\nabla \mathcal{L}(\boldsymbol{\theta})\|_{2} & \geq|\langle\nabla \mathcal{L}(\boldsymbol{\theta}), \overline{\mathbf{u}}\rangle|=\left|\lim _{h \rightarrow 0} \frac{1}{h n} \sum_{j=1}^{n}\left(\ell\left(y_{j} \Phi\left(\boldsymbol{\theta}+h \overline{\mathbf{u}} ; x_{j}\right)\right)-\ell\left(y_{j} \Phi\left(\boldsymbol{\theta} ; x_{j}\right)\right)\right)\right| \\
& =\lim _{h \rightarrow 0} \frac{1}{h n} \sum_{j=1}^{n}\left(\ell\left(y_{j} \Phi\left(\boldsymbol{\theta} ; x_{j}\right)\right)-\ell\left(y_{j} \Phi\left(\boldsymbol{\theta}+h \overline{\mathbf{u}} ; x_{j}\right)\right)\right) \\
& \geq \lim _{h \rightarrow 0} \frac{1}{h n}\left(\ell\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right)-\ell\left(y_{i} \Phi\left(\boldsymbol{\theta}+h \overline{\mathbf{u}} ; x_{i}\right)\right)\right) \\
& \geq \lim _{h \rightarrow 0} \frac{1}{h n}\left(\ell\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right)-\ell\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)+h \frac{\gamma q m^{3}}{90 Q M^{2}}\right)\right) \\
& =-\frac{\gamma q m^{3}}{90 n Q M^{2}} \ell^{\prime}\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right) . \tag{16}
\end{align*}
$$

In the above, the first inequality is by Cauchy-Schwartz; the second equality is due to Eq. (14), which guarantees that each summand is non-positive; the second inequality is another application of Eq. (14) which guarantees that we're omitting only non-negative terms; and the last inequality is due to Eq. (15). We now lower bound the above expression depending on whether $\ell(\cdot)$ is the exponential or the logistic loss.

First assume that $\ell(\cdot)$ is the exponential loss. Then we have that $-\ell^{\prime}(x)=\ell(x)$ for all $x \in \mathbb{R}$, and therefore we can directly use the inequality $\ell\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right)>1 / 2 n$.
In the case where $\ell(\cdot)$ is the logistic loss, we have that

$$
-\ell^{\prime}(x)=\frac{1}{1+\exp (x)}=1-\frac{1}{1+\exp (-x)}
$$

By the inequality $\ell\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right)>1 / 2 n$, we have $1+\exp \left(-y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right)>\exp (1 / 2 n)$, implying that

$$
-\ell^{\prime}\left(y_{i} \Phi\left(\boldsymbol{\theta} ; x_{i}\right)\right)>1-\exp \left(-\frac{1}{2 n}\right) \geq \frac{1}{4 n}
$$

where we used the inequality $\exp (-x) \leq 1-0.5 x$ which holds for all $x \in[0,1]$. Combining both loss cases and Eq. (16), we arrived at

$$
\|\nabla \mathcal{L}(\boldsymbol{\theta})\|_{2} \geq \frac{\gamma q m^{3}}{360 n^{2} Q M^{2}}
$$

Squaring the above and dividing by 2 , the lemma follows.
The following proposition establishes the separability (Def. C.1) of a neighborhood in weight space around our initialization point (Def. C.2) from the dataset $S$.
Proposition C. 1 (Bounded Gradient with High Probability). Under Assumptions 2.1 and 2.2, given any $\delta \in(0,1)$, suppose that the following hold

$$
\begin{equation*}
k \geq 6144 \cdot \frac{R^{4} \log \left(\frac{24 r}{\delta}\right)}{\rho} \quad \text { and } \quad \Delta=\frac{\delta \rho \sigma_{h}}{24 n k C R^{3}} \tag{17}
\end{equation*}
$$

Then with probability at least $1-\delta$, for all $\boldsymbol{\theta} \in U_{\Delta}(\boldsymbol{\theta}(0)) \cap L_{1 / 2 n}^{+}(\mathcal{L})$, we have that

$$
\frac{1}{2}\|\nabla \mathcal{L}(\boldsymbol{\theta})\|_{2}^{2} \geq 3 \cdot 10^{-11} \frac{\delta^{2} \rho^{2}}{n^{6} r^{2} C^{2} R^{8}} \sigma_{h}^{2}
$$

Proof. To prove the proposition, we will show that for all $\boldsymbol{\theta} \in U_{\Delta}(\boldsymbol{\theta}(0)) \cap L_{1 / 2 n}^{+}(\mathcal{L}), \boldsymbol{\theta}$ is separable from $S$ with high probability. By Lemma C.2, this would imply the proposition. We will begin with proving items 1-4 jointly, and then we will show item 5 separately.

1-4. Under Assumption 2.1, suppose we are given some $a \in[0, R]$ and $\boldsymbol{\xi}:=\left(\xi_{1}, \xi_{2}\right) \in\{-1,1\}^{2}$, and sufficiently small $\varepsilon>0$ such that

$$
\begin{equation*}
\varepsilon \leq \frac{512 R^{4} \log \left(\frac{24 r}{\delta}\right)}{k} \leq \frac{\rho}{6} \leq \frac{R}{6} \tag{18}
\end{equation*}
$$

where the first inequality is by assumption on $\varepsilon$, the second inequality is by the lower bound on $k$ in Eq. (17), and the last inequality follows from $\rho \leq \frac{2 R}{r+1} \leq R$ since $\rho$ must be smaller than the average length of an interval and since we assume $r \geq 1$. We now consider the event denoted by $E_{a, \varepsilon, \xi}$ where the weights $w_{i}, b_{i}$ of the $i$-th neuron satisfy

$$
\begin{equation*}
w_{i} \in\left(\xi_{1} \frac{\sigma_{\mathrm{h}}}{R}, \xi_{1} \frac{2 \sigma_{\mathrm{h}}}{R}\right) \quad \text { and } \quad-\frac{b_{i}}{w_{i}} \in\left(\xi_{2} a, \xi_{2}(a+\varepsilon)\right) \tag{19}
\end{equation*}
$$

Since such an event is symmetric about 0 , it is unaffected by the signs of $\boldsymbol{\xi}$. We can therefore assume without loss of generality that both intervals in Eq. (19) are contained in the positive real line and omit $\boldsymbol{\xi}$ from our notation. Under this assumption, the probability of $E_{a, \varepsilon}$ can be given in terms of Owen's T function which is defined by

$$
T(h, a):=\frac{1}{2 \pi} \int_{0}^{a} \frac{\exp \left(-\frac{1}{2} h^{2}\left(1+x^{2}\right)\right)}{1+x^{2}} d x
$$

(see Owen [1956]), yielding

$$
\mathbb{P}\left[E_{a, \varepsilon}\right]=T\left(\frac{1}{R}, a+\varepsilon\right)-T\left(\frac{1}{R}, a\right)-\left(T\left(\frac{2}{R}, a+\varepsilon\right)-T\left(\frac{2}{R}, a\right)\right) .
$$

Using the definition of $T(\cdot, \cdot)$, the above can be simplified to

$$
\begin{align*}
\mathbb{P}\left[E_{a, \varepsilon}\right] & =\frac{1}{2 \pi} \int_{a}^{a+\varepsilon} \frac{\exp \left(-\frac{1}{2 R^{2}}\left(1+x^{2}\right)\right)\left(1-\exp \left(-\frac{3}{2 R^{2}}\left(1+x^{2}\right)\right)\right)}{1+x^{2}} d x \\
& \geq \frac{1}{2 \pi}\left(1-\exp \left(-\frac{3}{2 R^{2}}\right)\right) \int_{a}^{a+\varepsilon} \frac{\exp \left(-\frac{1}{2 R^{2}}\left(1+x^{2}\right)\right)}{1+x^{2}} d x \\
& \geq \frac{3}{8 \pi R^{2}} \int_{a}^{a+\varepsilon} \frac{\exp \left(-\frac{1}{2 R^{2}}\left(1+x^{2}\right)\right)}{1+x^{2}} d x \\
& \geq \frac{3}{8 \pi R^{2}} \int_{a}^{a+\varepsilon} \frac{\exp \left(-\frac{1}{2 R^{2}}\left(1+(R+\varepsilon)^{2}\right)\right)}{1+(R+\varepsilon)^{2}} d x \geq \frac{\varepsilon}{512 R^{4}} . \tag{20}
\end{align*}
$$

In the above, the second inequality follows from the inequality $1-\exp (-x) \geq 0.5 x$ which holds for all $x \in[0,1.5]$ and from the fact that $1 \leq R$; the third inequality follows from the fact that the integrand is a monotonically decreasing function and $|a| \leq R$; and the last inequality follows from $\varepsilon \leq R$ which is implied by Eq. (18) and allows us to lower bound the numerator of the integrand by $\exp (-2.5)$ and upper bound the denominator by $5 R^{2}$, and the fact that $3 /(40 \pi \exp (2.5)) \geq 1 / 512$.

Next, given some interval $I_{j}:=\left(x_{i_{j}}, x_{i_{j+1}+1}\right), j \in[r+1]$, where the classification does not change signs on the data, we consider the three sub-intervals given by

$$
\begin{aligned}
I_{j_{1}} & :=\left(x_{i_{j}}, x_{i_{j}}+\varepsilon\right) \\
I_{j_{2}} & :=\left(\frac{x_{i_{j}}+x_{i_{j+1}+1}}{2}-\frac{\varepsilon}{2}, \frac{x_{i_{j}}+x_{i_{j+1}+1}}{2}+\frac{\varepsilon}{2}\right), \\
I_{j_{3}} & :=\left(x_{i_{j+1}+1}-\varepsilon, x_{i_{j+1}+1}\right) .
\end{aligned}
$$

We remark that due to Eq. (18), the above sub-intervals are all disjoint and the distance between the intervals is positive. We now wish to show that Items 1 and 3 hold. We have from Eq. (20) that the probability that a given sub-interval of length $\varepsilon$ contains no breakpoint is at most

$$
\left(1-\mathbb{P}\left[E_{a, \varepsilon}\right]\right)^{k} \leq\left(1-\frac{\log \left(\frac{24 r}{\delta}\right)}{k}\right)^{k} \leq \exp \left(-\log \left(\frac{24 r}{\delta}\right)\right)=\frac{\delta}{24 r}
$$

where we used the inequality $(1-x / y)^{y} \leq \exp (-x)$ which holds for all $x, y>0$. There are exactly $3 \cdot(r+1) \leq 6 r$ sub-intervals, therefore by a union bound we have that Items 1 and 3 hold for some positive $q, Q, m, M$ with probability at least $1-\frac{\delta}{4}$.
Next, we show Item 2. By Eq. (20), we have

$$
\mathbb{P}\left[E_{R, R / 6}\right] \geq \frac{1}{3072 R^{3}}
$$

Thus, the probability of initializing a neuron with breakpoint in $\left(R, \frac{7}{6} R\right)$ which is active on all the data points is at least $\frac{1}{6144 R^{3}}$, since there's an independent 0.5 probability that it has the correct orientation. This entails that the probability of initializing at most one neuron which is active on all the data points and has a breakpoint in $\left(R, \frac{7}{6} R\right)$ is upper bounded by

$$
\begin{gathered}
\left(1-\mathbb{P}\left[E_{R, R / 6}\right]\right)^{k}+k\left(1-\mathbb{P}\left[E_{R, R / 6}\right]\right)^{k-1} \mathbb{P}\left[E_{R, R / 6}\right] \leq 2\left(1-\frac{\mathbb{P}\left[E_{R, R / 6}\right]}{2}\right)^{k} \\
\leq 2\left(1-\frac{1}{6144 R^{3}}\right)^{k} \leq 2 \exp \left(-\frac{k}{6144 R^{3}}\right) \\
\leq 2 \exp \left(-\frac{R \log (24 r / \delta)}{\rho}\right) \leq \frac{2 \delta}{24 r} \leq \frac{\delta}{12}
\end{gathered}
$$

In the above, the first inequality follows from the inequality $(1-x)^{k}+k(1-x)^{k-1} x \leq$ $2(1-x / 2)^{k}$ which holds for any natural $k$ and all $x \in[0,1],{ }^{6}$ the third inequality follows from $(1-1 / x)^{x} \leq \exp (-1)$ for all $x>0$, the fourth inequality follows from our lower bound on $k$ in Eq. (17), and the penultimate inequality holds due to Eq. (18) which entails $\rho \leq R$. Therefore, by the above and a union bound on the symmetric event where two neurons are initialized in $\left(-\frac{7}{6} R,-R\right)$, we have that Item 2 holds for some $m, M$ with probability at least $1-\frac{\delta}{6} \geq 1-\frac{\delta}{4}$.
We will now derive explicit bounds on the constants $m, M, q, Q$, and in addition we will show that Item 4 holds. Applying a union bound on the two previous cases, we have that Items 1-3 hold with probability at least $1-\delta / 2$. In such a case, we get an explicit lower bound on $q$ as follows

$$
q \geq \frac{\left|I_{j}\right|-3 \varepsilon}{2} \geq \frac{\rho-3 \varepsilon}{2} \geq \frac{\rho}{4}
$$

where in the second inequality we used the fact that $\left|I_{j}\right| \geq \rho$ for all $j \in[r+1]$ which holds by the definition of $\rho$ and in the last inequality we used Eq. (18). To bound $Q$ in Item 4 , we first argue that under the realization of $E_{R, R / 6}$, the four neurons that are active on all the data points have a breakpoint with absolute value at most $\frac{7}{6} R \leq 2 R$. To upper bound $Q$ in Item 3, observe that under the realization of the previous events we have that

$$
\beta_{i+1}\left(I_{j}\right)-\beta_{i}\left(I_{j}\right) \leq \frac{\left|I_{j}\right|+\varepsilon}{2} \leq R+\frac{\varepsilon-\rho}{2} \leq R
$$

for all $j \in[r+1]$ and $i \in[2]$, which follows from $\left|I_{j}\right| \leq 2 R-\rho$ since $r \geq 1$ and $I_{j} \subseteq[-R, R]$ (i.e. there exists at least one interval other than $I_{j}$ which has length at least $\rho$ ), and from the inequality $\varepsilon \leq \rho / 6$ which holds by Eq. (18). To upper bound $Q$ in Item 4 , we bound the term $\left|\beta_{1}\left(I_{j+1}\right)-\beta_{3}\left(I_{j}\right)\right|$. Observe that under the realization of the above events we have that $\beta_{1}\left(I_{j+1}\right) \in(x, x+\varepsilon)$ and $\beta_{3}\left(I_{j}\right) \in\left(x^{\prime}-\varepsilon, x^{\prime}\right)$ where $x<x^{\prime}$ are the largest and smallest data instances in $I_{j}, I_{j+1}$, respectively. We therefore have

$$
\left|\beta_{1}\left(I_{j+1}\right)-\beta_{3}\left(I_{j}\right)\right| \leq \max \left\{\left|x-x^{\prime}\right|,\left|x-x^{\prime}+2 \varepsilon\right|\right\} \leq\left|x-x^{\prime}\right|+2 \varepsilon \leq \frac{7}{3} R
$$

where the second inequality follows from the triangle inequality and the last inequality follows from the fact that $x_{i} \in[-R, R]$ for all $i \in[n]$ and from Eq. (18). Turning to bound

[^1]$w_{i}, b_{i}$, we have by Eqs. $(18,19)$ that when $E_{a, \varepsilon}$ or $E_{R, R / 6}$ occur then the $i$-th neuron satisfies $\left|w_{i}\right| \geq \sigma / R$ and
$$
\left|b_{i}\right| \leq\left(R+\frac{1}{6} R\right) w_{i} \leq \frac{7}{6} R \cdot \frac{2 \sigma_{\mathrm{h}}}{R}=\frac{7}{3} \sigma_{\mathrm{h}}
$$
concluding the derivation of Items 1-4.
5. It will suffice to lower bound the probability of the event denoted by $A$ where $x_{i} \notin$ $\left(\alpha_{j}-\gamma / 2, \alpha_{j}+\gamma / 2\right)$ for all $i \in[n]$ and $j \in[r]$, where $\alpha_{j}$ is the $j$-th sign change of the ground truth function labelling $y_{i}$. The set $\cup_{j=1}^{r}\left(\alpha_{j}-\gamma / 2, \alpha_{j}+\gamma / 2\right)$ has Lebesgue measure of at most $\gamma r$, and since by Assumption 2.2 we have that $\mu(x) \leq C$ for all $x \in[-R, R]$, we lower bound the probability of the event by the expression
$$
\mathbb{P}[A] \geq(1-\gamma r C)^{n}
$$

Plugging $\gamma=\frac{\delta}{4 n r C}$ in the above which entails $\gamma r C \leq 1$ and using Bernoulli's inequality we have

$$
\mathbb{P}[A] \geq 1-\frac{\delta}{4}
$$

To conclude the derivation so far, using another union bound, we have shown that with probability at least $1-0.75 \delta, \boldsymbol{\theta}(0)$ is separable from $S$ with constants

$$
\begin{equation*}
\gamma_{0}=\frac{\delta}{4 n r C}, \quad q_{0}=\frac{\rho}{4}, \quad Q_{0}=\frac{7}{3} R, \quad m_{0}=\frac{\sigma_{\mathrm{h}}}{R}, \quad M_{0}=\frac{7}{3} \sigma_{\mathrm{h}} . \tag{21}
\end{equation*}
$$

We will now show that the separability also holds in a $\Delta$-hidden neighborhood of $\boldsymbol{\theta}(0)$ for an appropriately chosen $\Delta>0$. To this end, we first establish that

$$
\mathbb{P}\left[\min _{i \in[n], j \in[k]}\left|\beta_{j}-x_{i}\right|>\frac{\delta}{8 n k C}\right] \geq 1-\frac{\delta}{4}
$$

Suppose we have $n$ data instances in $[-R, R]$, then a cover of radius $\frac{\delta \pi}{8 k n}$ over these data points has a (one-dimensional Lebesgue) measure of at most $\frac{\delta \pi}{4 k}$. Thus, the probability that a breakpoint will not be initialized within distance less than $\frac{\delta \pi}{8 k n}$ from any point is at least $\left(1-\frac{\delta}{4 k}\right)$. This is true since Assumption 2.1 implies that the distribution of a breakpoint is a standard Cauchy distribution with density at most $1 / \pi$. We thus have that

$$
\begin{equation*}
\mathbb{P}\left[\min _{i \in[n], j \in[k]}\left|\beta_{j}-x_{i}\right|>\frac{\delta}{8 n k C}\right] \geq\left(1-\frac{\delta}{4 k}\right)^{k} \geq 1-\frac{\delta}{4}, \tag{22}
\end{equation*}
$$

where the last inequality follows from Bernoulli's inequality. A final union bound now implies that the above bound holds with the previous implications with probability at least $1-\delta$. Define

$$
\Delta:=\frac{\delta \rho \sigma_{\mathrm{h}}}{24 n k C R^{3}}
$$

we will now show that this implies the uniform separability of any $\boldsymbol{\theta} \in U_{\Delta}(\boldsymbol{\theta}(0))$ from $S$, by proving Items 1-4 jointly and Item 5 separately.

1-4. First, by Assumption 2.2, we have $1=\int_{-R}^{R} \mu(x) d x \leq 2 R C$ which with Eq. (18) implies that $\Delta \leq \frac{\delta \sigma_{\mathrm{h}}}{24 n k R} \leq \frac{\sigma_{\mathrm{h}}}{24 R}$. Since the weight and bias of each neuron in $U_{\Delta}(\boldsymbol{\theta}(0))$ change by at most $\Delta$, we have

$$
m \geq \frac{\sigma_{\mathrm{h}}}{R}-\frac{\sigma_{\mathrm{h}}}{24 R}=\frac{23 \sigma_{\mathrm{h}}}{24 R} \quad \text { and } \quad M \leq \frac{7}{3} \sigma_{\mathrm{h}}+\frac{\sigma_{\mathrm{h}}}{24 R} \leq \frac{57}{24} \sigma_{\mathrm{h}}
$$

To bound $q$ and $Q$, we will first show that under our assumptions the breakpoints cannot move much. To this end, we show that for each neuron, the function $f(w, b):=-\frac{b}{w}$ is Lipschitz on $U_{\Delta}(\boldsymbol{\theta}(0))$. We have

$$
\nabla f(w, b)=\left(\frac{b}{w^{2}},-\frac{1}{w}\right)
$$

and therefore for any neuron $(w, b) \in \boldsymbol{\theta}$ such that $\boldsymbol{\theta} \in U_{\Delta}(\boldsymbol{\theta}(0))$ we get

$$
\|\nabla f(w, b)\|=\sqrt{\frac{b^{2}}{w^{4}}+\frac{1}{w^{2}}} \leq \sqrt{\frac{M^{2}}{m^{4}}+\frac{1}{m^{2}}} \leq \frac{24 R}{23 \sigma_{\mathrm{h}}} \sqrt{\frac{57^{2}}{23^{2}} R^{2}+1}<\frac{3 R^{2}}{\sigma_{\mathrm{h}}}
$$

where the last inequality follows from $1 \leq R^{2}$. This implies that

$$
\left|-\frac{b}{w}+\frac{b_{0}}{w_{0}}\right| \leq\|\nabla f(w, b)\| \cdot\left\|(w, b)-\left(w_{0}, b_{0}\right)\right\|<\frac{3 R^{2}}{\sigma_{\mathrm{h}}} \Delta \leq \frac{\delta \rho}{8 n k C R} .
$$

That is, we have that the breakpoint of each neuron moves a distance strictly less than $\frac{\delta \rho}{8 n k C R} \leq \frac{\delta}{8 n k C}$, which along with Eq. (22) guarantees that $\mathcal{L}(\cdot)$ is differentiable on $U_{\Delta}(\boldsymbol{\theta}(0))$ since no ReLU crosses a data instance. Since $Q$ is the upper bound on the difference between two breakpoints where each moves by at most $\frac{\delta}{8 n k C}$, this also yields a bound on $Q$ as follows

$$
Q \leq Q_{0}+2 \frac{\delta}{8 n k C} \leq \frac{7}{3} R+\frac{1}{8} R \leq 2.5 R
$$

where we used the upper bound on $Q_{0}$ from Eq. (21), Eq. (17) which implies $k \geq 4$ (since $\rho \leq R$ by Eq. (18)), and $1 / C \leq 2 R$. Likewise, to lower bound $q$, compute

$$
q \geq q_{0}-2 \frac{\delta \rho}{8 n k C R} \geq \frac{\rho}{4}-\frac{\rho}{20 R} \geq \frac{\rho}{5}
$$

where again we used Eq. (21), Eq. (17) which implies $k \geq 10$, and $1 / C \leq 2 R$.
5. Since $\gamma$ depends on $S$ and not on $\boldsymbol{\theta}$, it remains unchanged and we have $\gamma=\gamma_{0}$.

We can now use the assumption $\boldsymbol{\theta} \in L_{1 / 2 n}^{+}(\mathcal{L})$ and Lemma C. 2 to conclude

$$
\frac{1}{2}\|\nabla \mathcal{L}(\boldsymbol{\theta})\|_{2}^{2} \geq \frac{\gamma^{2} q^{2} m^{6}}{259200 n^{4} Q^{2} M^{4}} \geq \frac{1}{259200 n^{4}} \cdot \frac{\delta^{2}}{4^{2} n^{2} r^{2} C^{2}} \cdot \frac{4 \rho^{2}}{25^{2} R^{2}} \cdot \frac{23^{6}}{57^{4} 24^{2} R^{6}} \cdot \sigma_{\mathrm{h}}^{2}
$$

Simplifying the above, the proposition follows.

Having established the required machinery for proving Theorem 3.1, we now turn to do so.
Proof of Theorem 3.1. We begin with bounding the loss upon initialization with high probability. First, consider $3 k$ i.i.d. random variables $X_{j} \sim \mathcal{N}(0,1)$. We have that

$$
\mathbb{P}\left[\max _{j \in[3 k]}\left|X_{j}\right| \leq x\right]=\left(\operatorname{erf}\left(\frac{x}{\sqrt{2}}\right)\right)^{3 k} \geq\left(1-\exp \left(-0.5 x^{2}\right)\right)^{3 k} \geq 1-3 k \exp \left(-0.5 x^{2}\right)
$$

where the first inequality follows from $1-\operatorname{erf}(x)<\exp \left(-x^{2}\right)$ for all $x \geq 0$ (see Eq. (7.8.3) in DLMF) and the second inequality follows from Bernoulli's inequality since $\exp \left(-0.5 x^{2}\right)<1$. Plugging $x=\sqrt{2 \log (6 k / \delta)}$ in the above, we have

$$
\mathbb{P}\left[\max _{j \in[3 k]}\left|X_{j}\right| \leq \sqrt{2 \log (6 k / \delta)}\right] \geq 1-\frac{3 k \delta}{6 k}=1-\frac{\delta}{2}
$$

Thus, with probability at least $1-\frac{\delta}{2}$, we have that all the weights of $\boldsymbol{\theta}(0)$ are at most $\sqrt{2 \log (6 k / \delta)}$ standard deviations away from zero. With this bound, we can derive for all $x \in[-R, R]$

$$
\mathcal{N}_{\boldsymbol{\theta}(0)}(x) \leq \sum_{j \in[k]}\left|v_{j}\right| \sigma\left(\left|w_{j}\right| \cdot|x|+\left|b_{j}\right|\right) \leq 4 k R \sigma_{\mathrm{h}} \sigma_{\mathrm{o}} \log \left(\frac{6 k}{\delta}\right)
$$

which for both the exponential and logistic losses implies

$$
\begin{equation*}
\mathcal{L}(\boldsymbol{\theta}(0)) \leq \exp \left(4 k R \sigma_{\mathrm{h}} \sigma_{\mathrm{o}} \log \left(\frac{6 k}{\delta}\right)\right) \leq e \tag{23}
\end{equation*}
$$

where the last inequality is by our assumption $\sigma_{\mathrm{o}} \leq \frac{1}{4 k R \sigma_{\mathrm{h}} \log \left(\frac{6 k}{\delta}\right)}$. Letting

$$
\lambda:=10^{-11} \frac{\delta^{2} \rho^{2}}{n^{6} r^{2} C^{2} R^{8}}
$$

and observing that our lower bound assumption on $\sigma_{\mathrm{h}}$ in Eq. (3) implies it's at least 0.5 since $\rho \leq R$ and $C \geq 1 / 2 R$, we can invoke Proposition C. 1 with confidence $\frac{\delta}{2}$ to obtain

$$
\begin{align*}
\frac{1}{2}\|\nabla \mathcal{L}(\boldsymbol{\theta}(t))\|_{2}^{2} & \geq 3 \cdot 10^{-11} \frac{\delta^{2} \rho^{2}}{n^{6} r^{2} C^{2} R^{8}} \sigma_{\mathrm{h}}^{2} \\
& \geq 3 \cdot 10^{-11} \frac{\delta^{2} \rho^{2}}{n^{6} r^{2} C^{2} R^{8}} \sigma_{\mathrm{h}}^{2} \cdot \frac{\mathcal{L}(\boldsymbol{\theta}(t))}{\mathcal{L}(\boldsymbol{\theta}(0))} \\
& \geq \lambda \sigma_{\mathrm{h}}^{2} \cdot \mathcal{L}(\boldsymbol{\theta}(t)), \tag{24}
\end{align*}
$$

where the second inequality holds since $\frac{\mathcal{L}(\boldsymbol{\theta}(t))}{\mathcal{L}(\boldsymbol{\theta}(0))} \leq 1$ because the flow is non-increasing, and the last inequality holds due to Eq. (23) which implies $\frac{1}{\mathcal{L}(\boldsymbol{\theta}(0))} \geq \exp (-1) \geq \frac{1}{3}$. By a union bound, the above holds with probability at least $1-\delta$.
Denote $D:=U_{\Delta}(\boldsymbol{\theta}(0)) \cap L_{\frac{1}{2 n}}^{+}(\mathcal{L})$ where $\Delta$ is defined in Eq. (17), and define $t^{\prime} \in[0, \infty)$ to be the smallest time such that $\boldsymbol{\theta}\left(t^{\prime}\right)$ is on the boundary of $U_{\Delta}(\boldsymbol{\theta}(0))$ (where $t^{\prime}=\infty$ if there exists no such time). We will now show that the flow attains loss at most $\frac{1}{2 n}$ in time $t_{0}:=\frac{\log (2 n \mathcal{L}(\boldsymbol{\theta}(0)))}{2 \lambda \sigma_{\mathrm{h}}^{2}}$, by analyzing several different cases.

- Suppose that $t^{\prime}>t_{0}$.
- If $\left\{\boldsymbol{\theta}(t): t \in\left[0, t_{0}\right]\right\} \subseteq D$, then by the PL-condition shown in Eq. (24) we have for all $t \in\left[0, t_{0}\right]$ that GF enjoys a convergence rate of

$$
\mathcal{L}(\boldsymbol{\theta}(t)) \leq \exp \left(-2 \lambda \sigma_{\mathrm{h}}^{2} t\right) \cdot \mathcal{L}(\boldsymbol{\theta}(0))
$$

Plugging $t=t_{0}$ in the above and simplifying, we have $\mathcal{L}\left(\boldsymbol{\theta}\left(t_{0}\right)\right) \leq \frac{1}{2 n}$.

- If $\left\{\boldsymbol{\theta}(t): t \in\left[0, t_{0}\right]\right\} \nsubseteq D$, then there exists a time $t^{\prime \prime} \leq t_{0}$ such that $\boldsymbol{\theta}\left(t^{\prime \prime}\right) \notin D$. Since $t^{\prime \prime} \leq t_{0}<t^{\prime}$, it must hold that $t^{\prime \prime} \notin L_{\frac{1}{2 n}}^{+}(\mathcal{L})$, and therefore $\mathcal{L}\left(\boldsymbol{\theta}\left(t^{\prime \prime}\right)\right)<\frac{1}{2 n}$ which implies $\mathcal{L}\left(\boldsymbol{\theta}\left(t_{0}\right)\right)<\frac{1}{2 n}$ since the flow is non-increasing.
- Suppose that $t^{\prime} \leq t_{0}$. Assume by contradiction that $\left\{\boldsymbol{\theta}(t): t \in\left[0, t^{\prime}\right]\right\} \subseteq D$. We will now show that the length of the trajectory of GF cannot have been long enough to reach the boundary of $D$, which will result in a contradiction. To this end, we use a similar technique as in Gupta et al. [2021, Thm. 9]. Define the potential function $\varepsilon(t)=\sqrt{\mathcal{L}(\boldsymbol{\theta}(t))}$. Taking the derivative of $\varepsilon(t)$ with respect to $t$ and using the chain rule we have

$$
\dot{\varepsilon}(t)=\frac{\frac{d \mathcal{L}(\boldsymbol{\theta}(t))}{d t}}{2 \sqrt{\mathcal{L}(\boldsymbol{\theta}(t))}}=-\frac{\|\nabla \mathcal{L}(\boldsymbol{\theta}(t))\|_{2}^{2}}{2 \sqrt{\mathcal{L}(\boldsymbol{\theta}(t))}} \leq-\sqrt{\frac{\lambda}{2}} \sigma_{\mathrm{h}} \cdot\|\nabla \mathcal{L}(\boldsymbol{\theta}(t))\|_{2},
$$

where the inequality follows from Eq. (24). We can now bound the length of the trajectory up until time $t^{\prime}$ by using the fundamental theorem of calculus and obtain

$$
\begin{aligned}
\int_{0}^{t^{\prime}}\|\nabla \mathcal{L}(\boldsymbol{\theta}(t))\|_{2} d t & \leq-\frac{1}{\sigma_{\mathrm{h}}} \sqrt{\frac{2}{\lambda}} \int_{0}^{t^{\prime}} \dot{\varepsilon}(t) d t \leq-\frac{1}{\sigma_{\mathrm{h}}} \sqrt{\frac{2}{\lambda}}[\sqrt{\mathcal{L}(\boldsymbol{\theta}(t))}]_{0}^{t^{\prime}} \\
& \leq \frac{1}{\sigma_{\mathrm{h}}} \sqrt{\frac{2 \mathcal{L}(\boldsymbol{\theta}(0))}{\lambda}} \leq \frac{1}{\sigma_{\mathrm{h}}} \sqrt{\frac{2 e}{\lambda}}<\Delta
\end{aligned}
$$

where in the second line, the first inequality uses the fact that $\mathcal{L}(\cdot)>0$, the second inequality follows from Eq. (23), and the last inequality follows from our bound on $\sigma_{\mathrm{h}}$ assumed in Eq. (3) and the definition of $\Delta$ in Eq. (17). In contrast, since $\boldsymbol{\theta}\left(t^{\prime}\right)$ is on the boundary of $U_{\Delta}(\boldsymbol{\theta}(0))$, this implies that there exists some neuron with weight and bias $w(t), b(t)$ at time
$t \geq 0$ such that $\left\|\left(w\left(t^{\prime}\right), b\left(t^{\prime}\right)\right)-(w(0), b(0))\right\|_{2}=\Delta$. From this and the path length upper bound we have

$$
\Delta \leq\left\|\boldsymbol{\theta}\left(t^{\prime}\right)-\boldsymbol{\theta}(0)\right\|_{2} \leq \int_{0}^{t^{\prime}}\|\nabla \mathcal{L}(\boldsymbol{\theta}(t))\|_{2} d t<\Delta
$$

which is a contradiction. We therefore must have that $\left\{\boldsymbol{\theta}(t): t \in\left[0, t^{\prime}\right]\right\} \nsubseteq D$. When this holds, there exists a time $t^{\prime \prime} \leq t^{\prime} \leq t_{0}$ such that $\boldsymbol{\theta}\left(t^{\prime \prime}\right) \notin D$. Since $t^{\prime \prime} \leq t^{\prime}$, it must hold that $t^{\prime \prime} \notin L_{\frac{1}{2 n}}^{+}(\mathcal{L})$, and therefore $\mathcal{L}\left(\boldsymbol{\theta}\left(t^{\prime \prime}\right)\right)<\frac{1}{2 n}$ which implies $\mathcal{L}\left(\boldsymbol{\theta}\left(t_{0}\right)\right)<\frac{1}{2 n}$ since the flow is non-increasing.

## D Over-parameterization is necessary

In this appendix, we further discuss and formally prove Theorem 3.2, which establishes that in general under Assumption 2.1, an over-parameterization of magnitude at least $1.3 r$ is necessary for achieving population loss below a constant. Our analysis is based on the following specific construction, where the labels are determined by a function $f_{r}$ parameterized by a natural number $r$ for all $x \in[-1,1]$, expressible by the sign of a teacher network of width $r$ and defined as

$$
\begin{equation*}
f_{r}(x):=\operatorname{sign}(\sin (0.5 \pi(r+1)(x+1))) . \tag{25}
\end{equation*}
$$

That is, $f_{r}$ changes value $r$ times between -1 and 1 on the interval $[-1,1]$, and is constant along intervals of length $\frac{2}{r+1}$. We now define the distribution $\mathcal{D}$ over the inputs of the dataset used in our lower bound and its corresponding labelling rule. We have

$$
\begin{equation*}
x \sim U[-1,1] \quad \text { and } \quad y=f_{r}(x) \tag{26}
\end{equation*}
$$

Recall the statement of Theorem 3.2, we have for example that if $\alpha=1.3$, then the width of the network being trained is no more than $1.3 r$ and GF attains loss at least $\frac{1}{160}$ in this case. While a lower bound of $r$ neurons for the construction specified in Eq. (26) is trivially implied by function approximation considerations, our lower bound merely improves upon this quantity by a constant multiplicative factor. Nevertheless, it is interesting to compare our lower bound to other similar settings in the literature, since it is typically difficult to derive lower bounds that require strictly more than $r$ neurons. For example, in a teacher-student setting where networks of the form $\mathbf{x} \mapsto$ $\sum_{i=1}^{r} \sigma\left(\mathbf{w}_{i}^{\top} \mathbf{x}\right)$ are considered, it is known that there are spurious (non-global) minima already when $r \geq 6$ [Safran and Shamir, 2018, Arjevani and Field, 2020, 2021], and that empirically we are more likely to get stuck in those minima the larger $r$ is [Safran and Shamir, 2018], but in spite of this ample empirical evidence, there is no proof that optimization will fail for any natural number $r \geq 6$. In contrast, in our univariate setting, it is possible to show a non-trivial lower bound since we utilize bias terms. This highlights the difference between settings that omit and include biases, which impacts the associated optimization problem in a non-trivial manner.

The proof of our lower bound, which appears below in Appendix D.1, relies on the observation that under Assumption 2.1, the breakpoints of the trained network upon initialization follow a standard Cauchy distribution. In such a case, neurons with a breakpoint outside the support of the data and with the wrong orientation will remain dormant throughout the optimization process, which requires initializing at least a fraction more of the minimal number of neurons required so that sufficiently many will be optimized and could improve the approximation of the target function. While one can circumvent this issue by scaling the breakpoints to the support of the data, this would require (i) an initialization scheme which is different than Assumption 2.1, which is used in our upper bounds; and (ii) this may even prove detrimental to optimization, as our positive result requires neurons that are active on all the data instances. We stress that our lower bound given here applies to training over a sample of any size, since it relies on approximation arguments. Additionally, we remark that by scaling the distribution $\mathcal{D}$ to be supported on a smaller interval we can increase the required magnitude of over-parameterization up to a factor of $\alpha=2$, however due to the common practice of scaling the data to have unit norm, we assume it is supported on $[-1,1]$. We also remark that a common initialization scheme is to set the bias terms to zero [He et al., 2015]. This results in breakpoints that are initialized at the origin and circumvents the issue of dormant neurons upon
initialization, however the main motivation for using such an initialization scheme is to achieve numerical stability and avoid exploding gradients when training very deep networks, which is not an issue for the shallow architecture we consider here. In any case, we stress that the goal of our lower bounds is to exemplify that over-parameterization is necessary in a setting complementary to our upper bound in Theorem 3.1, and we leave the derivation of stronger lower bounds under more general initialization schemes as a tantalizing future work direction.

## D. 1 Proof of Theorem 3.2

To prove the theorem, we would need the following auxiliary lemmas. The first lemma below establishes that if we approximate the function $f_{r}$ which is defined in Eq. (25) by a function which does not change its sign over an interval of length larger than $\frac{2}{r+1}$, then this results in a strictly positive loss which is roughly proportional to the length of the approximation interval.
Lemma D.1. Let $f_{r}$ be as defined in Eq. (25), let $\beta_{1}, \beta_{2} \in[-1,1], \mathcal{N}:\left[\beta_{1}, \beta_{2}\right] \rightarrow \mathbb{R}$ such that its sign is fixed on $\left[\beta_{1}, \beta_{2}\right]$, and let $\ell$ be either the exponential or logistic loss. Then

$$
\int_{\beta_{1}}^{\beta_{2}} \ell\left(\mathcal{N}(x) \cdot f_{r}(x)\right) d x \geq \frac{1}{2} \ell(0)\left(\beta_{2}-\beta_{1}-\frac{2}{r+1}\right)
$$

Proof. If $\beta_{2}-\beta_{1} \leq \frac{2}{r+1}$, then the right-hand side is non-positive and the lemma follows since $\ell(\cdot)>0$. If $\beta_{2}-\beta_{1}>\frac{2}{r+1}$, then the (one-dimensional Lebesgue) measure of the set $A:=\{x \in$ $\left.\left[\beta_{1}, \beta_{2}\right]: \mathcal{N}(x) \cdot f_{r}(x) \leq 0\right\}$ is at least

$$
\frac{1}{2}\left(\beta_{2}-\beta_{1}-\frac{2}{r+1}\right)
$$

since the measure of the complementary set $\left\{x \in\left[\beta_{1}, \beta_{2}\right]: \mathcal{N}(x) \cdot f_{r}(x)>0\right\}$ is at most $\frac{1}{2}\left(\beta_{2}-\beta_{1}+\frac{2}{r+1}\right)$, where the upper bound is attained when $\mathcal{N}(x) \cdot f_{r}(x)>0$ for all

$$
x \in\left[\beta_{1}, \beta_{1}+\frac{2}{r+1}\right] \cup\left[\beta_{2}-\frac{2}{r+1}, \beta_{2}\right]
$$

We can therefore lower bound the integral in the lemma by

$$
\int_{\beta_{1}}^{\beta_{2}} \ell\left(\mathcal{N}(x) \cdot f_{r}(x)\right) d x \geq \int_{x \in A} \ell\left(\mathcal{N}(x) \cdot f_{r}(x)\right) d x \geq \int_{x \in A} \ell(0) d x \geq \frac{1}{2} \ell(0)\left(\beta_{2}-\beta_{1}-\frac{2}{r+1}\right)
$$

The following lemma shows that approximating $f_{r}$ using a ReLU network with just $r^{\prime}$ neurons results in loss proportional to $1-r^{\prime} / r$.
Lemma D.2. Suppose that $f_{r}$ as defined in Eq. (25). Then for any ReLU network $\mathcal{N}_{\boldsymbol{\theta}}$ of width at most $r^{\prime}$, we have

$$
\mathcal{L}_{\mathcal{D}}(\boldsymbol{\theta}) \geq \frac{1}{4}\left(1-\frac{r^{\prime}}{r}\right)
$$

Proof. Denote by $\beta_{1}, \ldots, \beta_{r^{\prime}}$ the set of points where $\mathcal{N}$ changes sign in $(-1,1)$. Note that this set is always of size at most $r^{\prime}$, and we may assume without loss of generality that it is of size exactly $r^{\prime}$ (since otherwise we can prove a stronger claim, where the lemma holds for some $r^{\prime \prime}<r^{\prime}$ ). Further define the boundaries $\beta_{0}:=-1$ and $\beta_{r^{\prime}+1}:=1$. We compute

$$
\begin{aligned}
\mathcal{L}_{\mathcal{D}}(\boldsymbol{\theta}) & =\int_{-1}^{1} \frac{1}{2} \ell\left(f_{r}(x) \cdot \mathcal{N}_{\boldsymbol{\theta}}(x)\right) d x=\frac{1}{2} \sum_{i=0}^{r^{\prime}} \int_{\beta_{i}}^{\beta_{i+1}} \ell\left(f_{r}(x) \cdot \mathcal{N}_{\boldsymbol{\theta}}(x)\right) d x \\
& \geq \frac{1}{4} \ell(0) \sum_{i=0}^{r^{\prime}}\left(\beta_{i+1}-\beta_{i}-\frac{2}{r+1}\right)=\frac{1}{4}\left(2-2 \frac{r^{\prime}+1}{r+1}\right) \\
& =\frac{1}{2} \cdot \frac{r-r^{\prime}}{r+1} \geq \frac{1}{4}\left(1-\frac{r^{\prime}}{r}\right)
\end{aligned}
$$

where the first inequality uses Lemma D.1, the equality that follows is due to the sum telescoping and since $\ell(0)=1$ for both the exponential and logistic losses, and the final inequality is due to $r \geq 1$.

With the above auxiliary lemmas, we can now turn to the proof of the theorem.

Proof of Theorem 3.2. By Assumption 2.1, the breakpoints of $\mathcal{N}$ at initialization follow a standard Cauchy distribution. Since $\mathcal{D}$ is supported on $[-1,1]$, we have with probability exactly 0.5 that the breakpoint of a given neuron falls outside of $[-1,1]$. Moreover, with an independent probability of 0.5 , the orientation of the neuron is such that it is off on all the data instances. I.e., such a neuron remains dormant throughout the optimization process of GF with probability 0.25 . It can be verified that for any integer $k \geq 1$, at least $\lceil 0.25 k\rceil$ neurons will be dormant upon initialization with probability at least $0.25 .^{7}$ Thus, with probability at least 0.25 , we have that out of $\alpha r$ neurons, there are at most $r^{\prime}=\lfloor 0.75 \alpha r\rfloor$ neurons that are effectively being trained with breakpoints in $[-1,1]$. By Lemma D.2, this results in a lower bound on the population loss of

$$
\mathcal{L}_{\mathcal{D}}\left(\mathcal{N}_{\boldsymbol{\theta}(t)}\right) \geq \frac{1}{4}\left(1-\frac{\lfloor 0.75 \alpha r\rfloor}{r}\right) \geq \frac{1}{4}(1-0.75 \alpha),
$$

for any time $t \geq 0$.

## E Proof of Theorem 4.2

By Theorem 4.1, if there exists time $t_{0}$ such that $\mathcal{L}\left(\boldsymbol{\theta}\left(t_{0}\right)\right)<\frac{1}{n}$ then GF converges to zero loss, and converges in direction to a KKT point of Problem (4). We denote $\mathcal{N}_{\boldsymbol{\theta}}(x)=\sum_{j \in[k]} v_{j} \sigma\left(w_{j} x+b_{j}\right)$. Thus, $\mathcal{N}_{\boldsymbol{\theta}}$ is a network of width $k$, where the weights in the first layer are $w_{1}, \ldots, w_{k}$, the bias terms are $b_{1}, \ldots, b_{k}$, and the weights in the second layer are $v_{1}, \ldots, v_{k}$. We denote $J:=\left\{j \in[k]: v_{j} \neq 0\right\}$, $J^{+}:=\left\{j \in J: v_{j}>0\right\}$, and $J^{-}:=\left\{j \in J: v_{j}<0\right\}$. Since neurons with output weight $v_{j}=0$ do not affect the function that the network computes, then in this proof we ignore them. We also denote $I:=[n]$, and $I^{\prime}=\left\{i \in I: y_{i} \mathcal{N}_{\boldsymbol{\theta}}\left(x_{i}\right)=1\right\}$. Thus, $I^{\prime}$ are the indices of the examples where $\mathcal{N}_{\boldsymbol{\theta}}$ attains margin of exactly 1.
Assume that $\mathcal{N}_{\boldsymbol{\theta}}$ satisfies the KKT conditions of Problem (4). Thus, there are $\lambda_{1}, \ldots, \lambda_{n}$ such that for every $j \in J$ we have

$$
\begin{equation*}
w_{j}=\sum_{i \in I} \lambda_{i} \frac{\partial}{\partial w_{j}}\left(y_{i} \mathcal{N}_{\boldsymbol{\theta}}\left(x_{i}\right)\right)=\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime} x_{i} \tag{27}
\end{equation*}
$$

where $\sigma_{i, j}^{\prime}$ is a subgradient of $\sigma$ at $w_{j} \cdot x_{i}+b_{j}$, i.e., if $w_{j} \cdot x_{i}+b_{j} \neq 0$ then $\sigma_{i, j}^{\prime}=\mathbb{1}\left[w_{j} \cdot x_{i}+b_{j}>0\right]$, and otherwise $\sigma_{i, j}^{\prime}$ is some value in $[0,1]$ (we emphasize that in this case $\sigma_{i, j}^{\prime}$ may be any value in $[0,1]$ and in this proof we do not have any further assumptions on it). Also we have $\lambda_{i} \geq 0$ for all $i \in I$, and $\lambda_{i}=0$ if $i \notin I^{\prime}$. Likewise, we have

$$
\begin{equation*}
b_{j}=\sum_{i \in I} \lambda_{i} \frac{\partial}{\partial b_{j}}\left(y_{i} \mathcal{N}_{\boldsymbol{\theta}}\left(x_{i}\right)\right)=\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime} \tag{28}
\end{equation*}
$$

We say that $\mathcal{N}_{\boldsymbol{\theta}}$ has an activation point at $x$ if there is $j \in[J]$ with $w_{j} \neq 0$ such that $w_{j} \cdot x+b_{j}=0$. In this case we say that the activation point $x$ corresponds to the neuron $j$. Note that if $w_{j}=0$ then the neuron computes a constant function and thus it does not affect the number of linear regions in $\mathcal{N}_{\theta}$.
Lemma E.1. We denote $I^{\prime}=\left\{i_{1}, \ldots, i_{q}\right\}$ where $1 \leq i_{1}<\ldots<i_{q} \leq[n]$. For every $\ell \in[q-1]$ the network $\mathcal{N}_{\boldsymbol{\theta}}$ has at most two activation points in the open interval $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$. Moreover, $\mathcal{N}_{\boldsymbol{\theta}}$ has at most one activation point in $\left(-\infty, i_{1}\right)$ and at most one activation points in $\left(i_{q}, \infty\right)$.

[^2]Proof. Let $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ be an activation point, and let $j \in J$ such that $w_{j} \neq 0$ and $w_{j} \cdot x+b_{j}=0$. Suppose first that $w_{j}>0$. Since $w_{j} \cdot x+b_{j}=0$ then for every $x^{\prime}>x$ we have $w_{j} \cdot x^{\prime}+b_{j}>0$, and for every $x^{\prime}<x$ we have $w_{j} \cdot x^{\prime}+b_{j}<0$. By Eq. (27) we have

$$
w_{j}=\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime} x_{i}=\sum_{i \in I^{\prime}} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime} x_{i}
$$

Since $x \neq x_{i}$ for all $i \in I^{\prime}$ then $w_{j} \cdot x_{i}+b_{j} \neq 0$, and we have $\sigma_{i, j}^{\prime}=\mathbb{1}\left(w_{j} \cdot x_{i}+b_{j}>0\right)=\mathbb{1}\left[x_{i}>x\right]$. Therefore, the above displayed equation equals

$$
\begin{equation*}
\sum_{i \in I^{\prime}} \lambda_{i} y_{i} v_{j} \mathbb{1}\left[x_{i}>x\right] x_{i}=\sum_{i \in I^{\prime}, i \geq i_{\ell+1}} \lambda_{i} y_{i} v_{j} x_{i} \tag{29}
\end{equation*}
$$

Likewise, by Eq. (28) we have

$$
\begin{equation*}
b_{j}=\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime}=\sum_{i \in I^{\prime}} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime}=\sum_{i \in I^{\prime}} \lambda_{i} y_{i} v_{j} \mathbb{1}\left[x_{i}>x\right]=\sum_{i \in I^{\prime}, i \geq i_{\ell+1}} \lambda_{i} y_{i} v_{j} . \tag{30}
\end{equation*}
$$

By Eq. (29) and Eq. (30), the activation point $x$ satisfies

$$
x=\frac{-b_{j}}{w_{j}}=\frac{-\sum_{i \in I^{\prime}, i \geq i_{\ell+1}} \lambda_{i} y_{i} v_{j}}{\sum_{i \in I^{\prime}, i \geq i_{\ell+1}} \lambda_{i} y_{i} v_{j} x_{i}}=\frac{-\sum_{i \in I^{\prime}, i \geq i_{\ell+1}} \lambda_{i} y_{i}}{\sum_{i \in I^{\prime}, i \geq i_{\ell+1}} \lambda_{i} y_{i} x_{i}} .
$$

Therefore if $x$ and $x^{\prime}$ are two activation points in $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ that correspond to $w_{j}>0$ and $w_{j^{\prime}}>0$ respectively, then $x=x^{\prime}$. Thus, there is at most one activation point $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ that corresponds to some $w_{j}>0$.
Moreover, by similar arguments there is at most one activation point $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ that corresponds to $w_{j}<0$. Overall, in the interval $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ there are at most two activation points.
If $x \in\left(-\infty, i_{1}\right)$ then from similar argument we get that there is at most one activation point that corresponds to a neuron $j$ with $w_{j}>0$. Also, an activation point in $\left(-\infty, i_{1}\right)$ that corresponds to a neuron with $w_{j}<0$ does not exist, since such neuron is not active for any input $x_{i}$ with $i \in I^{\prime}$, and hence by Eq. (27) we must have $w_{j}=0$. The proof of the claim for the interval $\left(i_{q}, \infty\right)$ is similar.

Let $1 \leq a<b \leq n$ be indices such that for every $a \leq i<i^{\prime} \leq b$ we have $y_{i}=y_{i^{\prime}}$. Thus the labels do not switch signs for the inputs $x_{a}, x_{a+1}, \ldots, x_{b}$. Intuitively, the proof follows by showing that in the interval $\left[x_{a}, x_{b}\right]$ the network $\mathcal{N}_{\boldsymbol{\theta}}$ has a constant number of linear regions, and then concluding that the overall number of linear regions in $\mathcal{N}_{\boldsymbol{\theta}}$ must be $\mathcal{O}(r)$. We first consider the case where $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=1$. In the following lemmas we analyze the activation points in this case and obtain a bound on the number of linear regions. Then, we will extend this result also to the cases where $y_{i}=-1$ and where $x_{a}<x_{b} \leq 0$. For a given activation point $x$ we say that the derivative of the network increases (respectively, decreases) in $x$ if for every sufficiently small $\varepsilon>0$ the derivative of the network at $x-\varepsilon$ is smaller (respectively, larger) than the derivative at $x+\varepsilon$.
Lemma E.2. Suppose that $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=1$. In the interval $\left[x_{a}, x_{b}\right]$ the network $\mathcal{N}_{\boldsymbol{\theta}}$ has at most two activation points where the derivative decreases.

Proof. If the derivative decreases at an activation point $x \geq 0$, then there is at least one neuron $j \in J$ where $w_{j} \cdot x+b_{j}=0$ and the derivative (of the function computed by this neuron) decreases in $x$. There are two types of such neurons: either (a) $w_{j}>0, b_{j} \leq 0$ and $v_{j}<0$; or (b) $w_{j}<0, b_{j} \geq 0$ and $v_{j}<0$.

We now show that in the interval $\left[x_{a}, x_{b}\right]$ there is at most one activation point that corresponds to a neuron of type (a) and at most one activation point that corresponds to a neuron of type (b). We note that an activation point might correspond to multiple neurons, namely, to a set $Q \subseteq J$ of neurons of size larger than 1 . However, we show that if $x, x^{\prime}$ are activation points in $\left[x_{a}, x_{b}\right]$ that correspond to sets $Q_{x}$ and $Q_{x^{\prime}}$ of neurons (respectively) and both sets $Q_{x}, Q_{x^{\prime}}$ contain neurons of type (a) then $x=x^{\prime}$. Likewise, if both sets $Q_{x}, Q_{x^{\prime}}$ contain neurons of type (b) then we also have $x=x^{\prime}$.
Suppose towards contradiction that $x \in\left[x_{a}, x_{b}\right]$ is an activation point that corresponds to a neuron $j$ of type (a), and $x^{\prime} \in\left[x_{a}, x_{b}\right]$ is an activation point with $x^{\prime}>x$ that corresponds to a neuron $j^{\prime}$
of type (a). Since both neurons $j, j^{\prime}$ are of type (a), then we have $w_{j} \cdot z+b_{j}>0$ iff $z>x$ and $w_{j^{\prime}} \cdot z+b_{j^{\prime}}>0$ iff $z>x^{\prime}$. By Eq. (27) we have

$$
\frac{1}{v_{j^{\prime}}} \cdot w_{j^{\prime}}=\frac{1}{v_{j^{\prime}}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j^{\prime}} \sigma_{i, j^{\prime}}^{\prime} x_{i}\right)=\sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j^{\prime}}^{\prime} x_{i} \leq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \geq x^{\prime}\right] x_{i}
$$

where the last inequality is since $\sigma_{i, j^{\prime}}^{\prime}=\mathbb{1}\left[x_{i} \geq x^{\prime}\right]$ if $x_{i} \neq x^{\prime}$, and when $x_{i}=x^{\prime}$ we have $\sigma_{i, j^{\prime}}^{\prime} \leq \mathbb{1}\left[x_{i} \geq x^{\prime}\right]$ (and recall that for $a \leq i \leq b$ we have $y_{i}=1, x_{i} \geq 0$ and $\lambda_{i} \geq 0$ ). The above RHS equals

$$
\begin{aligned}
\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i}>x\right] x_{i}-\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x<x_{i}<x^{\prime}\right] x_{i} & \leq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{\mathbb { 1 }}\left[x_{i}>x\right] x_{i} \\
& \leq \sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j}^{\prime} x_{i} \\
& =\frac{1}{v_{j}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime} x_{i}\right) \\
& =\frac{1}{v_{j}} \cdot w_{j}
\end{aligned}
$$

Since $v_{j}<0$ we conclude that

$$
\begin{equation*}
w_{j} \leq \frac{v_{j}}{v_{j^{\prime}}} \cdot w_{j^{\prime}} \tag{31}
\end{equation*}
$$

Likewise, by Eq. (28) we have

$$
\begin{aligned}
\frac{1}{v_{j^{\prime}}} \cdot b_{j^{\prime}} & =\frac{1}{v_{j^{\prime}}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j^{\prime}} \sigma_{i, j^{\prime}}^{\prime}\right)=\sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j^{\prime}}^{\prime} \leq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \geq x^{\prime}\right] \\
& =\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i}>x\right]-\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x<x_{i}<x^{\prime}\right] \\
& \leq \sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j}^{\prime}=\frac{1}{v_{j}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime}\right)=\frac{1}{v_{j}} \cdot b_{j}
\end{aligned}
$$

Hence, we conclude that

$$
\begin{equation*}
b_{j} \leq \frac{v_{j}}{v_{j^{\prime}}} \cdot b_{j^{\prime}} \tag{32}
\end{equation*}
$$

Since $0 \leq x<x^{\prime}$, then by using Eq. (31) and (32) we have

$$
0<w_{j} \cdot x^{\prime}+b_{j} \leq \frac{v_{j}}{v_{j^{\prime}}} \cdot w_{j^{\prime}} \cdot x^{\prime}+\frac{v_{j}}{v_{j^{\prime}}} \cdot b_{j^{\prime}}=\frac{v_{j}}{v_{j^{\prime}}}\left(w_{j^{\prime}} \cdot x^{\prime}+b_{j^{\prime}}\right)=0
$$

Thus, we reached a contradiction.
Next, suppose that $x \in\left[x_{a}, x_{b}\right]$ is an activation point that corresponds to a neuron $j$ of type (b), and $x^{\prime} \in\left[x_{a}, x_{b}\right]$ is an activation point with $x^{\prime}>x$ that corresponds to a neuron $j^{\prime}$ of type (b). We will reach a contradiction using similar arguments to the case of type (a) neurons, with some required modifications.
Since both neurons $j, j^{\prime}$ are of type (b), then we have $w_{j} \cdot z+b_{j}>0$ iff $z<x$ and $w_{j^{\prime}} \cdot z+b_{j^{\prime}}>0$ iff $z<x^{\prime}$. By Eq. (27) we have

$$
\frac{1}{v_{j^{\prime}}} \cdot w_{j^{\prime}}=\frac{1}{v_{j^{\prime}}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j^{\prime}} \sigma_{i, j^{\prime}}^{\prime} x_{i}\right)=\sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j^{\prime}}^{\prime} x_{i} \geq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{\mathbb { 1 }}\left[x_{i}<x^{\prime}\right] x_{i}
$$

where the last inequality is since $\sigma_{i, j^{\prime}}^{\prime}=\mathbb{1}\left[x_{i}<x^{\prime}\right]$ if $x_{i} \neq x^{\prime}$, and when $x_{i}=x^{\prime}$ we have $\sigma_{i, j^{\prime}}^{\prime} \geq \mathbb{1}\left[x_{i}<x^{\prime}\right]$. The above RHS equals

$$
\begin{aligned}
\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \leq x\right] x_{i}+\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x<x_{i}<x^{\prime}\right] x_{i} & \geq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \leq x\right] x_{i} \\
& \geq \sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j}^{\prime} x_{i} \\
& =\frac{1}{v_{j}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime} x_{i}\right) \\
& =\frac{1}{v_{j}} \cdot w_{j}
\end{aligned}
$$

Since $v_{j}<0$ we conclude that

$$
\begin{equation*}
w_{j} \geq \frac{v_{j}}{v_{j^{\prime}}} \cdot w_{j^{\prime}} \tag{33}
\end{equation*}
$$

Likewise, by Eq. (28) we have

$$
\begin{aligned}
\frac{1}{v_{j^{\prime}}} \cdot b_{j^{\prime}} & =\frac{1}{v_{j^{\prime}}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j^{\prime}} \sigma_{i, j^{\prime}}^{\prime}\right)=\sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j^{\prime}}^{\prime} \geq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i}<x^{\prime}\right] \\
& =\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \leq x\right]+\sum_{i \in I} \lambda_{i} y_{i} \mathbb{\mathbb { 1 }}\left[x<x_{i}<x^{\prime}\right] \\
& \geq \sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j}^{\prime}=\frac{1}{v_{j}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime}\right)=\frac{1}{v_{j}} \cdot b_{j}
\end{aligned}
$$

Hence, we conclude that

$$
\begin{equation*}
b_{j} \geq \frac{v_{j}}{v_{j^{\prime}}} \cdot b_{j^{\prime}} \tag{34}
\end{equation*}
$$

Since $0 \leq x<x^{\prime}$, and by using Eq. (33) and (34), we have

$$
0>w_{j} \cdot x^{\prime}+b_{j} \geq \frac{v_{j}}{v_{j^{\prime}}} \cdot w_{j^{\prime}} \cdot x^{\prime}+\frac{v_{j}}{v_{j^{\prime}}} \cdot b_{j^{\prime}}=\frac{v_{j}}{v_{j^{\prime}}}\left(w_{j^{\prime}} \cdot x^{\prime}+b_{j^{\prime}}\right)=0
$$

Thus, we reached a contradiction.
We denote $\mathcal{I}_{a, b}:=\{a, a+1, \ldots, b\} \subseteq I$ and $\mathcal{I}_{a, b}^{\prime}:=\left\{i \in \mathcal{I}_{a, b}: y_{i} \mathcal{N}_{\boldsymbol{\theta}}\left(x_{i}\right)=1\right\}=\mathcal{I}_{a, b} \cap I^{\prime}$. Thus, $\mathcal{I}_{a, b}^{\prime}$ are the indices of the examples in the interval $\left[x_{a}, x_{b}\right]$ where $\mathcal{N}_{\boldsymbol{\theta}}$ attains margin of exactly 1 . We denote $\mathcal{I}_{a, b}^{\prime}=\left\{i_{1}, \ldots, i_{m}\right\}$, where $a \leq i_{1}<\ldots<i_{m} \leq b$.
Lemma E.3. Suppose that $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=1$. There are at most 2 indices $\ell \in[m-1]$ such that $\mathcal{N}_{\boldsymbol{\theta}}(x)>1$ for some $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$.

Proof. Assume that $\mathcal{N}_{\boldsymbol{\theta}}(x)>1$ for $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$. Since $\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell}}\right)=\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+1}}\right)=1$, then we have $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$. Now, since $\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell}}\right)=\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+1}}\right)=1$ and $\mathcal{N}_{\boldsymbol{\theta}}(x)>1$ for some $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$, then there must be an activation point in $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ where the derivative decreases. Since by Lemma E. 2 there are at most two such activation points in $\left[x_{a}, x_{b}\right]$ then the lemma follows.

Lemma E.4. Suppose that $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=1$. There are at most 5 indices $\ell \in[m-1]$ such that $\mathcal{N}_{\boldsymbol{\theta}}(x)<1$ for some $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$.

Proof. Assume that $\mathcal{N}_{\boldsymbol{\theta}}(x)<1$ for $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$. Since $\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell}}\right)=\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+1}}\right)=1$, then $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$. If $\ell \neq m-1$ then we have $\mathcal{N}_{\boldsymbol{\theta}}(x)<1$ and $\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+1}}\right)=\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+2}}\right)=1$. Hence, the interval $\left(x, x_{i_{\ell+1}}\right)$ contains a point with positive derivative, and the interval ( $x_{i_{\ell+1}}, x_{i_{\ell+2}}$ ) contains a point with non-positive derivative. Thus, there must be an activation point in $\left(x, x_{i_{\ell+2}}\right)$ where the derivative decreases. Therefore, there is an activation point with decreasing derivative either in the interval $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ or in the interval $\left[x_{i_{\ell+1}}, x_{i_{\ell+2}}\right)$ (and possibly in both). Thus, an interval
$\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ for $\ell \neq m-1$ might contain some $x$ with $\mathcal{N}_{\boldsymbol{\theta}}(x)<1$ only if there is an activation point with decreasing derivative in $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ or $\left[x_{i_{\ell+1}}, x_{i_{\ell+2}}\right)$. Since by Lemma E. 2 there are at most two such activation points in $\left[x_{a}, x_{b}\right]$, then there are at most 4 intervals $\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ with $\ell \neq m-1$ that contain some $x$ with $\mathcal{N}_{\boldsymbol{\theta}}(x)<1$. The interval $\left[x_{i_{m-1}}, x_{i_{m}}\right]$ might also contain such $x$. Overall, there are at most 5 indices $\ell \in[m-1]$ such that $\mathcal{N}_{\boldsymbol{\theta}}(x)<1$ for some $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$.
Lemma E.5. Suppose that $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=1$. There are at most 30 boundaries between linear regions in $\left[x_{i_{1}}, x_{i_{m}}\right]$.

Proof. By Lemmas E. 3 and E. 4 there are at most 7 indices $\ell \in[m-1]$ such that $\mathcal{N}_{\boldsymbol{\theta}}(x) \neq 1$ for some $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$. We denote the set of these indices by $R$. Let $x \in\left(x_{i_{1}}, x_{i_{m}}\right)$ be a boundary between two linear regions. Note that if $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ for some $\ell \in[m-1]$ then $\ell \in R$. Also, if $x=x_{i_{\ell}}$ for some $2 \leq \ell \leq m-1$ then either $\ell \in R$ or $\ell-1 \in R$. In any case, we have $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ for some $\ell \in R$. Note that each boundary between linear regions is also an activation point. Therefore, the number of boundaries between linear regions in $\left(x_{i_{1}}, x_{i_{m}}\right)$ is at most the number of activation points in the intervals $\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ with $\ell \in R$. By Lemma E. 1 each interval $\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right.$ ] contains at most 4 activation points: two points in $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ and two in $\left\{x_{i_{\ell}}, x_{i_{\ell+1}}\right\}$. Overall, there are at most $|R| \cdot 4 \leq 28$ boundaries between linear regions in $\left(x_{i_{1}}, x_{i_{m}}\right)$. Thus, there are at most 30 boundaries between linear regions in $\left[x_{i_{1}}, x_{i_{m}}\right]$.

In the above lemmas we considered the case where $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=1$, and proved that $\mathcal{I}_{a, b}^{\prime}$ is such that there are at most 30 boundaries between linear regions in [ $x_{i_{1}}, x_{i_{m}}$ ]. In Subsection E. 1 we show analogous results for the case where $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=-1$. Thus, if $x_{b}>x_{a} \geq 0$ and the labels do not switch sign in the interval $\left[x_{a}, x_{b}\right]$ (i.e., either all labels are 1 or all labels are -1 ) then there are at most 30 boundaries between linear regions in $\left[x_{i_{1}}, x_{i_{m}}\right]$. The case where $x_{a}<x_{b} \leq 0$ (and the labels do not switch sign in the interval $\left[x_{a}, x_{b}\right]$ ) can be handled in a similar manner. Thus, even where the inputs are negative, $\mathcal{I}_{a, b}^{\prime}$ is such that there are at most 30 boundaries between linear regions in $\left[x_{i_{1}}, x_{i_{m}}\right]$. The proof for this case is similar and for conciseness we do not repeat it.
We are now ready to finish the proof of the theorem. Consider the set $I^{\prime}$ of indices where $\mathcal{N}_{\boldsymbol{\theta}}$ attains margin 1 and denote $I^{\prime}=\left\{i_{1}, \ldots, i_{q}\right\}$. Note that if $I^{\prime}$ is an empty set, then by Eq. (27) and (28) all neurons have $w_{j}=b_{j}=0$ and hence the network $\mathcal{N}_{\boldsymbol{\theta}}$ is the zero function. Let $\ell \leq \ell^{\prime}$ be such that the labels of the examples in the dataset do not change sign in the interval $\left[x_{i_{\ell}}, x_{i_{\ell}}\right]$, and either $0 \leq x_{i_{\ell}} \leq x_{i_{\ell^{\prime}}}$ or $x_{i_{\ell}} \leq x_{i_{\ell^{\prime}}} \leq 0$. Thus, the interval [ $x_{i_{\ell}}, x_{i_{\ell^{\prime}}}$ ] contains at most 30 boundaries between linear regions. Also, by Lemma E. 1 the interval $\left(x_{i_{\ell-1}}, x_{i_{\ell}}\right)$ (or $\left(-\infty, x_{i_{\ell}}\right)$ if $\ell=1$ ) contains at most two boundaries between linear regions. Likewise, the interval ( $x_{i_{\ell^{\prime}}}, x_{i_{\ell^{\prime}+1}}$ ) (or $\left(x_{i_{\ell^{\prime}}}, \infty\right)$ if $\ell^{\prime}=q$ ) contains at most two boundaries between linear regions. Recall that the labels in the dataset switch sign at most $r$ times. Overall, we get that the number of boundaries between linear regions in the whole domain $\mathbb{R}$ is at most $30(r+2)+2(r+3)=32 r+66$. Indeed, if one of the $r+1$ intervals where $\mathcal{N}_{\boldsymbol{\theta}}$ do not switch sign contains 0 then we split it into two intervals, and thus we obtain $r+2$ intervals. Each of these intervals includes at most 30 boundaries, and outside of these intervals there are at most $2(r+3)$ boundaries. Thus, that are at most $32 r+67$ linear regions.

## E. 1 Lemmas for the case $y_{i}=-1$

Lemma E.6. Suppose that $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=-1$. In the interval $\left[x_{a}, x_{b}\right]$ the network $\mathcal{N}_{\boldsymbol{\theta}}$ has at most two activation points where the derivative increases.

Proof. If the derivative increases at an activation point $x \geq 0$, then there is a least one neuron $j \in J$ where $w_{j} \cdot x+b_{j}=0$ and the derivative (of the function computed by this neuron) increases in $x$. There are two types of such neurons: either (a) $w_{j}>0, b_{j} \leq 0$ and $v_{j}>0$; or (b) $w_{j}<0, b_{j} \geq 0$ and $v_{j}>0$.
We now show that in the interval $\left[x_{a}, x_{b}\right]$ there is at most one activation point that corresponds to a neuron of type (a) and at most one activation point that corresponds to a neuron of type (b). We note that an activation point might correspond to multiple neurons, namely, to a set $Q \subseteq J$ of neurons of size larger than 1 . However, we show that if $x, x^{\prime}$ are activation points in $\left[x_{a}, x_{b}\right]$ that correspond to sets $Q_{x}$ and $Q_{x^{\prime}}$ of neurons (respectively) and both sets $Q_{x}, Q_{x^{\prime}}$ contain neurons of type (a) then $x=x^{\prime}$. Likewise, if both sets $Q_{x}, Q_{x^{\prime}}$ contain neurons of type (b) then we also have $x=x^{\prime}$.

Suppose towards contradiction that $x \in\left[x_{a}, x_{b}\right]$ is an activation point that corresponds to a neuron $j$ of type (a), and $x^{\prime} \in\left[x_{a}, x_{b}\right]$ is an activation point with $x^{\prime}>x$ that corresponds to a neuron $j^{\prime}$ of type (a). Since both neurons $j, j^{\prime}$ are of type (a), then we have $w_{j} \cdot z+b_{j}>0$ iff $z>x$ and $w_{j^{\prime}} \cdot z+b_{j^{\prime}}>0$ iff $z>x^{\prime}$. By Eq. (27) we have

$$
\frac{1}{v_{j^{\prime}}} \cdot w_{j^{\prime}}=\frac{1}{v_{j^{\prime}}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j^{\prime}} \sigma_{i, j^{\prime}}^{\prime} x_{i}\right)=\sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j^{\prime}}^{\prime} x_{i} \geq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{\mathbb { 1 }}\left[x_{i} \geq x^{\prime}\right] x_{i}
$$

where the last inequality is since $\sigma_{i, j^{\prime}}^{\prime}=\mathbb{1}\left[x_{i} \geq x^{\prime}\right]$ if $x_{i} \neq x^{\prime}$, and when $x_{i}=x^{\prime}$ we have $\sigma_{i, j^{\prime}}^{\prime} \leq \mathbb{1}\left[x_{i} \geq x^{\prime}\right]$ (and recall that for $a \leq i \leq b$ we have $y_{i}=-1, x_{i} \geq 0$ and $\lambda_{i} \geq 0$ ). The above RHS equals

$$
\begin{aligned}
\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i}>x\right] x_{i}-\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x<x_{i}<x^{\prime}\right] x_{i} & \geq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i}>x\right] x_{i} \\
& \geq \sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j}^{\prime} x_{i} \\
& =\frac{1}{v_{j}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime} x_{i}\right) \\
& =\frac{1}{v_{j}} \cdot w_{j}
\end{aligned}
$$

We conclude that

$$
\begin{equation*}
w_{j} \leq \frac{v_{j}}{v_{j^{\prime}}} \cdot w_{j^{\prime}} \tag{35}
\end{equation*}
$$

Likewise, by Eq. (28) we have

$$
\begin{aligned}
\frac{1}{v_{j^{\prime}}} \cdot b_{j^{\prime}} & =\frac{1}{v_{j^{\prime}}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j^{\prime}} \sigma_{i, j^{\prime}}^{\prime}\right)=\sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j^{\prime}}^{\prime} \geq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \geq x^{\prime}\right] \\
& =\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i}>x\right]-\sum_{i \in I} \lambda_{i} y_{i} \mathbb{\mathbb { 1 }}\left[x<x_{i}<x^{\prime}\right] \\
& \geq \sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j}^{\prime}=\frac{1}{v_{j}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime}\right)=\frac{1}{v_{j}} \cdot b_{j}
\end{aligned}
$$

Hence, we conclude that

$$
\begin{equation*}
b_{j} \leq \frac{v_{j}}{v_{j^{\prime}}} \cdot b_{j^{\prime}} \tag{36}
\end{equation*}
$$

Since $0 \leq x<x^{\prime}$, then by using Eq. (35) and (36) we have

$$
0<w_{j} \cdot x^{\prime}+b_{j} \leq \frac{v_{j}}{v_{j^{\prime}}} \cdot w_{j^{\prime}} \cdot x^{\prime}+\frac{v_{j}}{v_{j^{\prime}}} \cdot b_{j^{\prime}}=\frac{v_{j}}{v_{j^{\prime}}}\left(w_{j^{\prime}} \cdot x^{\prime}+b_{j^{\prime}}\right)=0 .
$$

Thus, we reached a contradiction.
Next, suppose that $x \in\left[x_{a}, x_{b}\right]$ is an activation point that corresponds to a neuron $j$ of type (b), and $x^{\prime} \in\left[x_{a}, x_{b}\right]$ is an activation point with $x^{\prime}>x$ that corresponds to a neuron $j^{\prime}$ of type (b). We will reach a contradiction using similar arguments to the case of type (a) neurons, with some required modifications.
Since both neurons $j, j^{\prime}$ are of type (b), then we have $w_{j} \cdot z+b_{j}>0$ iff $z<x$ and $w_{j^{\prime}} \cdot z+b_{j^{\prime}}>0$ iff $z<x^{\prime}$. By Eq. (27) we have

$$
\frac{1}{v_{j^{\prime}}} \cdot w_{j^{\prime}}=\frac{1}{v_{j^{\prime}}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j^{\prime}} \sigma_{i, j^{\prime}}^{\prime} x_{i}\right)=\sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j^{\prime}}^{\prime} x_{i} \leq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{\mathbb { 1 }}\left[x_{i}<x^{\prime}\right] x_{i}
$$

where the last inequality is since $\sigma_{i, j^{\prime}}^{\prime}=\mathbb{1}\left[x_{i}<x^{\prime}\right]$ if $x_{i} \neq x^{\prime}$, and when $x_{i}=x^{\prime}$ we have $\sigma_{i, j^{\prime}}^{\prime} \geq \mathbb{1}\left[x_{i}<x^{\prime}\right]$ (and $y_{i}=-1$ ). The above RHS equals

$$
\begin{aligned}
\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \leq x\right] x_{i}+\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x<x_{i}<x^{\prime}\right] x_{i} & \leq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \leq x\right] x_{i} \\
& \leq \sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j}^{\prime} x_{i} \\
& =\frac{1}{v_{j}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime} x_{i}\right) \\
& =\frac{1}{v_{j}} \cdot w_{j}
\end{aligned}
$$

We conclude that

$$
\begin{equation*}
w_{j} \geq \frac{v_{j}}{v_{j^{\prime}}} \cdot w_{j^{\prime}} \tag{37}
\end{equation*}
$$

Likewise, by Eq. (28) we have

$$
\begin{aligned}
\frac{1}{v_{j^{\prime}}} \cdot b_{j^{\prime}} & =\frac{1}{v_{j^{\prime}}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j^{\prime}} \sigma_{i, j^{\prime}}^{\prime}\right)=\sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j^{\prime}}^{\prime} \leq \sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i}<x^{\prime}\right] \\
& =\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x_{i} \leq x\right]+\sum_{i \in I} \lambda_{i} y_{i} \mathbb{1}\left[x<x_{i}<x^{\prime}\right] \\
& \leq \sum_{i \in I} \lambda_{i} y_{i} \sigma_{i, j}^{\prime}=\frac{1}{v_{j}}\left(\sum_{i \in I} \lambda_{i} y_{i} v_{j} \sigma_{i, j}^{\prime}\right)=\frac{1}{v_{j}} \cdot b_{j}
\end{aligned}
$$

Hence, we conclude that

$$
\begin{equation*}
b_{j} \geq \frac{v_{j}}{v_{j^{\prime}}} \cdot b_{j^{\prime}} \tag{38}
\end{equation*}
$$

Since $0 \leq x<x^{\prime}$, and by using Eq. (37) and (38), we have

$$
0>w_{j} \cdot x^{\prime}+b_{j} \geq \frac{v_{j}}{v_{j^{\prime}}} \cdot w_{j^{\prime}} \cdot x^{\prime}+\frac{v_{j}}{v_{j^{\prime}}} \cdot b_{j^{\prime}}=\frac{v_{j}}{v_{j^{\prime}}}\left(w_{j^{\prime}} \cdot x^{\prime}+b_{j^{\prime}}\right)=0
$$

Thus, we reached a contradiction.
We use the notations $\mathcal{I}_{a, b}:=\{a, a+1, \ldots, b\} \subseteq I$ and $\mathcal{I}_{a, b}^{\prime}:=\left\{i \in \mathcal{I}_{a, b}: y_{i} \mathcal{N}_{\boldsymbol{\theta}}\left(x_{i}\right)=1\right\}=\mathcal{I} \cap I^{\prime}$. Thus, $\mathcal{I}^{\prime}$ are the indices of the examples in the interval $\left[x_{a}, x_{b}\right]$ where $\mathcal{N}_{\boldsymbol{\theta}}$ attains margin of exactly 1 . We denote $\mathcal{I}_{a, b}^{\prime}=\left\{i_{1}, \ldots, i_{m}\right\}$, where $a \leq i_{1}<\ldots<i_{m} \leq b$.
Lemma E.7. Suppose that $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=-1$. There are at most 2 indices $\ell \in[m-1]$ such that $\mathcal{N}_{\boldsymbol{\theta}}(x)<-1$ for some $\bar{x} \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$.

Proof. Assume that $\mathcal{N}_{\boldsymbol{\theta}}(x)<-1$ for $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$. Since $\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell}}\right)=\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+1}}\right)=-1$, then we have $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$. Now, since $\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell}}\right)=\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+1}}\right)=-1$ and $\mathcal{N}_{\boldsymbol{\theta}}(x)<-1$ for some $x \in$ $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$, then there must be an activation point in $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$ where the derivative increases. Since by Lemma E. 6 there are at most two such activation points in $\left[x_{a}, x_{b}\right]$ then the lemma follows.

Lemma E.8. Suppose that $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=-1$. There are at most 5 indices $\ell \in[m-1]$ such that $\mathcal{N}_{\boldsymbol{\theta}}(x)>-1$ for some $\bar{x} \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$.

Proof. Assume that $\mathcal{N}_{\boldsymbol{\theta}}(x)>-1$ for $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$. Since $\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell}}\right)=\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+1}}\right)=-1$, then $x \in\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right)$. If $\ell \neq m-1$ then we have $\mathcal{N}_{\boldsymbol{\theta}}(x)>-1$ and $\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+1}}\right)=\mathcal{N}_{\boldsymbol{\theta}}\left(x_{i_{\ell+2}}\right)=-1$. Hence, the interval ( $x, x_{i_{\ell+1}}$ ) contains a point with negative derivative, and the interval ( $x_{i_{\ell+1}}, x_{i_{\ell+2}}$ ) contains a point with non-negative derivative. Thus, there must be an activation point in ( $x, x_{i_{\ell+2}}$ ) where the derivative increases. Therefore, there is an activation point with increasing derivative either in the interval $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ or in the interval $\left[x_{i_{\ell+1}}, x_{i_{\ell+2}}\right)$ (and possibly in both). Thus, an interval $\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ for $\ell \neq m-1$ might contain some $x$ with $\mathcal{N}_{\boldsymbol{\theta}}(x)>-1$ only if there is an activation point
with increasing derivative in $\left(x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ or $\left[x_{i_{\ell+1}}, x_{i_{\ell+2}}\right)$. Since by Lemma E. 6 there are at most two such activation points in $\left[x_{a}, x_{b}\right]$, then there are at most 4 intervals $\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$ with $\ell \neq m-1$ that contain some $x$ with $\mathcal{N}_{\boldsymbol{\theta}}(x)>-1$. The interval $\left[x_{i_{m-1}}, x_{i_{m}}\right]$ might also contain such $x$. Overall, there are at most 5 indices $\ell \in[m-1]$ such that $\mathcal{N}_{\boldsymbol{\theta}}(x)>-1$ for some $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$.

Lemma E.9. Suppose that $x_{b}>x_{a} \geq 0$ and for all $a \leq i \leq b$ we have $y_{i}=-1$. There are at most 30 boundaries between linear regions in $\left[x_{i_{1}}, x_{i_{m}}\right]$.

Proof. The proof is similar to the proof of Lemma E.5. The only difference is that here we use Lemmas E. 6 and E. 7 in order to conclude that there are at most 7 indices $\ell \in[m-1]$ such that $\mathcal{N}_{\boldsymbol{\theta}}(x) \neq-1$ for some $x \in\left[x_{i_{\ell}}, x_{i_{\ell+1}}\right]$, and denote the set of these indices by $R$.


[^0]:    ${ }^{5}$ We note that these intervals overlap and thus do contain instances that change classification with respect to the teacher network, but not with respect to the sample.

[^1]:    ${ }^{6}$ To show this inequality holds, consider $k$ i.i.d. random variables $X_{j} \sim U([0,1])$. Then the left-hand side equals $\mathbb{P}\left[\left|\left\{x_{j}: x_{j} \in[0, x]\right\}\right| \leq 1\right]$. The occurrence of the complement of this event is implied if $x_{i} \in[0, x / 2]$ and $x_{i^{\prime}} \in[x / 2, x]$ hold for some $i \neq i^{\prime}$, therefore to upper bound the left-hand side it suffices to upper bound the complement of the event where $x_{i} \in[0, x / 2]$ and $x_{i^{\prime}} \in[x / 2, x]$ hold for some $i \neq i^{\prime}$. This in turn follows from applying a union bound on $\mathbb{P}\left[\left|\left\{x_{j}: x_{j} \in[0, x / 2]\right\}\right|=0\right]$ and $\mathbb{P}\left[\left|\left\{x_{j}: x_{j} \in[x / 2, x]\right\}\right|=0\right]$.

[^2]:    ${ }^{7}$ Essentially, this holds true since the median of a binomially-distributed random variable $B(n, k)$ is $\lceil n k\rceil$ or $\lfloor n k\rfloor$, and since deviating from the median by at most 1 never increases the probability of the tail to more than 0.75 .

