## A Detailed comparisons with related work

Here, we describe comparisons of our results to those in the literature and give detailed comments on the specific rates we achieve. In Table 1, we compare our agnostic learning results. We note the guarantees for the population risk in the fourth column, the marginal distributions over $x$ for which the bounds hold in the fifth column, and the sample complexity required to reach the specified level of risk plus some $\varepsilon>0$ in the final column. Our results in this setting come from Theorem3.3. The Big-O notation hides constants that may depend on the parameters of the distribution or activation function, but does not hide explicit dependence on the dimension $d$. However, the parameters of the distribution itself may have implicit dependence on the dimension. In particular, for bounded distributions that satisfy $\|x\|_{2} \leq B_{X}$, the $O()$ hides multiplicative factors that depend on $B_{X}$. This means that if $B_{X}$ depends on $d$, so will our bounds. For ReLU, the $O()$ hides polynomial factors in $B_{X}$. For non-ReLU, the worst-case activation functions under consideration in Assumption 3.1 (e.g. the sigmoid) can have $\gamma \sim \exp \left(-B_{X}\right)$, making the runtime and sample complexity depend on $\gamma^{-1} \sim \exp \left(B_{X}\right)$, in which case it is preferable for $B_{X}$ to be a constant independent of the dimension. We note that the sample complexity for Diakonikolas et al. [8] for the $(1+\delta)$ OPT guarantee is $O\left(\varepsilon^{-2}\left[d \delta^{-3} \nu^{-2}\right]^{\delta^{-3}}\right)$ when $\mathcal{D}_{x}$ is $\nu$ sub-Gaussian for some $\nu=O(1)$, and thus the exact dependence on the dimension depends on the sub-Gaussian norm and error threshold desired.
In Table 2, we provide comparisons of our noisy teacher network setting, where $y=\sigma\left(v^{\top} x\right)+\xi$ for some zero mean noise $\xi$. Our results in this setting come from Theorem 4.1. The complexity column here denotes the sample complexity required to reach population risk $\mathrm{OPT}+\varepsilon$. The subspace eigenvalue assumption given by Mukherjee and Muthukumar [26] is that $\mathbb{E}\left[x x^{\top} \mathbb{1}\left(v^{\top} x \geq 0\right)\right] \succ 0$; a similar assumption was used by Du et al. [10] in the realizable setting. We note that the result of Mukherjee and Muthukumar holds for any bounded noise distribution and thus is in the more general adversarial noise (but not agnostic ${ }^{3}$ ) setting.
Finally, in Table 3, we provide comparisons with results in the realizable setting $(\xi \equiv 0)$. (Our results in this setting are given in Theorem E.1 in Appendix E]) For G.D. and S.G.D., the complexity column denotes the sample complexity required to reach population risk $\varepsilon$. For G.D. or gradient flow on the population risk, it refers to the runtime complexity only as there are no samples in this setting. For Du et al. [10], the subspace eigenvalue assumption is that for any $w$ and for the target neuron $v$, it holds that $\mathbb{E}\left[x x^{\top} \mathbb{1}\left(w^{\top} x \geq 0, v^{\top} x \geq\right)\right] \succ 0$. This is a nondegeneracy assumption that is related to the marginal spread condition given in Assumption 3.2, in the sense that it allows for one to show that $H$ is an upper bound for $G$. Finally, we note that any result in the agnostic or noisy teacher network settings applies in the realizable setting as well.

## B Proof of Lemma 3.5

To prove Lemma 3.5, we use the following result of Yehudai and Shamir [35].
Lemma B. 1 (Lemma B.1, [35]). Under Assumption 3.2, for any two vectors $a, b \in \mathbb{R}^{2}$ satisfying $\theta(a, b) \leq \pi-\delta$ for $\delta \in(0, \pi]$, it holds that

$$
\inf _{u \in \mathbb{R}^{2}:\|u\|=1} \int\left(u^{\top} y\right)^{2} \mathbb{1}\left(a^{\top} y \geq 0, b^{\top} y \geq 0,\|y\| \leq \alpha\right) d y \geq \frac{\alpha^{4}}{8 \sqrt{2}} \sin ^{3}(\delta / 4)
$$

Proof of Lemma 3.5. We first consider the case when $\sigma$ satisfies Assumption 3.1. By assumption,

$$
H(w)=(1 / 2) \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x\right)\right] \leq \varepsilon
$$

[^0]Table 1: Comparison of results in the agnostic setting

| Algorithm | Activations | Pop. risk | $\mathcal{D}_{x}$ | Sample <br> Complexity |
| :--- | :--- | :--- | :--- | :--- |
| Halfspace <br> reduction [14] | ReLU | $O\left(\mathrm{OPT}^{2 / 3}\right)$ | standard <br> Gaussian | $O\left(\right.$ poly $\left.\left(d, \varepsilon^{-1}\right)\right)$ |
| Convex <br> surrogate G.D. [8] | ReLU | $O(\mathrm{OPT})$ | isotropic <br> + log-concave | $O\left(d \varepsilon^{-2}\right)$ |
| Convex <br> surrogate G.D. <br> + Domain Partition [8] | ReLU | $(1+\delta) \mathrm{OPT}$ | sub-Gaussian | $O\left(d^{c} \varepsilon^{-2}\right)$ |
| Gradient Descent <br> (This paper) | strictly <br> increasing <br> + Lipschitz | $O(\mathrm{OPT})$ | bounded | $O\left(\varepsilon^{-2}\right)$ |
| Gradient Descent <br> ReLU | $O\left(\mathrm{OPT}^{1 / 2}\right)$ | bounded |  | marginal <br> spread |

Table 2: Comparison of results in the noisy teacher network setting

| Paper | Algorithm | Activations | $\mathcal{D}_{x}$ | Sample <br> Complexity |
| :--- | :--- | :--- | :--- | :--- |
| K. et al. [20] | GLMTron | increasing <br> + Lipschitz | bounded | $O\left(\varepsilon^{-2}\right)$ |
| M. \& M. [26] | Modified <br> S.G.D. | ReLU | bounded <br> + subspace eigenvalue <br> bounded | $O(\log (1 / \varepsilon))$ |
| F. et al. [11] | Meta-algo. | strictly <br> increasing <br> + Lipschitz <br> $+\sigma^{\prime}$ Lipschitz |  | $O\left(\varepsilon^{-2} \wedge d \varepsilon^{-1}\right)$ |
|  |  | strictly <br> increasing <br> + diff'ble <br> + Lipschitz <br> $+\sigma^{\prime}$ Lipschitz <br> $+\sigma^{\prime \prime}$ Lipschitz | centered <br> + sub-Gaussian <br> $+\mathbb{E}\left[x x^{\top}\right] \succ 0$ | $O\left(d \varepsilon^{-1}\right)$ |

Since the term in the expectation is nonnegative, restricting the integral to a smaller set only decreases its value, so that

$$
\begin{equation*}
(1 / 2) \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x\right) \mathbb{1}\left(\left|w_{t}^{\top} x\right| \leq \rho\right)\right] \leq \varepsilon \tag{B.1}
\end{equation*}
$$

For $\rho=B W$, since $\|w\|_{2} \leq W$, the inclusion $\left\{\|x\|_{2} \leq \rho / W\right\}=\left\{\|x\|_{2} \leq B\right\} \subset\left\{\left|w_{t}^{\top} x\right| \leq \rho\right\}$ holds. This means we can lower bound (B.1) by substituting the indicator $\mathbb{1}\left(\left|w_{t}^{\top} x\right| \leq \rho\right)$ with $\mathbb{1}\left(\|x\|_{2} \leq B\right)$, which is identically one by assumption. Since $H(w) \leq \varepsilon$, this implies

$$
\frac{\gamma}{2} \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2}\right] \leq(1 / 2) \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x\right) \mathbb{1}\left(\|x\|_{2} \leq B\right)\right] \leq \varepsilon
$$

Table 3: Comparison of results in the realizable setting

| Paper | Algorithm | Activations | $\mathcal{D}_{x}$ | Sample Complexity |
| :---: | :---: | :---: | :---: | :---: |
| D. et al. [10] | S.G.D. | ReLU | bounded <br> + subspace eigenvalue | $O(\log (1 / \varepsilon))$ |
| S. [30] | Projected Regularized G.D. | ReLU | standard <br> Gaussian | $O(\log (1 / \varepsilon))$ |
| Y. \& S. [35] | Pop. G.D. | leaky ReLU | $\begin{aligned} & \text { bounded } \\ & +\mathbb{E}\left[x x^{\top}\right] \succ 0 \end{aligned}$ | $O(\log (1 / \varepsilon))$ |
| Y. \& S | Pop. G.D. | $\begin{aligned} & \inf _{0<z<\alpha} \sigma^{\prime}(z)>0 \\ & + \text { Lipschitz } \end{aligned}$ | bounded <br> + marginal spread | $O(\log (1 / \varepsilon))$ |
| Y. \& S | Population <br> Gradient <br> Flow | ReLU | marginal spread <br> + spherical symmetry | $O(\log (1 / \varepsilon))$ |
| Y. \& S | S.G.D. | $\begin{aligned} & \inf _{0<z<\alpha} \sigma^{\prime}(z)>0 \\ & + \text { Lipschitz } \end{aligned}$ | bounded <br> + marginal spread | $\tilde{O}\left(\varepsilon^{-2}\right)$ |
| This paper | $\begin{aligned} & \text { Pop. G.D. } \\ & \text { + S.G.D. } \end{aligned}$ | strictly increasing + Lipschitz | bounded | $O\left(\varepsilon^{-1}\right)$ |
| This paper | $\begin{aligned} & \text { Pop. G.D. } \\ & \text { + S.G.D. } \end{aligned}$ | ReLU | bounded <br> + marginal spread | $O\left(\varepsilon^{-1}\right)$ |

Dividing both sides by $\gamma$ completes this part of the proof.
For ReLU, let us assume that $H(w) \leq \varepsilon$, and denote the event

$$
K_{w, v}:=\left\{w^{\top} x \geq 0, v^{\top} x \geq 0\right\}
$$

and define $\zeta:=\beta \alpha^{4} / 8 \sqrt{2}$. Since $H(w)=\mathbb{E}\left[\left(\sigma\left(w^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \mathbb{1}\left(w^{\top} x \geq 0\right)\right] \leq \zeta \varepsilon$, it holds that

$$
\begin{equation*}
\mathbb{E}\left[\left(\sigma\left(w^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \mathbb{1}\left(K_{w, v}\right)\right] \leq \zeta \varepsilon . \tag{B.2}
\end{equation*}
$$

Denote $\widehat{w}$ and $\widehat{v}$ as the projections of $w$ and $v$ respectively onto the two dimensional subspace $\operatorname{span}(w, v)$. Using a proof similar to that of Yehudai and Shamir [35], we have

$$
\begin{align*}
& \mathbb{E}_{x \sim \mathcal{D}}\left[\left(w^{\top} x-v^{\top} x\right)^{2} \mathbb{1}\left(K_{w, v}\right)\right]=\|w-v\|_{2}^{2} \mathbb{E}_{x \sim \mathcal{D}}\left[\left(\left(\frac{w-v}{\|w-v\|_{2}}\right)^{\top} x\right)^{2} \mathbb{1}\left(K_{w, v}\right)\right] \\
& \geq\|w-v\|_{2}^{2} \inf _{u \in \operatorname{span}(w, v),\|u\|=1} \mathbb{E}_{x}\left[\mathbb{1}\left(u^{\top} x\right)^{2} \mathbb{1}\left(K_{w, v}\right)\right] \\
& =\|w-v\|_{2}^{2} \inf _{u \in \mathbb{R}^{2},\|u\|=1} \mathbb{E}_{y \sim \mathcal{D}_{w, v}}\left[\left(u^{\top} y\right)^{2} \mathbb{1}\left(\widehat{w}^{\top} y \geq 0, \widehat{v}^{\top} y \geq 0\right)\right] \\
& \geq\|w-v\|_{2}^{2} \inf _{u \in \mathbb{R}^{2},\|u\|=1} \int\left(u^{\top} y\right)^{2} \mathbb{1}\left(\widehat{w}^{\top} y \geq 0, \widehat{v}^{\top} y \geq 0,\|y\|_{2} \leq \alpha\right) p_{w, v}(y) d y \\
& \geq \beta\|w-v\|_{2}^{2} \inf _{u \in \mathbb{R}^{2},\|u\|=1} \int\left(u^{\top} y\right)^{2} \mathbb{1}\left(\widehat{w}^{\top} y \geq 0, \widehat{v}^{\top} y \geq 0,\|y\|_{2} \leq \alpha\right) d y \tag{B.3}
\end{align*}
$$

By assumption, $\|w-v\|_{2} \leq 1$. Since

$$
1 \geq\|w-v\|_{2}^{2}=\|w\|_{2}\left(\|w\|_{2}-2 \cos \theta(w, v)\right)+1
$$

we must have either $w=0$ or $\theta(w, v) \in[0, \pi / 2]$. To see that $w=0$ is impossible, suppose for the contradiction that $w=0$ and so $H(w)=H(0) \leq \zeta \varepsilon$. Let $z$ be any vector orthogonal to $v$, so that
$\theta(v, z)=\pi / 2$. Then,

$$
\begin{align*}
\zeta \varepsilon & \geq H(0) \\
& =\mathbb{E}_{x \sim \mathcal{D}}\left[\left(v^{\top} x\right)^{2} \mathbb{1}\left(v^{\top} x \geq 0\right)\right] \\
& =\mathbb{E}_{y \sim \mathcal{D}_{0, v}}\left[\left(\widehat{v}^{\top} y\right)^{2} \mathbb{1}\left(\widehat{v}^{\top} y \geq 0\right]\right. \\
& \geq \inf _{u:\|u\|=1} \int\left(u^{\top} x\right)^{2} \mathbb{1}\left(v^{\top} x \geq 0, z^{\top} x \geq 0,\|y\|_{2} \leq \alpha\right) p_{0, v}(y) d y \\
& \geq \beta \inf _{u:\|u\|=1} \int\left(u^{\top} x\right)^{2} \mathbb{1}\left(v^{\top} x \geq 0, z^{\top} x \geq 0,\|y\|_{2} \leq \alpha\right) d y \\
& \geq \frac{\beta \alpha^{4}}{8 \sqrt{2}} . \tag{B.4}
\end{align*}
$$

The last line follows by using Lemma B.1. For $\varepsilon<1$, this is impossible by the definition of $\zeta$. This shows that $\theta(w, v) \leq \pi / 2$. We can therefore apply Lemma B.1 to (B.3) to get

$$
\begin{aligned}
\zeta \varepsilon & \geq \beta\|w-v\|_{2}^{2} \inf _{u \in \mathbb{R}^{2},\|u\|=1} \int\left(u^{\top} y\right)^{2} \mathbb{1}\left(\widehat{w}^{\top} y \geq 0, \widehat{v}^{\top} y \geq 0,\|y\|_{2} \leq \alpha\right) d y \\
& \geq \frac{\beta \alpha^{4}}{8 \sqrt{2}}\|w-v\|_{2}^{2} \\
& =\zeta B^{2}\|w-v\|_{2}^{2} .
\end{aligned}
$$

This shows that $\|w-v\|_{2}^{2} \leq B^{-2} \varepsilon$. Since $\sigma$ is 1-Lipschitz, Hölder's inequality and $\mathbb{E}\|x\|_{2}^{2} \leq B^{2}$ imply that $G(w) \leq \varepsilon$.

## C Proofs for ReLU activation in agnostic setting

In this section we prove Lemma 3.10, which we reproduce below for the reader's convenience.
Lemma C.1. Suppose that $\|x\|_{2} \leq B_{X}$ a.s. under $\mathcal{D}_{x}$. Suppose $\sigma$ is non-decreasing and $L$-Lipschitz. Assume $\widehat{F}(v) \in(0,1)$. Gradient descent with fixed step size $\eta \leq(1 / 4) L^{-2} B_{X}^{-2}$ initialized at $w_{0}=0$ finds weights $w_{t}$ satisfying $\widehat{H}\left(w_{t}\right) \leq 2 L^{2} B_{X} \widehat{F}(v)^{1 / 2}$ within $T=\left\lceil\eta^{-1} L^{-1} B_{X}^{-1} \widehat{F}(v)^{-1 / 2}\right\rceil$ iterations, with $\left\|w_{t}-v\right\|_{2} \leq 1$ for each $t=0, \ldots, T-1$.

As mentioned in the main section, the proof of Lemma 3.6 heavily relied upon Fact 3.7 , which is only satisfied if the activation is strictly increasing. When $\sigma$ is only non-decreasing, we can get a similar lower bound as we did in (3.8) if we use the following fact.
Fact C.2. If $\sigma$ is non-decreasing and $L$-Lipschitz, then for any $z_{1}, z_{2}$ in the domain of $\sigma$, it holds that $\left(\sigma\left(z_{1}\right)-\sigma\left(z_{2}\right)\right)\left(z_{1}-z_{2}\right) \geq L^{-1}\left(\sigma\left(z_{1}\right)-\sigma\left(z_{2}\right)\right)^{2}$.

With this fact we can present the analogue to Lemma 3.6 that holds for a general non-decreasing and Lipschitz activation and hence includes the ReLU.

Proof. Just as in the proof of Lemma 3.6, the lemma is proved if we can show that for every $t \in \mathbb{N}$, either (a) $\widehat{H}\left(w_{\tau}\right) \leq 2 L^{2} B_{X} \widehat{F}(v)^{1 / 2}$ for some $\tau<t$, or (b) $\left\|w_{t}-v\right\|_{2}^{2} \leq\left\|w_{t-1}-v\right\|_{2}^{2}-$ $\eta L B_{X} \widehat{F}(v)^{1 / 2}$ holds. To this end we assume the induction hypothesis holds for some $t \in \mathbb{N}$, and since we are done if (a) holds, we assume (a) does not hold and thus for every $\tau \leq t$, we have $\widehat{H}\left(w_{\tau}\right)>2 L^{2} B_{X} \widehat{F}(v)^{1 / 2}$. Since (a) does not hold, $\left\|w_{\tau}-v\right\|_{2}^{2} \leq\left\|w_{\tau-1}-v\right\|_{2}^{2}-\eta L B_{X} \widehat{F}(v)^{1 / 2}$ holds for each $\tau=1, \ldots, t$ and hence the identity

$$
\begin{equation*}
\left\|w_{\tau}-v\right\|_{2} \leq 1 \quad \forall \tau \leq t \tag{C.1}
\end{equation*}
$$

holds. We now proceed with showing the analogues of (3.8) and (3.9) We begin with the lower bound,

$$
\left\langle\nabla \widehat{F}\left(w_{t}\right), w_{t}-v\right\rangle=(1 / n) \sum_{i=1}^{n}\left(\sigma\left(w_{t}^{\top} x_{i}\right)-\sigma\left(v^{\top} x_{i}\right)\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right)\left(w_{t}^{\top} x_{i}-v^{\top} x_{i}\right)
$$

$$
\begin{align*}
& +\left\langle(1 / n) \sum_{i=1}^{n}\left(\sigma\left(v^{\top} x_{i}\right)-y_{i}\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) x_{i}, w_{t}-v\right\rangle  \tag{C.2}\\
\geq & (1 / L n) \sum_{i=1}^{n}\left(\sigma\left(w_{t}^{\top} x_{i}\right)-\sigma\left(v^{\top} x_{i}\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) \\
& -\left\|w_{t}-v\right\|_{2}\left\|(1 / n) \sum_{i=1}^{n}\left(\sigma\left(v^{\top} x_{i}\right)-y_{i}\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) x_{i}\right\|_{2} \\
\geq & 2 L^{-1} \widehat{H}\left(w_{t}\right)-L B_{X} \widehat{F}(v)^{1 / 2} . \tag{C.3}
\end{align*}
$$

In the first inequality, we have used Fact C. 2 and that $\sigma^{\prime}(z) \geq 0$ for the first term. For the second term, we use Cauchy-Schwarz. The last inequality is a consequence of (C.1) Cauchy-Schwarz, and that $\sigma^{\prime}(z) \leq L$ and $\|x\|_{2} \leq B_{X}$. As for the gradient upper bound at $w_{t}$, the bound (3.9) still holds since it only uses that $\sigma$ is $L$-Lipschitz. The choice of $\eta \leq(1 / 4) L^{-2} B_{X}^{-2}$ then ensures

$$
\begin{align*}
\left\|w_{t}-v\right\|_{2}^{2}-\left\|w_{t+1}-v\right\|_{2}^{2} \geq & 2 \eta\left(2 L^{-1} \widehat{H}\left(w_{t}\right)-L B_{X} \widehat{F}(v)^{1 / 2}\right) \\
& -\eta^{2}\left(4 B_{X}^{2} L \widehat{H}\left(w_{t}\right)+4 L^{2} B_{X}^{2} \widehat{F}(v)\right) \\
\geq & \eta\left(3 L^{-1} \widehat{H}\left(w_{t}\right)-3 L B_{X}\left(\widehat{F}(v) \vee \widehat{F}(v)^{1 / 2}\right)\right) \\
\geq & \eta L B_{X} \widehat{F}(v)^{1 / 2} \tag{C.4}
\end{align*}
$$

where the last line comes from the induction hypothesis that $\widehat{H}\left(w_{t}\right) \geq 2 L^{2} B_{X} \widehat{F}(v)^{1 / 2}$ and since $\widehat{F}(v) \in(0,1)$. This completes the proof.

With the above lemma, we can prove the ReLU case of Theorem 3.3 .

Proof of Theorem 3.3 for ReLU. We highlight here the main technical differences with the proof for the strictly increasing case. Although Lemma 3.9 applies to the loss function $\ell(w ; x)=$ $(1 / 2)\left(\sigma\left(w^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2}$, the same results hold for the loss function $\tilde{\ell}(w ; x)=\ell(w ; x) \sigma^{\prime}\left(w^{\top} x\right)$ for ReLU, since $\nabla \sigma^{\prime}\left(w^{\top} x\right) \equiv 0$ a.e. Thus $\tilde{\ell}$ is still $B_{X}$-Lipschitz, and we have

$$
\begin{equation*}
\mathbb{E}_{S \sim \mathcal{D}^{n}} \mathfrak{R}_{S}(\tilde{\ell} \circ \sigma \circ \mathcal{G}) \leq \frac{2 B_{X}^{2}}{\sqrt{n}} \tag{C.5}
\end{equation*}
$$

With this in hand, the proof is essentially identical: By Lemmas 3.10 and 3.8 , with probability at least $1-\delta / 2$ gradient descent finds a point with

$$
\begin{equation*}
\widehat{H}\left(w_{t}\right) \leq 2 B_{X} \widehat{F}(v)^{1 / 2} \leq 2 B_{X}\left(\mathrm{OPT}^{1 / 2}+\frac{\sqrt{3 a} \log ^{1 / 4}(4 / \delta)}{n^{1 / 4}}\right) \tag{C.6}
\end{equation*}
$$

We can then use (C.5) to get that with probability at least $1-\delta$,

$$
\begin{equation*}
H\left(w_{t}\right) \leq 2 B_{X}\left(\mathrm{OPT}^{1 / 2}+\frac{\sqrt{3 a} \log ^{1 / 4}(4 / \delta)}{n^{1 / 4}}\right)+\frac{2 B_{X}^{2}}{\sqrt{n}}+2 B_{X}^{2} \sqrt{\frac{2 \log (8 / \delta)}{n}} \tag{C.7}
\end{equation*}
$$

Since $\mathcal{D}_{x}$ satisfies Assumption 3.2 and $\left\|w_{t}-v\right\|_{2} \leq 1$, Lemma 3.5 yields $G\left(w_{t}\right) \leq$ $8 \sqrt{2} \alpha^{-4} \beta^{-1} H\left(w_{t}\right)$. Then applying Claim 3.4 completes the proof.

## D Noisy teacher network proofs

As in the agnostic case, we have a key lemma that shows $\widehat{H}$ is small when we run gradient descent for a sufficiently large time. Note that one difference with the proof in the agnostic case is that we do not need to consider different auxiliary errors for the strictly increasing and ReLU cases; $H$ alone suffices.

Lemma D.1. Suppose that $\|x\|_{2} \leq B_{X}$ a.s. under $\mathcal{D}_{x}$. Let $\sigma$ be non-decreasing and $L$-Lipschitz. Suppose that the bound

$$
\begin{equation*}
\left\|(1 / n) \sum_{i=1}^{n}\left(\sigma\left(v^{\top} x_{i}\right)-y_{i}\right) \alpha_{i} x_{i}\right\|_{2} \leq K \leq 1 \tag{D.1}
\end{equation*}
$$

holds for scalars satisfying $\alpha_{i} \in[0, L]$. Then gradient descent run with fixed step size $\eta \leq$ $(1 / 4) L^{-2} B_{X}^{-2}$ from initialization $w_{0}=0$ finds weights $w_{t}$ satisfying $\hat{H}\left(w_{t}\right) \leq 4 L K$ within $T=\left\lceil\eta^{-1} K^{-1}\right\rceil$ iterations, with $\left\|w_{t}-v\right\|_{2} \leq 1$ for each $t=0, \ldots, T-1$.

Proof. Just as in the proof of Lemma 3.6, the theorem can be shown by proving the following induction statement. We claim that for every $t \in \mathbb{N}$, either (a) $\widehat{H}\left(w_{\tau}\right) \leq 4 L K$ for some $\tau<t$, or (b) $\left\|w_{t}-v\right\|_{2}^{2} \leq\left\|w_{t-1}-v\right\|_{2}^{2}-\eta K$. If the induction hypothesis holds, then until gradient descent finds a point where $\widehat{H}\left(w_{t}\right) \leq 4 L K$, the squared distance $\left\|w_{t}-v\right\|_{2}^{2}$ decreases by $\eta K$ at every iteration. Since $\left\|w_{0}-v\right\|_{2}^{2}=1$, this means there can be at most $\eta^{-1} K^{-1}$ iterations until we reach $\widehat{H}\left(w_{t}\right) \leq 4 L K$. This shows the induction statement implies the theorem.
We begin with the proof by supposing the induction hypothesis holds for $t$, and considering the case $t+1$. If (a) holds, then we are done. So now consider the case that for every $\tau \leq t$, we have $\widehat{H}\left(w_{\tau}\right)>4 L K$. Since (a) does not hold, $\left\|w_{\tau}-v\right\|_{2}^{2} \leq\left\|w_{\tau-1}-v\right\|_{2}^{2}-\eta K$ holds for each $\tau=1, \ldots, t$. Since $\left\|w_{0}-v\right\|_{2}=1$, this implies

$$
\begin{equation*}
\left\|w_{\tau}-v\right\|_{2} \leq 1 \forall \tau \leq t \tag{D.2}
\end{equation*}
$$

We can therefore bound

$$
\begin{align*}
\left\langle\nabla \widehat{F}\left(w_{t}\right), w_{t}-v\right\rangle= & \left\langle\frac{1}{n} \sum_{1=1}^{n}\left(\sigma\left(w_{t}^{\top} x_{i}\right)-y_{i}\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) x_{i}, w_{t}-v\right\rangle \\
= & \frac{1}{n} \sum_{i=1}^{n}\left(\sigma\left(w_{t}^{\top} x_{i}\right)-\sigma\left(v^{\top} x_{i}\right)\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right)\left(w_{t}^{\top} x_{i}-v^{\top} x_{i}\right) \\
& +\left\langle\frac{1}{n} \sum_{i=1}^{n}\left(\sigma\left(v^{\top} x_{i}\right)-y_{i}\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) x_{i}, w_{t}-v\right\rangle \\
\geq & \frac{L^{-1}}{n} \sum_{i=1}^{n}\left(\sigma\left(w_{t}^{\top} x_{i}\right)-\sigma\left(v^{\top} x_{i}\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right)-K\left\|w_{t}-v\right\|_{2} \\
\geq & 2 L^{-1} \widehat{H}\left(w_{t}\right)-K . \tag{D.3}
\end{align*}
$$

In the first inequality, we have used Fact C. 2 for the first term. For the second term, we use (D.1) and that $\alpha_{i}:=\sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) \in[0, L]$. The last inequality uses (D.2).
For the gradient upper bound, we have

$$
\begin{align*}
\left\|\nabla \widehat{F}\left(w_{t}\right)\right\|_{2}^{2}= & \left\|\frac{1}{n} \sum_{i=1}^{n}\left(\sigma\left(w_{t}^{\top} x_{i}\right)-\sigma\left(v^{\top} x_{i}\right)\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) x_{i}+\frac{1}{n} \sum_{i=1}^{n}\left(\sigma\left(v^{\top} x_{i}\right)-y_{i}\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) x_{i}\right\|_{2}^{2} \\
\leq & 2\left\|\frac{1}{n} \sum_{i=1}^{n}\left(\sigma\left(w_{t}^{\top} x_{i}\right)-\sigma\left(v^{\top} x_{i}\right)\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) x_{i}\right\|_{2}^{2} \\
& +2\left\|\frac{1}{n} \sum_{i=1}^{n}\left(\sigma\left(v^{\top} x_{i}\right)-y_{i}\right) \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right) x_{i}\right\|_{2}^{2} \\
\leq & \frac{2 L B_{X}^{2}}{n} \sum_{i=1}^{n}\left(\sigma\left(w^{\top} x_{i}\right)-\sigma\left(v^{\top} x_{i}\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x_{i}\right)+2 K^{2} \\
= & 4 L B_{X}^{2} \widehat{H}\left(w_{t}\right)+2 K^{2} \tag{D.4}
\end{align*}
$$

The first inequality uses Young's inequality. The second uses that $\sigma^{\prime}(z) \leq L$ and that $\|x\|_{2} \leq B_{X}$ a.s. and (D.1)

Putting (D.3) and (D.4) together, the choice of $\eta \leq(1 / 4) L^{-2} B_{X}^{-2}$ gives us

$$
\begin{aligned}
\left\|w_{t}-v\right\|_{2}^{2}-\left\|w_{t+1}-v\right\|_{2}^{2} & =2 \eta\left\langle\nabla \widehat{F}\left(w_{t}\right), w_{t}-v\right\rangle-\eta^{2}\left\|\nabla \widehat{F}\left(w_{t}\right)\right\|_{2}^{2} \\
& \geq 2 \eta\left(L^{-1} \widehat{H}\left(w_{t}\right)-K\right)-\eta^{2}\left(4 L B_{X}^{2} \widehat{H}\left(w_{t}\right)+2 K^{2}\right) \\
& \geq \eta L^{-1} \widehat{H}\left(w_{t}\right)-3 \eta K .
\end{aligned}
$$

In particular, this implies

$$
\begin{equation*}
\left\|w_{t+1}-v\right\|_{2}^{2} \leq\left\|w_{t}-v\right\|_{2}^{2}+3 \eta K-\eta L^{-1} \widehat{H}\left(w_{t}\right) \tag{D.5}
\end{equation*}
$$

Since $\widehat{H}\left(w_{t}\right)>4 K L$, this completes the induction. The base case follows easily since $\left\|w_{0}-v\right\|_{2}=$ 1 allows for us to deduce the desired bound on $\left\|w_{1}-v\right\|_{2}^{2}$ using (D.5)

To prove a concrete bound on the $K$ term of Lemma D.1, we will need the following definition of norm sub-Gaussian random vectors.
Definition D.2. A random vector $z \in \mathbb{R}^{d}$ is said to be norm sub-Gaussian with parameter $s>0$ if

$$
\mathbb{P}(\|z-\mathbb{E} z\| \geq t) \leq 2 \exp \left(-t^{2} / 2 s^{2}\right)
$$

A Hoeffding-type inequality for norm sub-Gaussian vectors was recently shown by Jin et al. [19].
Lemma D. 3 (Lemma 6, [19]). Suppose $z_{1}, \ldots, z_{n} \in \mathbb{R}^{d}$ are random vectors with filtration $\mathcal{F}_{t}:=$ $\sigma\left(z_{1}, \ldots, z_{t}\right)$ such that $z_{i} \mid \mathcal{F}_{i-1}$ is a zero-mean norm sub-Gaussian vector with parameter $s_{i} \in \mathbb{R}$ for each $i$. Then, there exists an absolute constant $c>0$ such that for any $\delta>0$, with probability at least $1-\delta$,

$$
\left\|\sum_{i=1}^{n} z_{i}\right\| \leq c \sqrt{\log (2 d / \delta) \sum_{i=1}^{n} s_{i}^{2}}
$$

Using this, we can show that if $\xi_{i}:=\sigma\left(v^{\top} x_{i}\right)-y_{i}$ is $s$ sub-Gaussian, then we can get a bound on $K$ at rate $n^{-1 / 2}$. We note that if we make the stronger assumption that $\xi_{i}$ is bounded a.s., we can get rid of the $\log (d)$ dependence by using concentration of bounded random variables in a Hilbert space (e.g. Pinelis and Sakhanenko [27], Corollary 2).

Lemma D.4. Suppose that $\|x\|_{2} \leq B_{X}$ a.s. under $\mathcal{D}_{x}$, and let $\sigma$ be any continuous function. Assume $\xi_{i}:=\sigma\left(v^{\top} x_{i}\right)-y_{i}$ is $s$ sub-Gaussian and satisfies $\mathbb{E}\left[\xi_{i} \mid x_{i}\right]=0$. Then there exists an absolute constant $c_{0}>0$ such that for constants $\alpha_{i} \in[0, L]$, with probability at least $1-\delta$, we have

$$
\left\|(1 / n) \sum_{i=1}^{n}\left(\sigma\left(v^{\top} x_{i}\right)-y_{i}\right) \alpha_{i} x_{i}\right\| \leq c_{0} L B_{X} s \sqrt{n^{-1} \log (2 d / \delta)}
$$

Proof of Lemma D.4. Define $z_{i}:=\left(\sigma\left(v^{\top} x_{i}\right)-y_{i}\right) \alpha_{i} x_{i}$. Using iterated expectations, we see that $\mathbb{E}\left[z_{i}\right]=0$. Since $\sigma\left(v^{\top} x_{i}\right)-y_{i}$ is $s$-sub-Gaussian and $\left\|\alpha_{i} x_{i}\right\|_{2} \leq L B_{X}$, it follows from the definition that $z_{i}$ is norm sub-Gaussian with parameter $L B_{X} s$ for each $i$. By Lemma D.3, we have with probability at least $1-\delta$,

$$
\left\|\sum_{i=1}^{n} z_{i}\right\| \leq c \sqrt{\log (2 d / \delta) L^{2} B_{X}^{2} n s^{2}}
$$

Dividing each side by $n$ proves the lemma.
Proof of Theorem 4.1. By Lemmas D.1 and D.4, there exists some $w_{t}, t<T$ and $\left\|w_{t}-v\right\|_{2} \leq 1$, such that

$$
\widehat{H}\left(w_{t}\right) \leq 4 L K \leq 4 c_{0} L^{2} B_{X} s \sqrt{\frac{\log (2 d / \delta)}{n}} .
$$

Consider $\sigma$ satisfying Assumption 3.1 first, with $\gamma$ corresponding to $\rho=2 B_{X}$. Since $\left\|w_{t}\right\|_{2} \leq 2$, we can use Lemma 3.5 to transform the above bound for $\widehat{H}$ into one for $\widehat{G}$,

$$
\widehat{G}\left(w_{t}\right) \leq 4 c_{0} \gamma^{-1} L^{2} B_{X} s \sqrt{\frac{\log (2 d / \delta)}{n}}
$$

Since $\|w-v\|_{2} \leq 1$ implies $G(w) \leq L^{2} B_{X}^{2} / 2$, standard results from Rademacher complexity imply (e.g. Theorem 26.5 of [28]) that with probability at least $1-\delta$,

$$
G\left(w_{t}\right) \leq \widehat{G}\left(w_{t}\right)+\mathbb{E}_{S \sim \mathcal{D}^{n}} \Re_{S}(\ell \circ \sigma \circ \mathcal{G})+2 L^{2} B_{X}^{2} \sqrt{\frac{2 \log (4 / \delta)}{n}}
$$

where $\ell(w ; x)=(1 / 2)\left(\sigma\left(w^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2}$ and $\mathcal{G}$ are from Lemma 3.9 For the second term above, Lemma 3.9 and rescaling $\delta$ yields that

$$
G\left(w_{t}\right) \leq \frac{2 L^{3} B_{X}^{2}}{\sqrt{n}}+\frac{2 L^{2} B_{X}^{2} \sqrt{2 \log (8 / \delta)}}{\sqrt{n}}+\frac{4 c_{0} \gamma^{-1} L^{2} B_{X} s \sqrt{\log (4 d / \delta)}}{\sqrt{n}}
$$

Then Claim 3.4 completes the proof for strictly increasing $\sigma$.
When $\sigma$ is ReLU, the proof follows the same argument given in the proof of Theorem 3.3. Denoting the loss function $\tilde{\ell}(w ; x)=(1 / 2)\left(\sigma\left(w^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w^{\top} x\right)$, we have

$$
\begin{equation*}
\mathbb{E}_{S \sim \mathcal{D}^{n}} \mathfrak{R}_{S}(\tilde{\ell} \circ \sigma \circ \mathcal{G}) \leq \frac{2 B_{X}^{2}}{\sqrt{n}} \tag{D.6}
\end{equation*}
$$

By Lemmas D.1 and D.4, there exists some $w_{t}, t<T$ and $\left\|w_{t}-v\right\|_{2} \leq 1$, such that

$$
\begin{equation*}
\widehat{H}\left(w_{t}\right) \leq 4 L K \leq 4 c_{0} L^{2} B_{X} s \sqrt{\frac{\log (2 d / \delta)}{n}} . \tag{D.7}
\end{equation*}
$$

Using standard results from Rademacher complexity,

$$
H\left(w_{t}\right) \leq \widehat{H}\left(w_{t}\right)+\mathbb{E}_{S \sim \mathcal{D}^{n}} \Re_{S}(\tilde{\ell} \circ \sigma \circ \mathcal{G})+2 B_{X}^{2} \sqrt{\frac{2 \log (4 / \delta)}{n}}
$$

By (D.6), this means

$$
H\left(w_{t}\right) \leq \frac{4 c_{0} B_{X} s \sqrt{\log (4 d / \delta)}}{\sqrt{n}}+\frac{2 B_{X}^{2}}{\sqrt{n}}+\frac{2 B_{X}^{2} \sqrt{2 \log (8 / \delta)}}{\sqrt{n}}
$$

Since $\mathcal{D}$ satisfies Assumption 3.2 and $\left\|w_{t}-v\right\|_{2} \leq 1$, Lemma 3.5 shows that $G\left(w_{t}\right) \leq$ $8 \sqrt{2} \alpha^{-4} \beta^{-1} H\left(w_{t}\right)$. Then Claim 3.4 translates the bound for $G\left(w_{t}\right)$ into one for $F\left(w_{t}\right)$.

## E Realizable setting

In this section we assume $y=\sigma\left(v^{\top} x\right)$ a.s. for some $\|v\|_{2} \leq 1$. As in the agnostic and noisy teacher network setting, we use the auxiliary loss

$$
H(w):=(1 / 2) \mathbb{E}_{x \sim \mathcal{D}}\left[\left(\sigma\left(w^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w^{\top} x\right)\right]
$$

Note that in the realizable setting, the previous auxiliary loss $G$ defined in (3.4) coincides with the true objective $F$, i.e. we have

$$
F(w):=(1 / 2) \mathbb{E}_{x \sim \mathcal{D}}\left[\left(\sigma\left(w^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2}\right]
$$

For purpose of comparison with Yehudai and Shamir [35], we provide analyses for two settings in the realizable case: in the first setting, we consider gradient descent on the population loss,

$$
\begin{equation*}
w_{t+1}=w_{t}-\eta \nabla F\left(w_{t}\right) \tag{E.1}
\end{equation*}
$$

and return $w_{t^{*}}:=\operatorname{argmin}_{0 \leq t<T} F\left(w_{t}\right)$. The second setting is online SGD with samples $x_{t} \sim$ $\mathcal{D}$. Here we compute unbiased estimates (conditional on $w_{t}$ ) of the population risk $F_{t}\left(w_{t}\right):=$ $(1 / 2)\left(\sigma\left(w_{t}^{\top} x_{t}\right)-\sigma\left(v^{\top} x_{t}\right)\right)^{2}$ and update the weights by

$$
\begin{equation*}
w_{t+1}=w_{t}-\eta \nabla F_{t}\left(w_{t}\right) \tag{E.2}
\end{equation*}
$$

For SGD, we output $w_{t^{*}}=\operatorname{argmin}_{0 \leq t<T} F_{t}\left(w_{t}\right)$.
We summarize our results in the realizable case in Theorem E. 1

Theorem E.1. Suppose $\|x\|_{2} \leq B$ a.s. and $\sigma$ is non-decreasing and $L$-Lipschitz. Let $\eta \leq L^{-2} B^{-2}$ be the step size.
(a) Let $\sigma$ satisfy Assumption 3.1, and let $\gamma$ be the constant corresponding to $\rho=4 B$. For any initialization satisfying $\left\|w_{0}\right\|_{2} \leq 2$, if we run gradient descent on the population risk $T=$ $\left\lceil 2 \varepsilon^{-1} L \eta^{-1} \gamma^{-1}\left\|w_{0}-v\right\|_{2}^{2}\right\rceil$ iterations, then there exists $t<T$ such that $F\left(w_{t}\right) \leq \varepsilon$. For stochastic gradient descent, for any $\delta>0$, running SGD for $\tilde{T}=6 T \log (1 / \delta)$ guarantees there exists $w_{t}, t<T$, such that w.p. at least $1-\delta, F\left(w_{t}\right) \leq \varepsilon$.
(b) Let $\sigma$ be ReLU and further assume that $\mathcal{D}$ satisfies Assumption 3.2 for constants $\alpha, \beta>0$ and that $w_{0}=0$. Let $\nu=\alpha^{4} \beta / 8 \sqrt{2}$. If we run gradient descent on the population risk $T=\left\lceil 2 \varepsilon^{-1} L \eta^{-1} \nu^{-1}\left\|w_{0}-v\right\|_{2}^{2}\right\rceil$ iterations, then there exists $t<T$ such that $F\left(w_{t}\right) \leq \varepsilon$. For stochastic gradient descent, for any $\delta>0$, running SGD for $\tilde{T}=6 T \log (1 / \delta)$ guarantees there exists $w_{t}, t<T$, such that w.p. at least $1-\delta, F\left(w_{t}\right) \leq \varepsilon$.

A few remarks on the above theorem: first, in comparison with our noisy neuron result in Theorem 4.1. we are able to achieve $\mathrm{OPT}+\varepsilon=\varepsilon$ population risk with sample complexity and runtime of order $\varepsilon^{-1}$ rather than $\varepsilon^{-2}$ using the same assumptions by invoking a martingale Bernstein inequality rather than Hoeffding. Second, although TheoremE. 1 requires the distribution to be bounded almost surely, we show in Section E. 1 below that for GD on the population loss, we can accomodate essentially any distribution with finite expected squared norm.
Yehudai and Shamir [35] used the marginal spread assumption to show that with probability $1 / 2$, a single neuron in the realizable setting can be learned using gradient-based optimization with random initialization for Lipschitz activation functions satisfying $\inf _{0<z<\alpha} \sigma^{\prime}(z)>0$, where $\alpha$ is the same constant in Assumption 3.2, and thus includes essentially all neural network activation functions like softplus, sigmoid, tanh, and ReLU. Under the additional assumption of spherical symmetry, they showed that this can be improved to a high probability guarantee for the ReLU activation. For gradient descent on the population risk, they proved linear convergence, i.e. a runtime of order $O(\log (1 / \varepsilon))$, while for SGD their runtime and sample complexity is of order $O\left(\varepsilon^{-2} \log (1 / \varepsilon)\right)$. In comparison, our result for the non-ReLU activations requires only boundedness of the distributions and holds with high probability over random initializations, with runtime and sample complexity of order $O\left(\varepsilon^{-1}\right)$ for both gradient descent on the population risk and SGD. Our results for ReLU use the same marginal spread assumption as Yehudai and Shamir, but our proof technique differs in that we do not require the angle $\theta\left(w_{t}, v\right)$ between the weights in the GD trajectory and the target neuron be decreasing. As they pointed out, angle monotonicity fails to hold for the trajectory of gradient descent even when the distribution is a non-centered Gaussian, so that proofs based on angle monotonicity will not translate to more general distributions. Indeed, our proofs in the agnostic and noisy teacher network setting use essentially the same proof technique as the realizable case without relying on angle monotonicity. Instead, we show a type of inductive bias of gradient descent in the sense that if initialized at the origin, the angle between the target vector and the population risk minimizer cannot become larger than $\pi / 2$, even in the agnostic setting.

## E. 1 Gradient descent on population loss

The key lemma for the proof is as follows.
Lemma E.2. Consider gradient descent on the population risk given in (E.1). Let $w_{0}$ be the initial point of gradient descent and assume $\left\|w_{0}\right\|_{2} \leq 2$. Suppose that $\mathcal{D}$ satisfies $\mathbb{E}_{x}\left[\|x\|_{2}^{2}\right] \leq B^{2}$. Let $\sigma$ be non-decreasing and $L$-Lipschitz. Assume the step size satisfies $\eta \leq L^{-2} B^{-2}$. Then for any $T \in \mathbb{N}$, we have for all $t=0, \ldots, T-1,\left\|w_{t}-v\right\|_{2} \leq\left\|w_{0}-v\right\|_{2}$, and

$$
\left\|w_{0}-v\right\|_{2}^{2}-\left\|w_{T}-v\right\|_{2}^{2} \geq \eta L^{-1} \sum_{t=0}^{T-1} H\left(w_{t}\right)
$$

Proof. We begin with the identity, for $t<T$,

$$
\begin{equation*}
\left\|w_{t}-v\right\|_{2}^{2}-\left\|w_{t+1}-v\right\|_{2}^{2}=2 \eta\left\langle\nabla F\left(w_{t}\right), w_{t}-v\right\rangle-\eta^{2}\left\|\nabla F\left(w_{t}\right)\right\|_{2}^{2} \tag{E.3}
\end{equation*}
$$

First, we have

$$
\begin{aligned}
\left\|\nabla F\left(w_{t}\right)\right\|_{2} & \leq \mathbb{E}_{x}\left\|\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right) \sigma^{\prime}\left(w_{t}^{\top} x\right) x\right\|_{2} \\
& \leq \sqrt{\mathbb{E}_{x}\left[\sigma^{\prime}\left(w_{t}^{\top} x\right)\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2}\right]} \sqrt{\mathbb{E}_{x} \sigma^{\prime}\left(w_{t}^{\top} x\right)\|x\|_{2}^{2}} \\
& \leq B \sqrt{L} \sqrt{\mathbb{E}_{x}\left[\sigma^{\prime}\left(w_{t}^{\top} x\right)\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2}\right]}
\end{aligned}
$$

The first inequality is by Jensen. The second inequality uses that $\sigma^{\prime}(z) \geq 0$ and Hölder, and the third inequality uses that $\sigma$ is $L$-Lipschitz and that $\mathbb{E}\left[\|x\|_{2}^{2}\right] \leq B^{2}$. We therefore have the gradient upper bound

$$
\begin{equation*}
\left\|\nabla F\left(w_{t}\right)\right\|_{2}^{2} \leq 2 B^{2} L H\left(w_{t}\right) \tag{E.4}
\end{equation*}
$$

For the inner product term of (E.3), since $\sigma^{\prime}(z) \geq 0$, we can use Fact C. 2 to get

$$
\begin{equation*}
\left\langle\nabla F\left(w_{t}\right), w_{t}-v\right\rangle \geq L^{-1} \mathbb{E}_{x}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x\right)\right]=2 L^{-1} H\left(w_{t}\right) \tag{E.5}
\end{equation*}
$$

Putting (E.5) and (E.4) into (E.3) we get

$$
\left\|w_{t}-v\right\|_{2}^{2}-\left\|w_{t+1}-v\right\|_{2}^{2} \geq 4 \eta L^{-1} H\left(w_{t}\right)-2 \eta^{2} B^{2} L H\left(w_{t}\right) \geq 2 \eta L^{-1} H\left(w_{t}\right)
$$

where we have used $\eta \leq L^{-2} B^{-2}$. Telescoping the above over $t<T$ gives

$$
\left\|w_{0}-v\right\|_{2}^{2}-\left\|w_{T}-v\right\|_{2}^{2} \geq 2 \eta L^{-1} \sum_{t=0}^{T-1} H\left(w_{t}\right)
$$

Dividing each side by $\eta T$ shows the desired bound.
We will show that if $\sigma$ satisfies Assumption 3.1, then Lemma E. 2 allows for a population risk bound for essentially any distribution with $\mathbb{E}\left[\|x\|_{2}^{2}\right] \leq B^{2}$. In particular, we consider distributions with finite expected norm squared and the possible types of tail bounds for the norm.
Assumption E.3. (a) Bounded distributions: there exists $B>0$ such that $\|x\|_{2} \leq B$ a.s.
(b) Exponential tails: there exist $a_{0}, C_{e}>0$ such that $\mathbb{P}\left(\|x\|_{2}^{2} \geq a\right) \leq C_{e} \exp (-a)$ holds for all $a \geq a_{0}$.
(c) Polynomial tails: there exist $a_{0}, C_{p}>0$ and $\beta>1$ such that $\mathbb{P}\left(\|x\|_{2}^{2} \geq b\right) \leq C_{p} a^{-\beta}$ holds for all $a \geq a_{0}$.

If either (a), (b), or (c) holds, there exists $B>0$ such that $\mathbb{E}\|x\|_{2}^{2} \leq B^{2}$. One can verify that for (b), taking $B^{2}=2\left(a_{0} \vee C_{e}\right)$ suffices, and for (c), $B^{2}=2\left(a_{0} \vee C_{p}^{1 / \beta} /(1-\beta)\right)$ suffices. In fact, any distribution that satisfies $\mathbb{E}\|x\|_{2}^{2}<\infty$ cannot have a tail bound of the form $\mathbb{P}\left(\|x\|_{2}^{2} \geq a\right)=\Omega\left(a^{-1}\right)$, since in this case we would have

$$
\mathbb{E}\|x\|_{2}^{2}=\int_{0}^{\infty} \mathbb{P}\left(\|x\|_{2}^{2}>t\right) d t \geq C \int_{a_{0}}^{\infty} t^{-1} d t=\infty
$$

So the polynomial tail assumption (c) is tight up to logarithmic factors for distributions with finite $\mathbb{E}\|x\|_{2}^{2}$.
Theorem E.4. Let $\mathbb{E}\left[\|x\|_{2}^{2}\right] \leq B^{2}$ and assume $\mathcal{D}$ satisfies one of the conditions in Assumption E. 3 Let $\sigma$ satisfy Assumption 3.1
(a) Under Assumption E.3a, let $\gamma$ be the constant corresponding to $\rho=4 B$ in Assumption 3.1. Running gradient descent for $T=\left\lceil 2 \varepsilon^{-1} L \eta^{-1} \gamma^{-1}\left\|w_{0}-v\right\|_{2}^{2}\right\rceil$ guarantees there exists $t \in$ [T-1] such that $F\left(w_{t}\right) \leq \varepsilon$.
(b) Under Assumption E. 3 , let $\gamma$ be the constant corresponding to $\rho=4 \sqrt{\log \left(18 C_{e} / \varepsilon\right)}$. Running gradient descent for $T=\left\lceil 2 \varepsilon^{-1} L \eta^{-1} \gamma^{-1}\left\|w_{0}-v\right\|_{2}^{2}\right\rceil$ guarantees there exists $t \in[T-1]$ such that $F\left(w_{t}\right) \leq \varepsilon$.
(c) Under Assumption E.3c, let $\gamma$ be the constant corresponding to $\rho=4\left(18 C_{p} / \varepsilon(\beta-1)\right)^{(1-\beta) / 2}$. Running gradient descent for $T=\left\lceil 2 \varepsilon^{-1} L \eta^{-1} \gamma^{-1}\left\|w_{0}-v\right\|_{2}^{2}\right\rceil$ guarantees there exists $t \in$ $[T-1]$ such that $F\left(w_{t}\right) \leq \varepsilon$.

Proof. First, note that the conditions of Lemma E.2 hold, so that we have for all $t=0, \ldots, T-1$, $\left\|w_{t}\right\|_{2} \leq 4$ and

$$
\begin{equation*}
\eta \sum_{t=0}^{T-1} H\left(w_{t}\right) \leq L\left\|w_{0}-v\right\|_{2}^{2}-L\left\|w_{T}-v\right\|_{2}^{2} \tag{E.6}
\end{equation*}
$$

By taking $T=\zeta^{-1} L \varepsilon^{-1} \eta^{-1}\left\|w_{0}-v\right\|_{2}^{2}$ for arbitrary $\zeta>0$,(E.6) implies that there exists $t \in[T-1]$ such that

$$
\begin{equation*}
H\left(w_{t}\right)=\mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x\right)\right] \leq \frac{L\left\|w_{0}-v\right\|_{2}^{2}}{\eta T} \leq \zeta \varepsilon \tag{E.7}
\end{equation*}
$$

It therefore suffices to bound $F\left(w_{t}\right)$ in terms of the left hand side of (E.7). We will do so by using the distributional assumptions given in Assumption E. 3 and by choosing $\zeta$ appropriately.
We begin by noting that (E.7) implies, for any $\rho>0$,

$$
\begin{equation*}
\mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x\right) \mathbb{1}\left(\left|w_{t}^{\top} x\right| \leq \rho\right)\right] \leq \zeta \varepsilon \tag{E.8}
\end{equation*}
$$

For any $\rho>0$, since $\left\|w_{t}\right\|_{2} \leq 4$, the inclusion

$$
\begin{equation*}
\left\{\|x\|_{2} \leq \rho / 4\right\} \subset\left\{\left|w_{t}^{\top} x\right| \leq \rho\right\} \tag{E.9}
\end{equation*}
$$

holds. Under Assumption E. 3 a, by taking $\rho=4 B$ and letting $\gamma$ be the corresponding constant from Assumption 3.1, eqs. (E.8) and (E.9) imply

$$
\gamma \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2}\right] \leq \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x\right) \mathbb{1}\left(\|x\|_{2} \leq \rho / 4\right)\right] \leq \zeta \varepsilon
$$

By taking $\zeta=\gamma / 2$, this implies $F\left(w_{t}\right) \leq \varepsilon$.
Under Assumption E. 3 b, by taking $\rho=4 \sqrt{a_{0}}$, we get

$$
\begin{align*}
\mathbb{E}\left[\|x\|_{2}^{2} \mathbb{1}\left(\|x\|_{2}^{2}>\rho^{2} / 4^{2}\right)\right] & =\int_{a_{0}}^{\infty} \mathbb{P}\left(\|x\|_{2}^{2}>t\right) d t \\
& \leq C_{e} \exp \left(-a_{0}\right) \tag{E.10}
\end{align*}
$$

Note that Assumption E. 3 b holds if we take $a_{0}$ larger. We can therefore let $a_{0}$ be large enough so that $a_{0} \geq \log \left(18 C_{e} / \varepsilon\right)$, so that then

$$
\begin{equation*}
\mathbb{E}\left[\|x\|_{2}^{2} \mathbb{1}\left(\|x\|_{2}^{2}>\rho^{2} / 4^{2}\right)\right] \leq \varepsilon / 18 \tag{E.11}
\end{equation*}
$$

Similarly, under Assumption E.3, we can let $\gamma$ be the constant corresponding to $\rho=4 \sqrt{a_{0}}$ and take $a_{0} \geq\left(\varepsilon(\beta-1) / 18 C_{p}\right)^{1 /(1-\beta)}$ so that

$$
\begin{aligned}
\mathbb{E}\left[\|x\|_{2}^{2} \mathbb{1}\left(\|x\|_{2}^{2}>\rho^{2} / 4^{2}\right)\right] & =\int_{a_{0}}^{\infty} \mathbb{P}\left(\|x\|_{2}^{2}>t\right) d t \\
& \leq C_{p} \frac{a_{0}^{1-\beta}}{\beta-1} \\
& \leq \varepsilon / 18 .
\end{aligned}
$$

and so (E.11) holds as well under Assumption E.3:. We can therefore bound

$$
\begin{aligned}
\mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \mathbb{1}\left(\|x\|_{2}^{2}>\rho^{2} / 4^{2}\right)\right] & \leq \mathbb{E}\left[\left\|w_{t}-v\right\|_{2}^{2}\|x\|_{2}^{2} \mathbb{1}\left(\|x\|_{2}^{2}>\rho^{2} / 4^{2}\right)\right] \\
& \leq\left\|w_{0}-v\right\|_{2}^{2} \mathbb{E}\left[\|x\|_{2}^{2} \mathbb{1}\left(\|x\|_{2}^{2}>\rho^{2} / 4^{2}\right)\right]
\end{aligned}
$$

$$
\begin{align*}
& \leq\left\|w_{0}-v\right\|_{2}^{2} \varepsilon / 18 \\
& \leq \varepsilon / 2 \tag{E.12}
\end{align*}
$$

The first inequality uses that $\sigma$ is 1-Lipschitz and Cauchy-Schwarz. The second inequality uses (E.6) The third inequality uses (E.11). The final inequality uses that $\left\|w_{0}-v\right\|_{2} \leq\left\|w_{0}\right\|_{2}+\|v\|_{2} \leq 3$.
We can then guarantee

$$
\begin{aligned}
2 \gamma F\left(w_{t}\right)= & \gamma \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2}\right] \\
= & \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \gamma \mathbb{1}\left(\left|w_{t}^{\top} x\right| \leq \rho\right)\right] \\
& +\gamma \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \mathbb{1}\left(\left|w_{t}^{\top} x\right|>\rho\right)\right] \\
\leq & \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x\right) \mathbb{1}\left(\left|w_{t}^{\top} x\right| \leq \rho\right)\right] \\
& +\gamma \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x\right)-\sigma\left(v^{\top} x\right)\right)^{2} \mathbb{1}\left(\|x\|_{2}^{2}>\rho^{2} / 4^{2}\right)\right] \\
\leq & \zeta \varepsilon+\gamma \varepsilon / 2 \\
\leq & \gamma \varepsilon .
\end{aligned}
$$

The first inequality follows since Assumption 3.1 implies $\sigma^{\prime}(z) \mathbb{1}(|z| \leq \rho) \geq \gamma \mathbb{1}(|z| \leq \rho)$ and by (E.9). The second inequality uses (E.8) and (E.12) The final inequality takes $\zeta=\gamma / 2$.

Remark E.5. The precise runtime guarantee in Theorem E.1 will depend upon the activation function and tail distribution. The worst-case activation functions (like the sigmoid) can have $\gamma \sim \exp (-\rho)$, and so if one only has polynomial tails, the runtime can be exponential in $\varepsilon^{-1}$ in this case. If the distribution of $\|x\|_{2}^{2}$ has exponential tails, as is the case if the components of $x$ are sub-Gaussian, runtime will be polynomial in $\varepsilon^{-1}$. On the other hand, if the $\gamma$ in Assumption 3.1 is a fixed constant independent of $\rho$ (as it is for the leaky ReLU), any of the tail bounds under consideration will have runtime of order $\varepsilon^{-1}$.

## E. 2 Stochastic gradient descent proofs

We consider the online version of stochastic gradient descent, where we sample independent samples $x_{t} \sim \mathcal{D}$ at each step and compute stochastic gradient updates $g_{t}$, such that

$$
g_{t}=\left(\sigma\left(w_{t}^{\top} x_{t}\right)-\sigma\left(v^{\top} x_{t}\right)\right) \sigma^{\prime}\left(w_{t}^{\top} x_{t}\right) x_{t}, \quad w_{t+1}=w_{t}-\eta g_{t}
$$

As in the gradient descent case, we have a key lemma that relates the distance of the weights at iteration $t$ from the optimal $v$ with the distance from initialization and the cumulative loss.
Lemma E.6. Assume that $\sigma$ is non-decreasing and $L$-Lipschitz, and that $\mathcal{D}$ satisfies $\|x\|_{2} \leq B$ a.s. Assume the initialization satisfies $\left\|w_{0}\right\|_{2} \leq 2$. Let $T \in \mathbb{N}$ and run stochastic gradient descent for $T-1$ iterations at a fixed learning rate $\eta$ satisfying $\eta \leq L^{-2} B^{-2}$. Then with probability one over $\mathcal{D}$, we have $\left\|w_{t+1}-v\right\|_{2} \leq\left\|w_{t}-v\right\|_{2}$ for all $t<T$, and

$$
\left\|w_{0}-v\right\|_{2}^{2}-\left\|w_{T}-v\right\|_{2}^{2} \geq 2 \eta L^{-1} \sum_{t=0}^{T-1} H_{t}
$$

where $H_{t}:=\frac{1}{2}\left(\sigma\left(w_{t}^{\top} x_{t}\right)-\sigma\left(v^{\top} x_{t}\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x_{t}\right)$.
Proof. We begin with the decomposition

$$
\begin{equation*}
\left\|w_{t}-v\right\|_{2}^{2}-\left\|w_{t+1}-v\right\|_{2}^{2}=2 \eta\left\langle g_{t}, w_{t}-v\right\rangle-\eta^{2}\left\|g_{t}\right\|_{2}^{2} \tag{E.13}
\end{equation*}
$$

By Assumption 3.1, since $\|x\|_{2} \leq B$ a.s. it holds with probability one that

$$
\begin{equation*}
\left\|g_{t}\right\|_{2}^{2}=\left\|\left(\sigma\left(w_{t}^{\top} x_{t}\right)-\sigma\left(v^{\top} x_{t}\right)\right) \sigma^{\prime}\left(w_{t}^{\top} x_{t}\right) x_{t}\right\|_{2}^{2} \leq 2 L B^{2} H_{t} \tag{E.14}
\end{equation*}
$$

By Fact C.2. since $\sigma^{\prime}(z) \geq 0$, we have with probability one,

$$
\left\langle g_{t}, w_{t}-v\right\rangle=\left(\sigma\left(w_{t}^{\top} x_{t}\right)-\sigma\left(v^{\top} x_{t}\right)\right) \sigma^{\prime}\left(w_{t}^{\top} x_{t}\right)\left(w_{t}^{\top} x_{t}-v^{\top} x_{t}\right)
$$

$$
\begin{align*}
& \geq L^{-1}\left(\sigma\left(w_{t}^{\top} x_{t}\right)-\sigma\left(v^{\top} x_{t}\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x_{t}\right) \\
& =2 L^{-1} H_{t} \tag{E.15}
\end{align*}
$$

Putting (E.14) and (E.15) into (E.13), we get

$$
\begin{aligned}
\left\|w_{t}-v\right\|_{2}^{2}-\left\|w_{t+1}-v\right\|_{2}^{2} & \geq 4 \eta L^{-1} H_{t}-2 \eta^{2} L B^{2} H_{t} \\
& \geq 2 \eta L^{-1} H_{t}
\end{aligned}
$$

by taking $\eta \leq L^{-2} B^{-2}$. Telescoping over $t<T$ gives the desired bound.

We now want to translate the bound on the empirical error to that of the true error. For this we use a martingale Bernstein inequality of Beygelzimer et al. [5]. A similar analysis of SGD was used by Ji and Telgarsky [18] for a one-hidden-layer ReLU network.
Lemma E. 7 (Beygelzimer et al. [5], Theorem 1). Let $\left\{Y_{t}\right\}$ be a martingale adapted to the filtration $\mathcal{F}_{t}$, and let $Y_{0}=0$. Let $\left\{D_{t}\right\}$ be the corresponding martingale difference sequence. Define the sequence of conditional variance

$$
V_{t}:=\sum_{k=1}^{t} \mathbb{E}\left[D_{k}^{2} \mid \mathcal{F}_{k-1}\right]
$$

and assume that $D_{t} \leq R$ almost surely. Then for any $\delta \in(0,1)$, with probability greater than $1-\delta$,

$$
Y_{t} \leq R \log (1 / \delta)+(e-2) V_{t} / R
$$

Lemma E.8. Suppose that $\|x\|_{2} \leq B$ a.s., and let $\sigma$ be non-decreasing and $L$-Lipschitz. Assume that the trajectory of SGD satisfies $\left\|w_{t}-v\right\|_{2} \leq\left\|w_{0}-v\right\|_{2}$ for all $t$ a.s. We then have with probability at least $1-\delta$,

$$
\frac{1}{T} \sum_{t=0}^{T-1} H\left(w_{t}\right) \leq \frac{4}{T} \sum_{t=0}^{T-1} H_{t}+\frac{2}{T} B^{2} L^{3}\left\|w_{0}-v\right\|_{2}^{2} \log (1 / \delta)
$$

Proof. Let $\mathcal{F}_{t}=\sigma\left(x_{0}, \ldots, x_{t}\right)$ be the $\sigma$-algebra generated by the first $t+1$ draws from $\mathcal{D}$. Then the random variable $G_{t}:=\sum_{\tau=0}^{t}\left(H\left(w_{\tau}\right)-H_{\tau}\right)$ is a martingale with respect to the filtration $\mathcal{F}_{t}$ with martingale difference sequence $D_{t}:=H\left(w_{t}\right)-H_{t}$. We need bounds on $D_{t}$ and on $\mathbb{E}\left[D_{t}^{2} \mid \mathcal{F}_{t-1}\right]$ in order to apply Lemma E. 7 .
Since $\sigma$ is $L$-Lipschitz and $\|x\|_{2} \leq B$ a.s., with probability one we have

$$
\begin{equation*}
D_{t} \leq H\left(w_{t}\right) \leq \frac{1}{2} L^{3} B^{2}\left\|w_{t}-v\right\|_{2}^{2} \leq \frac{1}{2} L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2} \tag{E.16}
\end{equation*}
$$

The last inequality uses the assumption that $\left\|w_{t}-v\right\|_{2} \leq\left\|w_{0}-v\right\|_{2}$ a.s. Similarly,

$$
\begin{align*}
\mathbb{E}\left[H_{t}^{2} \mid \mathcal{F}_{t-1}\right] & =\frac{1}{4} \mathbb{E}\left[\left(\sigma\left(w_{t}^{\top} x_{t}\right)-\sigma\left(v^{\top} x_{t}\right)\right)^{4} \sigma^{\prime}\left(w_{t}^{\top} x_{t}\right)^{2} \mid \mathcal{F}_{t-1}\right] \\
& \leq \frac{1}{4} L^{3} B^{2}\left\|w_{t}-v\right\|_{2}^{2} \mathbb{E}_{x}\left[\left(\sigma\left(w_{t} x_{t}\right)-\sigma\left(v^{\top} x_{t}\right)\right)^{2} \sigma^{\prime}\left(w_{t}^{\top} x_{t}\right) \mid \mathcal{F}_{t-1}\right] \\
& \leq \frac{1}{2} L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2} H\left(w_{t}\right) \tag{E.17}
\end{align*}
$$

In the first inequality, we have used $\|x\|_{2}^{2} \leq B^{2}$ a.s. and $L$-Lipschitzness of $\sigma$. For the second, we use the assumption that $\left\|w_{t}-v\right\|_{2} \leq\left\|w_{0}-v\right\|_{2}$ together with the fact that $\mathbb{E}_{x}\left[H_{t} \mid \mathcal{F}_{t-1}\right]=H\left(w_{t}\right)$. We then can use (E.17) to bound the squared increments,

$$
\begin{align*}
\mathbb{E}\left[D_{t}^{2} \mid \mathcal{F}_{t-1}\right] & =H\left(w_{t}\right)^{2}-2 H\left(w_{t}\right) \mathbb{E}\left[H_{t} \mid \mathcal{F}_{t-1}\right]+\mathbb{E}\left[H_{t}^{2} \mid \mathcal{F}_{t-1}\right] \\
& =-H\left(w_{t}\right)^{2}+\mathbb{E}\left[H_{t}^{2} \mid \mathcal{F}_{t-1}\right] \\
& \leq \frac{1}{2} L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2} H\left(w_{t}\right) . \tag{E.18}
\end{align*}
$$

This allows for us to bound

$$
V_{T}:=\sum_{t=0}^{T-1} \mathbb{E}\left[D_{t}^{2} \mid \mathcal{F}_{t-1}\right] \leq \frac{1}{2} B^{2} L^{3}\left\|w_{0}-v\right\|_{2}^{2} \sum_{t=0}^{T-1} H\left(w_{t}\right)
$$

Since $D_{t} \leq H\left(w_{t}\right) \leq(1 / 2) L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2}$ a.s. by (E.16). Lemma E.7 implies that with probability at least $1-\delta$, we have

$$
\sum_{t=0}^{T-1}\left(H\left(w_{t}\right)-H_{t}\right) \leq(\exp (1)-2) \sum_{t=0}^{T-1} H\left(w_{t}\right)+\frac{1}{2} L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2} \log (1 / \delta)
$$

and using that $(1-\exp (1)+2)^{-1} \leq 4$, we divide each side by $T$ and get

$$
\begin{equation*}
\frac{1}{T} \sum_{t=0}^{T-1} H\left(w_{t}\right) \leq \frac{4}{T} \sum_{t=0}^{T-1} H_{t}+\frac{2}{T} L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2} \log (1 / \delta) \tag{E.19}
\end{equation*}
$$

With the above in hand, we can prove TheoremE.1]in the SGD setting.
Proof of Theorem E.1. SGD. By the assumptions in the theorem, Lemma E. 6 holds, so that we have for any $t=0, \ldots, T-1,\left\|w_{t}\right\|_{2} \leq 4$ and

$$
\begin{equation*}
\left\|w_{t}-v\right\|_{2}^{2}+2 \eta L^{-1} \sum_{\tau=0}^{t-1} H_{\tau} \leq\left\|w_{0}-v\right\|_{2}^{2} \tag{E.20}
\end{equation*}
$$

This shows that $\left\|w_{t}-v\right\|_{2} \leq\left\|w_{0}-v\right\|_{2}$ holds for all $t=0, \ldots, T-1$ a.s., allowing for the application of Lemma E. 8 to get

$$
\begin{equation*}
\frac{1}{T} \sum_{t=0}^{T-1} H\left(w_{t}\right) \leq \frac{4}{T} \sum_{t=1}^{T} H_{t}+\frac{2}{T} L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2} \log (1 / \delta) \tag{E.21}
\end{equation*}
$$

Dividing both sides of (E.20) by $\eta T L^{-1}$ yields

$$
\min _{t<T} H\left(w_{t}\right) \leq \frac{1}{T} \sum_{t=0}^{T-1} H\left(w_{t}\right) \leq \frac{L\left\|w_{0}-v\right\|_{2}^{2}}{\eta T}+\frac{2}{T} L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2} \log (1 / \delta)
$$

For arbitrary $\zeta>0$, taking $T=\left\lceil 2 \varepsilon^{-1} \zeta^{-1} \eta^{-1} L^{3} B^{2}\left\|w_{0}-v\right\|_{2}^{2} \log (1 / \delta)\right\rceil$ shows there exists $T$ such that $H\left(w_{t}\right) \leq \zeta \varepsilon$. When $\sigma$ satisfies Assumption 3.1, since $\left\|w_{t}\right\|_{2} \leq 4$ for all $t$, it holds that $H\left(w_{t}\right) \geq \gamma F\left(w_{t}\right)$, so that $\zeta=\gamma$ furnishes the desired bound.
When $\sigma$ is ReLU and $\mathcal{D}$ satisfies Assumption 3.2, we note that Lemma E.6 implies $\left\|w_{t}-v\right\|_{2} \leq$ $\left\|w_{0}-v\right\|_{2}$ a.s. Thus taking $\zeta=\alpha^{4} \beta / 8 \sqrt{2}$ and using Lemma 3.5 completes the proof.

## F Remaining Proofs

Proof of Lemma 3.8. Since $\sigma$ is non-decreasing, $\left|\sigma\left(v^{\top} x\right)-y\right| \leq\left|\sigma\left(B_{X}\right)\right|+B_{Y}$. In particular, each summand defining $\widehat{F}(v)$ is a random variable with absolute value at most $a=\left(\left|\sigma\left(B_{X}\right)\right|+B_{Y}\right)^{2}$. As $\mathbb{E}[\widehat{F}(v)]=F(v)=$ OPT, Hoeffding's inequality implies the lemma.

Proof of Lemma 3.9. The bound $\mathfrak{R}_{S}(\mathcal{G}) \leq 2 \max _{i}\left\|x_{i}\right\|_{2} / \sqrt{n}$ follows since $\|w\|_{2} \leq 2$ holds on $\mathcal{G}$ with standard results Rademacher complexity theory (e.g. Sec. 26.2 of [28]); this shows $\mathfrak{R}(\mathcal{G}) \leq$ $2 B_{X} / \sqrt{n}$. Using the contraction property of the Rademacher complexity, this implies $\mathfrak{R}(\sigma \circ \mathcal{G}) \leq$ $2 B_{X} L / \sqrt{n}$. Finally, note that if $\|w-v\|_{2} \leq 1$ and $\|x\|_{2} \leq B_{X}$, we have

$$
\begin{equation*}
\|\nabla \ell(w ; x)\|=\left\|\left(\sigma\left(w^{\top} x\right)-\sigma\left(v^{\top} x\right)\right) \sigma^{\prime}\left(w^{\top} x\right) x\right\| \leq L^{2}\|w-v\|\|x\| \leq L^{2} B_{X} \tag{F.1}
\end{equation*}
$$

In particular, $\ell$ is $L^{2} B_{X}$ Lipschitz. The result follows.


[^0]:    ${ }^{3}$ Agnostic learning results typically require i.i.d. samples, and adversarial noise may depend on other samples in malicious ways. Even in the i.i.d. case, trouble arises if one wishes to use parameter recovery to show that a given algorithm competes with the population risk minimizer. Consider the ReLU with labels given by $y=\sigma\left(v^{\top} x\right)+\xi$ where $\xi=-\sigma\left(v^{\top} x\right)$. The zero vector minimizes the population risk, and so any algorithm that returns the target neuron $\sigma\left(v^{\top} x\right)$ has large population risk. A similar phenomenon occurs for $\xi=\sigma\left(v^{\top} x\right)$.
    ${ }^{4}$ Although their result is stated for the ReLU and isotropic log-concave distributions, their results also apply for $L$-Lipschitz activations satisfying $\inf _{z} \sigma^{\prime}(z) \geq \gamma>0$ for isotropic distributions that satisfy our Assumption 3.2 In this setting, one can show that the Chow parameters satisfy $\left\|\chi\left(\sigma_{u}\right)-\chi\left(\sigma_{w}\right)\right\|^{2} \geq \gamma L^{-1} \mathbb{E}\left[\left(\sigma\left(u^{\top} x\right)-\right.\right.$ $\left.\left.\sigma\left(v^{\top} x\right)\right)^{2}\right]$, from which the result follows easily.

