## Appendix of the paper: Optimal Weak to Strong Learning

## A Weak to strong learning

The following theorem is essentially a restatement of Theorem 1 from the main paper.
Theorem A.1. Assume we are given access to a $\gamma$-weak learner for a $0<\gamma<1 / 2$, using base hypothesis set $\mathcal{H} \subseteq \mathcal{X} \rightarrow\{-1,1\}$ of VC-dimension d. Then there is a universal constant $\alpha>0$ and an algorithm $\mathcal{A}$, such that for every $0<\delta<1$ and every distribution $\mathcal{D}$ over $\mathcal{X} \times\{-1,1\}$, it holds with probability at least $1-\delta$ over a set of $m$ samples $S \sim \mathcal{D}^{m}$, that $\mathcal{A}$ on $S$ outputs a classifier $h_{S}=\mathcal{A}(S) \in \mathcal{X} \rightarrow\{-1,1\}$ with

$$
\mathcal{L}_{\mathcal{D}}\left(h_{S}\right) \leq \alpha \cdot \frac{d \gamma^{-2}+\ln (1 / \delta)}{m}
$$

Theorem 1 of the main paper follows by setting $\varepsilon=\mathcal{L}_{\mathcal{D}}\left(h_{S}\right)$ and solving for $m$ and letting the label in the distribution $\mathcal{D}$ be $c(x)$ for every $x \in \mathcal{X}$. The algorithm that obtains the guarantees has been described in the main paper. We thus only present (again) the two algorithms Algorithm 1 and Algorithm 2], as well as AdaBoost ${ }_{V}^{*}$ Algorithm A.1] by Rätsch et al. [20] that achieves almost optimal margins and is used in Algorithm 2
In the remainder of the section, we prove that Algorithm 2 has the guarantees of Theorem A. 1 The proof follows that of Hanneke [13] pretty much uneventfully, although carefully using that a generalization error of $1 / 200$ suffices. For simplicity, we assume $m$ is a power of 4 . This can easily be ensured by rounding $m$ down to the nearest power of 4 and ignoring all excess samples. This only affects the generalization bound by a constant factor. With $m$ being a power of 4 we can observe from Algorithm 1 that the cardinalities of all recursively generated sets $A_{0}$ (which are the input to the next level of the recursion) are also powers of 4 . Hence we can ignore all roundings.

## A. 1 Proof of Optimal Strong Learning

Let $C \subseteq \mathcal{X} \rightarrow\{-1,1\}$ be a concept class and assume there is a $\gamma$-weak learner for $C$ using hypothesis set $\mathcal{H}$ of VC-dimension $d$. Let $\mathcal{A}_{v}^{*}$, be an algorithm that on a sample $S$ consistent with a concept $c \in C$, computes a voting classifier $f \in \Delta(\mathcal{H})$ with $y f(x) \geq \gamma / 2$ for all $(x, y) \in S$ and returns as its output hypothesis $g(x)=\operatorname{sign}(f(x))$. We could e.g. let $\mathcal{A}_{v}^{*}$ be AdaBoost ${ }_{v}^{*}$. For a sample $S$, we use the notation $\mathcal{M}_{\gamma}(S)$ to denote the set of hypotheses $g(x)=\operatorname{sign}(f(x))$ for an $f \in \Delta(\mathcal{H})$ satisfying $y f(x) \geq \gamma$ for all $(x, y) \in S$. The set $\mathcal{M}_{\gamma}(S)$ is thus the set of all voting classifiers obtained by taking the sign of a voter that has margins at least $\gamma$ on all samples in $S$. By definition, the output hypothesis $g$ of $\mathcal{A}_{v}^{*}$ on a set of samples $S$ always lies in $\mathcal{M}_{\gamma / 2}(S)$.
Let $c \in \mathcal{C}$ be an unknown concept in $C$ and let $\mathcal{D}$ be an arbitrary distribution over $\mathcal{X}$. Let $S=\left\{\left(x_{i}, c\left(x_{i}\right)\right)\right\}_{i=1}^{m} \in(\mathcal{X} \times\{-1,1\})^{m}$ be a set of $m$ samples with each $x_{i}$ an i.i.d. sample from $\mathcal{D}$. Let $S_{1: k}$ denote the first $k$ samples of $S$. Let $c^{\prime} \geq 4$ be a constant to be determined later. We will prove by induction that for every $m^{\prime} \in \mathbb{N}$ that is a power of 4 , for every $\delta^{\prime} \in(0,1)$, and every finite sequence $B^{\prime}$ of samples in $\mathcal{X} \times\{-1,1\}$ with $y_{i}=c\left(x_{i}\right)$ for each $\left(x_{i}, y_{i}\right) \in B^{\prime}$, with probability at least $1-\delta^{\prime}$, the classifier

$$
\hat{h}_{m^{\prime}, B^{\prime}}=\operatorname{sign}\left(\sum_{C_{i} \in \operatorname{Sub-Sample}\left(S_{1: m^{\prime}}, B^{\prime}\right)} \mathcal{A}_{v}^{*}\left(C_{i}\right)\right)
$$

satisfies

$$
\begin{equation*}
\mathcal{L}_{\mathcal{D}}\left(\hat{h}_{m^{\prime}, B^{\prime}}\right) \leq \frac{c^{\prime}}{m^{\prime}}\left(d \gamma^{-2}+\ln \left(1 / \delta^{\prime}\right)\right) . \tag{3}
\end{equation*}
$$

The conclusion of Theorem A.1 follows by letting $B^{\prime}=\emptyset$ and $m^{\prime}=m$ (and recalling that we assume $m$ is a power of 4). Thus what remains is to give the inductive proof.

As the base case, consider any $m^{\prime} \in \mathbb{N}$ with $m^{\prime} \leq c^{\prime}$ and $m^{\prime}$ a power of 4 . In this case, the bound $c^{\prime}\left(d \gamma^{-2}+\ln \left(1 / \delta^{\prime}\right)\right) / m^{\prime}$ is at least $d \gamma^{-2} \geq 1$ and $\mathcal{L}_{\mathcal{D}}\left(\hat{h}_{m^{\prime}, B^{\prime}}\right) \leq 1$ obviously holds.

For the inductive step, take as inductive hypothesis that, for some $m \in \mathbb{N}$ with $m>c^{\prime}$ and $m$ a power of 4 , it holds for all $m^{\prime} \in \mathbb{N}$ with $m^{\prime}<m$ and $m^{\prime}$ a power of 4 , that for every $\delta^{\prime} \in(0,1)$ and every

```
Algorithm A.1: AdaBoost* \({ }_{v}[20]\)
Input: training set \(S=\left\{\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)\right\}\)
    number of rounds \(T\)
    desired accuracy \(v\)
Result: An ensemble hypothesis \(H_{\text {out }}\) with almost optimal margins
\(\mathcal{D}^{(1)} \leftarrow\left(\frac{1}{m}, \ldots \frac{1}{m}\right) \quad\) // uniform initialization of \(\mathcal{D}\)
for \(t \in\{1, \ldots, T\}\) do
        \(h_{t} \leftarrow \mathrm{WL}\left(\mathcal{D}^{(t)}, S\right) \quad / /\) invoke weak learner
        \(\gamma_{t} \leftarrow \sum_{i=1}^{m} \mathcal{D}_{i}^{(t)} y_{i} h_{t}\left(x_{i}\right) \quad\) // average margin of \(h_{t}\)
        if \(\left|\gamma_{t}\right|=1\) then
            /* \(h_{t}\) is consistent \(\Rightarrow\) taking only \(h_{t}\) as 'ensemble' maximizes the margin */
            \(w_{1} \leftarrow \operatorname{sign}\left(\gamma_{t}\right), \quad h_{1} \leftarrow h_{t}, \quad T \leftarrow 1\)
            break
        \(\gamma_{t}^{\min } \leftarrow \min _{r \in[t]} \gamma_{r}, \quad\) // update assumed advantage
        \(\rho_{t}^{\min } \leftarrow \gamma_{t}^{\min }-v\)
        \(w_{t}=\frac{1}{2} \ln \frac{1+\gamma_{t}}{1-\gamma_{t}}-\frac{1}{2} \ln \frac{1+\rho_{t}}{1-\rho_{t}} \quad / /\) weight for the current hypothesis
        for \(i \in\{1, \ldots, m\}\) do
            \(\mathcal{D}_{i}^{(t+1)} \leftarrow \frac{\mathcal{D}_{i}^{(t)} \exp \left(-w_{t} y_{i} h_{t}\left(x_{i}\right)\right)}{\sum_{j=1}^{m} \mathcal{D}_{i}^{(t)} \exp \left(-w_{t} y_{j} h_{t}\left(x_{j}\right)\right)} \quad\) // update \(\mathcal{D}\)
return \(f_{\text {out }}(x)=\frac{1}{\sum_{i=1}^{T}\left|w_{i}\right|} \sum_{t=1}^{T} w_{t} h_{t}(x) \quad / /\) (normalized) weighted majority vote
```

finite sequence $B^{\prime}$ of samples in $\mathcal{X} \times\{-1,1\}$ with $y_{i}=c\left(x_{i}\right)$ for each $\left(x_{i}, y_{i}\right) \in B^{\prime}$, with probability at least $1-\delta^{\prime}$, Eq. (3) holds. We need to prove that the inductive hypothesis also holds for $m^{\prime}=m$.
Fix a $\delta \in(0,1)$ and any finite sequence $B$ of points in $\mathcal{X} \times\{-1,1\}$ with $y_{i}=c\left(x_{i}\right)$ for each $\left(x_{i}, y_{i}\right)$ in $B$. Since $m>c^{\prime} \geq 4$, we have that $\operatorname{Sub}-\operatorname{Sample}\left(S_{1: m}, B\right)$ returns in Step 5 of Algorithm 1. Let $A_{0}, A_{1}, A_{2}, A_{3}$ be as defined in Step 4 of Algorithm 1. Also define $B_{1}=A_{2} \cup A_{3} \cup B$, $B_{2}=A_{1} \cup A_{3} \cup B, B_{3}=A_{1} \cup A_{2} \cup B$, and for each $i \in\{1,2,3\}$, denote

$$
h_{i}=\operatorname{sign}\left(\sum_{C_{i} \in \operatorname{Sub-Sample}\left(A_{0}, B_{i}\right)} \mathcal{A}_{v}^{*}\left(C_{i}\right)\right) .
$$

Note that the $h_{i}$ 's correspond to the majority vote classifiers trained on the sub-samples of the three recursive calls in Algorithm 1. Moreover, notice that $h_{i}=\hat{h}_{m / 4, B_{i}}$. Therefore, the inductive hypothesis may be used on $h_{1}, h_{2}, h_{3}$ to conclude that for each $i \in\{1,2,3\}$, there is an event $E_{i}$ of probability at least $1-\delta / 9$, on which

$$
\begin{equation*}
\mathcal{L}_{\mathcal{D}}\left(h_{i}\right) \leq \frac{c^{\prime}}{\left|A_{0}\right|}\left(d \gamma^{-2}+\ln (9 / \delta)\right) \leq \frac{4 c^{\prime}}{m}\left(d \gamma^{-2}+\ln (1 / \delta)+3\right) \leq \frac{12 c^{\prime}}{m}\left(d \gamma^{-2}+\ln (1 / \delta)\right) . \tag{4}
\end{equation*}
$$

Here we chose the probability $1-\delta / 9$ in order to perform a union bound in the end of the induction step which is possible since the inductive hypothesis holds for every $\delta^{\prime}$. Next, define $\operatorname{Err}\left(h_{i}\right)$ as the set of points $x \in \mathcal{X}$ for which $h_{i}(x) \neq c(x)$. Now fix an $i \in\{1,2,3\}$ and denote by $\left\{\left(Z_{i, 1}, c\left(Z_{i, 1}\right)\right), \ldots,\left(Z_{i, N_{i}}, c\left(Z_{i, N_{i}}\right)\right\}=A_{i} \cap\left(\operatorname{Err}\left(h_{i}\right) \times\{-1,1\}\right)\right.$, where $N_{i}=\mid A_{i} \cap\left(\operatorname{Err}\left(h_{i}\right) \times\right.$ $\{-1,1\}) \mid$. Said in words, the set $\left\{\left(Z_{i, j}, c\left(Z_{i, j}\right)\right\}_{j=1}^{N_{i}}\right.$ is the subset of samples in $A_{i}$ on which $h_{i}$ makes a mistake. Notice that $h_{i}$ is not trained on any samples from $A_{i}\left(B_{i}\right.$ excludes $\left.A_{i}\right)$, hence $h_{i}$ and $A_{i}$ are independent. Therefore, given $h_{i}$ and $N_{i}$, the samples $Z_{i, 1}, \ldots, Z_{i, N_{i}}$ are conditionally independent samples with distribution $\mathcal{D}\left(\cdot \mid \operatorname{Err}\left(h_{i}\right)\right.$ ) (provided $\left.N_{i}>0\right)$. From Theorem 6 in the main paper, we get that there is an event $E_{i}^{\prime}$ of probability at least $1-\delta / 9$, such that if $N_{i} \geq c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right)$, then every $h \in \mathcal{M}_{\gamma / 2}\left(\left\{\left(Z_{i, j}, c\left(Z_{i, j}\right)\right)\right\}_{j=1}^{N_{i}}\right)$ satisfies

$$
\mathcal{L}_{\mathcal{D}\left(\cdot \mid \operatorname{Err}\left(h_{i}\right)\right)}(h) \leq \frac{1}{200} .
$$

Note that this is a key step where our proof differs from Hanneke's original proof since we exploit that a bound of $\frac{1}{200}$ on the generalization error suffices for the rest of the proof. We continue by observing
that for each $j \in\{1,2,3\} \backslash\{i\}$, the set $B_{j}$ contains $A_{i}$ and this remains the case in all recursive calls of $\operatorname{Sub}$-Sample $\left(A_{0}, B_{i}\right)$. Thus for $\left\{C_{1}, \ldots, C_{k}\right\}=\operatorname{Sub}-\operatorname{Sample}\left(A_{0}, B_{j}\right)$, it holds for all $C_{k}$ that $\mathcal{A}_{v}^{*}\left(C_{k}\right) \in \mathcal{M}_{\gamma / 2}\left(B_{j}\right) \Rightarrow \mathcal{A}_{v}^{*}\left(C_{k}\right) \in \mathcal{M}_{\gamma / 2}\left(A_{0}\right) \Rightarrow \mathcal{A}_{v}^{*}\left(C_{k}\right) \in \mathcal{M}_{\gamma / 2}\left(\left\{\left(Z_{i, j}, c\left(Z_{i, j}\right)\right)\right\}_{j=1}^{N_{i}}\right)$. Thus on the event $E_{i}^{\prime}$ defined above, if $N_{i}>c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right)$, then it holds for all $j \in\{1,2,3\} \backslash\{i\}$ and all $C_{k} \in \operatorname{Sub}-\operatorname{Sample}\left(A_{0}, B_{j}\right)$, that the hypothesis $h=\mathcal{A}_{v}^{*}\left(C_{k}\right)$ satisfies

$$
\begin{aligned}
\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{i}(x) \neq c(x) \wedge h(x) \neq c(x)\right] & =\mathcal{L}_{\mathcal{D}}\left(h_{i}\right) \cdot \mathcal{L}_{\mathcal{D}\left(\cdot \mid \operatorname{Err}\left(h_{i}\right)\right)}(h) \\
& \leq \frac{1}{200} \mathcal{L}_{\mathcal{D}}\left(h_{i}\right) .
\end{aligned}
$$

Assume now that $\mathcal{L}_{\mathcal{D}}\left(h_{i}\right) \geq\left((10 / 7) c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right)+23 \ln (9 / \delta)\right) /(m / 4) \geq 23 \ln (9 / \delta) /\left|A_{i}\right|$.
Using that $h_{i}$ and $A_{i}$ are independent, it follows by a Chernoff bound that

$$
\begin{aligned}
\operatorname{Pr}\left[N_{i} \geq(7 / 10) \mathcal{L}_{\mathcal{D}}\left(h_{i}\right)\left|A_{i}\right|\right] & \geq 1-\exp \left(-(3 / 10)^{2} \mathcal{L}_{\mathcal{D}}\left(h_{i}\right)\left|A_{i}\right| / 2\right) \\
& \geq 1-\exp \left(-(3 / 10)^{2} \cdot 23 \ln (9 / \delta) / 2\right) \\
& >1-\delta / 9
\end{aligned}
$$

Thus there is an event $E_{i}^{\prime \prime}$ of probability at least $1-\delta / 9$, on which, if

$$
\mathcal{L}_{\mathcal{D}}\left(h_{i}\right) \geq \frac{(10 / 7) c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right)+23 \ln (9 / \delta)}{m / 4}
$$

then

$$
\begin{aligned}
N_{i} & \geq(7 / 10) \mathcal{L}_{\mathcal{D}}\left(h_{i}\right)\left|A_{i}\right| \\
& =(7 / 10) \mathcal{L}_{\mathcal{D}}\left(h_{i}\right) m / 4 \\
& \geq c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right) .
\end{aligned}
$$

Combining it all, we have that on the event $E_{i} \cap E_{i}^{\prime} \cap E_{i}^{\prime \prime}$, which occurs with probability at least $1-\delta / 3$, if $\mathcal{L}_{\mathcal{D}}\left(h_{i}\right) \geq\left((10 / 7) c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right)+23 \ln (9 / \delta)\right) /(m / 4)$, then every $h=\mathcal{A}_{v}^{*}\left(C_{k}\right)$ for a $C_{k} \in \operatorname{Sub-Sample}\left(A_{0}, B_{j}\right)$ with $j \neq i$ has:

$$
\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{i}(x) \neq c(x) \wedge h(x) \neq c(x)\right] \leq \frac{1}{200} \mathcal{L}_{\mathcal{D}}\left(h_{i}\right)
$$

By Eq. (4), this is at most

$$
\begin{aligned}
\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{i}(x) \neq c(x) \wedge h(x) \neq c(x)\right] & \leq \frac{1}{200} \cdot \frac{12 c^{\prime}}{m}\left(d \gamma^{-2}+\ln (1 / \delta)\right) \\
& \leq \frac{c^{\prime}}{16 m}\left(d \gamma^{-2}+\ln (1 / \delta)\right)
\end{aligned}
$$

On the other hand, if $\mathcal{L}_{\mathcal{D}}\left(h_{i}\right)<\left(c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right)+23 \ln (9 / \delta)\right) /(m / 4)$, then

$$
\begin{aligned}
\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{i}(x) \neq c(x) \wedge h(x) \neq c(x)\right] & \leq \mathcal{L}_{\mathcal{D}}\left(h_{i}\right) \\
& \leq\left(c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right)+23 \ln (9 / \delta)\right) /(m / 4) \\
& \leq 4 c^{\prime \prime}\left(d \gamma^{-2}+24 \ln (1 / \delta)+23 \ln 9\right) / m
\end{aligned}
$$

Using that $23 \cdot \ln 9<51 \leq 51 d \gamma^{-2}$, the above is at most $204 c^{\prime \prime}\left(d \gamma^{-2}+\ln (1 / \delta)\right) / m$. Fixing the constant $c^{\prime}$ to $c^{\prime} \geq(16 \cdot 204) c^{\prime \prime}$, this is at most

$$
\frac{c^{\prime}}{16 m}\left(d \gamma^{-2}+\ln (1 / \delta)\right)
$$

We conclude that on the event $\bigcap_{i=1,2,3}\left\{E_{i} \cap E_{i}^{\prime} \cap E_{i}^{\prime \prime}\right\}$, which occurs with probability at least $1-\delta$ by a union bound, it holds for all $i$ and all $C_{k} \in \operatorname{Sub}-\operatorname{Sample}\left(A_{0}, B_{j}\right)$ with $j \neq i$ that the hypothesis $h=\mathcal{A}_{v}^{*}\left(C_{k}\right)$ satisfies:

$$
\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{i}(x) \neq c(x) \wedge h(x) \neq c(x)\right] \leq \frac{c^{\prime}}{16 m}\left(d \gamma^{-2}+\ln (1 / \delta)\right) .
$$

Now consider an $x$ on which $\hat{h}_{m, B}$ errs. On such an $x$, the majority among the classifiers

errs. For the majority to err, there must be an $i \in\{1,2,3\}$ for which the majority of

$$
\bigcup_{C_{k} \in \operatorname{Sub-Sample}\left(S_{1: m / 4}, B_{i}\right)}\left\{\mathcal{A}_{v}^{*}\left(C_{k}\right)\right\}
$$

errs. This is equivalent to $h_{i}(x) \neq c(x)$. Furthermore, even when all of the classifiers in

$$
\bigcup_{C_{k} \in \operatorname{Sub-Sample}\left(S_{1: m / 4}, B_{i}\right)}\left\{\mathcal{A}_{v}^{*}\left(C_{k}\right)\right\}
$$

err, there still must be another (1/6)-fraction of all the classifiers

that err. This follows since each of the three recursive calls in Sub-Sample generated equally many classifiers/samples. It follows that if we pick a uniform random $i \in\{1,2,3\}$ and a uniform random hypothesis $h$ in

then with probability at least $(1 / 3)(1 / 6)(3 / 2)=1 / 12$, we have that $h_{i}(x) \neq c(x) \wedge h(x) \neq c(x)$. It follows by linearity of expectation that on the event $\bigcap_{i=1,2,3}\left\{E_{i} \cap E_{i}^{\prime} \cap E_{i}^{\prime \prime}\right\}$, we have:

$$
\mathcal{L}_{\mathcal{D}}\left(\hat{h}_{m, B}\right) \leq 12 \cdot \frac{c^{\prime}}{16 m}\left(d \gamma^{-2}+\ln (1 / \delta)\right)<\frac{c^{\prime}}{m}\left(d \gamma^{-2}+\ln (1 / \delta)\right) .
$$

This completes the inductive proof and shows Theorem A. 1 .

## B Lower bound

In this section, we prove the following lower bound which directly implies Theorem 2 from the main paper:

Theorem B.1. There is a universal constant $\alpha>0$ such that for all integers $d \in \mathbb{N}$ and every $2^{-d}<\gamma<1 / 80$, there is a finite set $\mathcal{X}$, a concept class $C \subset \mathcal{X} \rightarrow\{-1,1\}$ and a hypothesis set $\mathcal{H} \subseteq \mathcal{X} \rightarrow\{-1,1\}$ of $V C$-dimension at most $d$, such that for every integer $m \in \mathbb{N}$ and $0<\delta<1 / 3$, there is a distribution $\mathcal{D}$ over $\mathcal{X}$ such that the following holds:

1. For every $c \in \mathcal{C}$ and every distribution $\mathcal{D}^{\prime}$ over $\mathcal{X}$, there is an $h \in \mathcal{H}$ with

$$
\operatorname{Pr}_{x \sim \mathcal{D}^{\prime}}[h(x) \neq c(x)] \leq 1 / 2-\gamma .
$$

2. For any algorithm $\mathcal{A}$, there is a concept $c \in C$ such that with probability at least $\delta$ over a set of $m$ samples $S \sim \mathcal{D}^{m}$, the classifier $\mathcal{A}(S) \in \mathcal{X} \rightarrow\{-1,1\}$ produced by $\mathcal{A}$ on $S$ and $c(S)$ must have

$$
\mathcal{L}_{\mathcal{D}}(\mathcal{A}(S)) \geq \alpha \cdot \frac{d \gamma^{-2}+\ln (1 / \delta)}{m}
$$

Theorem B. 1 immediately implies Theorem 2 by solving the equation in the second statement for $\varepsilon=\mathcal{L}_{\mathcal{D}}(\mathcal{A}(S))$.
The proof of the term $\ln (1 / \delta) / m$ in the lower bound follows from previous work. In particular, we could let $C=\mathcal{H}$ and invoke the tight lower bounds for PAC-learning in the realizable setting [5]. Thus, we focus on $\delta=1 / 3$ and only need to prove that the loss of $\mathcal{A}(S)$ is at least $\alpha d \gamma^{-2} / m$ with probability $1 / 3$ over $S$.
For the proof, we make use of the following lemma by Grønlund et al. [9] to construct the 'hard' hypothesis set $\mathcal{H}$ and concept class $C$ :

Lemma B. 1 (Grønlund et al. [9]). For every $\gamma \in(0,1 / 40), \delta \in(0,1)$ and integers $k \leq u$, there exists a distribution $\mu=\mu(u, d, \gamma, \delta)$ over a hypothesis set $\mathcal{H} \subset \mathcal{X} \rightarrow\{-1,1\}$, where $\mathcal{X}$ is a set of size $u$, such that the following holds.

1. For all $\mathcal{H} \in \operatorname{supp}(\mu)$, we have $|\mathcal{H}|=N$; and
2. For every labeling $\ell \in\{-1,1\}^{u}$, if no more than $k$ points $x \in \mathcal{X}$ satisfy $\ell(x)=-1$, then

$$
\operatorname{Pr}_{\mathcal{H} \sim \mu}[\exists f \in \Delta(\mathcal{H}): \forall x \in \mathcal{X}: \ell(x) f(x) \geq \gamma] \geq 1-\delta
$$

$$
\text { where } N=\Theta\left(\gamma^{-2} \ln u \ln \left(\gamma^{-2} \ln u \delta^{-1}\right) e^{\Theta\left(\gamma^{2} k\right)}\right)
$$

To prove Theorem B. 1 for a given $\gamma \in\left(2^{-d}, 1 / 80\right)$ and $m, d \in \mathbb{N}$, let $u=k$ for a $u$ to be determined. Invoke Lemma B.1 with $\delta=1 / 2$ and $\gamma^{\prime}=2 \gamma$ to conclude that there exists a hypothesis set $\mathcal{H}$ such that among all labelings $\ell \in\{-1,1\}^{u}$, at least half of them satisfy:

$$
\exists f \in \Delta(\mathcal{H}): \forall x \in \mathcal{X}: \ell(x) f(x) \geq 2 \gamma
$$

Moreover, we have $N=|\mathcal{H}|=\Theta\left(\gamma^{-2} \ln u \ln \left(\gamma^{-2} \ln u\right) e^{\Theta\left(\gamma^{2} u\right)}\right)$. Let the concept class $C$ be the set of such labelings.
For the given VC-dimension $d$, we need to bound the VC-dimension of $\mathcal{H}$ by $d$. For this, note that the VC-dimension is bounded by $\left.\lg |\mathcal{H}|=\Theta\left(\gamma^{2} u\right)+\lg \left(\gamma^{-2} \lg u\right)\right)$. Using that $\gamma \geq 2^{-d}$, this is at most $\Theta\left(\gamma^{2} u+d+\lg \lg u\right)$. We thus choose $u=\Theta\left(\gamma^{-2} d\right)$ which implies the claimed VC-dimension of $\mathcal{H}$.
Next, we have to argue that any concept $c \in C$ can be $\gamma$-weakly learned from $\mathcal{H}$. That is, the first statement of Theorem B. 1 holds for $\mathcal{H}, C$. To see this, we must show that for every distribution $\mathcal{D}$ over $\mathcal{X}$, there is a hypothesis $h \in \mathcal{H}$ such that $\operatorname{Pr}_{x \sim \mathcal{D}}[h(x)=c(x)] \geq 1 / 2+\gamma$. To argue that this is indeed the case, let $f \in \Delta(\mathcal{H})$ satisfy $\forall x \in \mathcal{X}: c(x) f(x) \geq 2 \gamma$. Such an $f$ exists by definition of $C$. Then, $\mathbb{E}_{x \sim \mathcal{D}}[c(x) f(x)] \geq 2 \gamma$. Since $f(x)$ is a convex combination of hypotheses from $\mathcal{H}$, it follows that there is a hypothesis $h \in \mathcal{H}$ also satisfying $\mathbb{E}_{x \sim \mathcal{D}}[c(x) h(x)] \geq 2 \gamma$. But

$$
\begin{aligned}
\underset{x \sim \mathcal{D}}{\mathbb{E}}[c(x) h(x)] & =\sum_{x \in \mathcal{X}} \mathcal{D}(x) c(x) h(x) \\
& =\sum_{x \in \mathcal{X}:} \mathcal{D}(x)=h(x) \\
& =\operatorname{Pr}_{x \sim \mathcal{D}}[c(x)=h(x)]-\sum_{x \in \mathcal{X}} \operatorname{Pr}[c(x) \neq h(x) \\
& =\operatorname{Pr}_{x \sim \mathcal{D}}[c(x)=h(x)]-\left(1-\operatorname{Pr}_{x \sim \mathcal{D}}[c(x)=h(x)]\right) \\
& =2 \operatorname{Pr}_{x \sim \mathcal{D}}[c(x)=h(x)]-1 .
\end{aligned}
$$

Hence, $2 \cdot \operatorname{Pr}_{x \sim \mathcal{D}}[c(x)=h(x)]-1 \geq 2 \gamma \Longrightarrow \operatorname{Pr}_{x \sim \mathcal{D}}[c(x)=h(x)] \geq 1 / 2+\gamma$ as claimed.
We have thus constructed $\mathcal{H}$ and $C$ satisfying the first statement of Theorem B.1, where $C$ contains at least half of all possible labelings of the points $\mathcal{X}=\left\{x_{1}, \ldots, x_{u}\right\}$ with $u=\Theta\left(\gamma^{-2} d\right)$. For the remainder of the proof, we assume $u$ is at least some large constant, which is true for $\gamma$ small enough.
What remains is to establish the second statement of Theorem B.1. For this, we first define the hard distribution $\mathcal{D}$ over $\mathcal{X}$. The distribution $\mathcal{D}$ returns the point $x_{1}$ with probability $1-(u-1) / 4 m$ and with the remaining probability $(u-1) / 4 m$ it returns a uniform random sample $x_{i}$ among $x_{2}, \ldots, x_{u}$. Also, let $c$ be a uniform random concept drawn from $C$.
Let $\mathcal{A}$ be any (possibly randomized) learning algorithm that on a set of samples $S$ from $\mathcal{X}$ and a labeling $\ell(S)$ of $S$ that is consistent with at least one concept $c \in C$ (i.e. $\ell(S)=c(S)$ ), outputs a hypothesis $h_{S, \ell(S)}$ in $\mathcal{X} \rightarrow\{-1,1\}$. The algorithm $\mathcal{A}$ is not constrained to output a hypothesis from $\Delta(\mathcal{H})$ or $\mathcal{H}$, but instead may output any desirable hypothesis in $\mathcal{X} \rightarrow\{-1,1\}$, using the full knowledge of $C, \ell(S), \mathcal{H}$ and the promise that $c \in C$. Our goal is to show that

$$
\begin{equation*}
\underset{c \sim C}{\mathbb{E}}\left[\operatorname{Pr}_{S \sim \mathcal{D}^{m}}\left[\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)}(x) \neq c(x)\right] \geq \alpha^{\prime} \frac{d \gamma^{-2}}{m}\right]\right] \geq 1 / 3 \tag{5}
\end{equation*}
$$

where $c \sim C$ denotes the uniform random choice of $c$. Notice that if this is the case, there must exist a concept $c$ for which

$$
\operatorname{Pr}_{S \sim \mathcal{D}^{m}}\left[\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)}(x) \neq c(x)\right] \geq \alpha^{\prime} \frac{d \gamma^{-2}}{m}\right] \geq 1 / 3 .
$$

To establish Eq. (5), we start by observing that for any randomized algorithm $\mathcal{A}$, there is a deterministic algorithm $\mathcal{A}^{\prime}$ obtaining a smaller than or equal value of the left hand side of Eq. (5) (by Yao's principle). Thus, we assume from here on that $\mathcal{A}$ is deterministic.

The main idea in our proof is to first show that conditioned on the set $S$ and label $c(S)$, the concept $c$ is still largely unknown. We formally measure this by arguing that the binary Shannon entropy of $c$ is large conditioned on $S$ and $c(S)$. Next, we argue that if a learning algorithm often manages to produce an accurate hypothesis from $S$ and $c(S)$, then that reveals a lot of information about $c$, i.e. the entropy of $c$ is small conditioned on $S$ and $c(S)$. This contradicts the first statement and thus the algorithm cannot produce an accurate hypothesis. We now proceed with the two steps.

Large Conditional Entropy. Consider the binary Shannon entropy of the uniform random $c$ conditioned on $S$ and $c(S)$, denoted $H(c \mid S, c(S))$. We know that $H(c)=\lg |C| \geq \lg \left(2^{u} / 2\right)=$ $u-1$. The random variable $c$ is independent of $S$, hence $H(c \mid S)=H(c)$. We therefore have $H(c \mid S, c(S)) \geq H(c \mid S)-H(c(S) \mid S)=u-1-H(c(S) \mid S)$. For a fixed $s \in X^{m}$, let $p_{s}=\operatorname{Pr}_{S \sim \mathcal{D}^{m}}[S=s]$. Then $H(c(S) \mid S)=\sum_{s \in X^{m}} p_{s} H(c(S) \mid S=s) \leq \sum_{s \in X^{m}} p_{s}|s|$, where the last step follows from the fact that, conditioned on $s$, the labeling $c(s)$ consists of $|s|$ signs. Note that the size of the set $|s|$ is possibly smaller than $m$ due to repetitions.
Now notice that $\operatorname{Pr}[|S|>u / 3]$ is exponentially small in $u$ since each of the $m$ samples from $\mathcal{D}$ is among $x_{2}, \ldots, x_{u}$ with probability only $(u-1) /(4 m)$. Therefore, we get $H(c(S) \mid S) \leq$ $u / 3+\exp (-\Omega(u)) u \leq u / 2-1$. It follows that

$$
\begin{equation*}
H(c \mid S, c(S)) \geq u-1-(u / 2-1)=u / 2 \tag{6}
\end{equation*}
$$

Accuracy Implies Low Entropy. Now assume that $h_{S, c(S)}$ is such that $\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)} \neq c(x)\right]<$ $\alpha^{\prime} d \gamma^{-2} / m$ for a sufficiently small constant $\alpha^{\prime}$. Any point $x_{i}$ where $c\left(x_{i}\right)$ disagrees with $h_{S, c(S)}\left(x_{i}\right)$ adds at least $1 /(4 m)$ to $\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)} \neq c(x)\right]$ (the point $x_{1}$ would add more), hence $h_{S, c(S)}$ makes a mistake on at most $\alpha^{\prime} d \gamma^{-2} / m \cdot(4 m)=4 \alpha^{\prime} d \gamma^{-2}$ points. Recalling that $u=\Theta\left(d \gamma^{-2}\right)$, we get that for $\alpha^{\prime}$ small enough, this is less than $u / 100$. Thus, conditioned on $\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)} \neq c(x)\right]<\alpha^{\prime} d \gamma^{-2} / m$ and $h_{S, c(S)}$, we get that the entropy of the concept $c$ is no more than $\lg \left(\sum_{i=0}^{u / 100}\binom{u}{i}\right)$ since $c$ is within a Hamming ball of radius $u / 100$ from $h_{S, c(S)}$. Now $\sum_{i=0}^{u / 100}\binom{u}{i} \leq 2^{H_{b}(1 / 100) u}$, where $H_{b}$ is the binary entropy of a Bernoulli random variable with success probability $1 / 100$. Numerical calculations give $H_{b}(1 / 100)=(1 / 100) \lg _{2}(100)+(99 / 100) \lg _{2}(100 / 99)<0.09$. Thus

$$
\begin{equation*}
H\left(c \mid h_{S, c(S)}, \operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)} \neq c(x)\right]<\alpha^{\prime} d \gamma^{-2} / m\right) \leq 0.09 u \tag{7}
\end{equation*}
$$

Now let $X_{S, c}$ be an indicator random variable for the event that $\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)} \neq c(x)\right]<\alpha^{\prime} d \gamma^{-2} / m$. Then $H(c \mid S, c(S)) \leq H\left(c \mid S, c(S), h_{S, c(S)}, X_{S, c}\right)+H\left(X_{S, c}\right)$. Here we remark that we add $h_{S, c(S)}$ in the conditioning for free since it depends only on $S$ and $c(S)$. Adding $X_{S, c}$ costs at most its entropy which satisfies $H\left(X_{S, c}\right) \leq 1$. Since removing variables that we condition on only increases entropy, we get $H(c \mid S, c(S)) \leq H\left(c \mid h_{S, c(S)}, X_{S, c}\right)+1$. Now observe that $H\left(c \mid h_{S, c(S)}, X_{S, c}\right)=\operatorname{Pr}\left[X_{S, c}=1\right] H\left(c \mid h_{S, c(S)}, X_{S, c}=1\right)+\operatorname{Pr}\left[X_{S, c}=0\right] H\left(c \mid h_{S, c(S)}, X_{S, c}=0\right)$. The latter entropy we simply bound by $u$ and the former is bounded by $0.09 u$ by Eq. (7) Thus $H(c \mid S, c(S)) \leq 1+\operatorname{Pr}\left[X_{S, c}=1\right] 0.09 u+\left(1-\operatorname{Pr}\left[X_{S, c}=1\right]\right) u$.

Combining the Bounds. Combining the above with Eq. (6) we conclude that

$$
1+\operatorname{Pr}\left[X_{S, c}=1\right] 0.09 u+\left(1-\operatorname{Pr}\left[X_{S, c}=1\right]\right) u \geq u / 2 .
$$

It follows that $\operatorname{Pr}\left[X_{S, c}=1\right] \leq 2 / 3$. This completes the proof since

$$
\underset{c \sim C}{\mathbb{E}}\left[\operatorname{Pr}_{S \sim D^{m}}\left[\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)}(x) \neq c(x)\right] \geq \alpha^{\prime} \frac{d \gamma^{-2}}{m}\right]\right]=\underset{c \sim C}{\mathbb{E}}\left[\underset{S \sim \mathcal{D}^{m}}{\mathbb{E}}\left[\left(1-X_{S, c}\right)\right]\right]=1-\operatorname{Pr}\left[X_{S, c}=1\right]
$$

and thus

$$
\underset{c \sim C}{\mathbb{E}}\left[\operatorname{Pr}_{S \sim \mathcal{D}^{m}}\left[\operatorname{Pr}_{x \sim \mathcal{D}}\left[h_{S, c(S)}(x) \neq c(x)\right] \geq \alpha^{\prime} \frac{d \gamma^{-2}}{m}\right]\right] \geq \frac{1}{3}
$$

This finishes the proof of Theorem B. 1 .

## C Proofs for the margin-based generalization bound for voting classifiers

This section covers the proofs of all lemmas needed to show the generalization bound for voting classifiers with large margins (Theorem 4 in the main paper) that did not fit into the main text.

## C. 1 Proofs of key properties of $\mathcal{D}_{f, t}$

First, we present the proofs of Lemma 1, 2, and 3 from the main paper covering different properties of the distribution $\mathcal{D}_{f, t}$.
Restatement of Lemma 1. For any $x \in \mathcal{X}$, any $f \in \Delta(\mathcal{H})$ and any $\mu>0$ :

$$
\operatorname{Pr}_{g \sim \mathcal{D}_{f, t}}[|f(x)-g(x)| \geq \mu]<5 \exp \left(-\mu^{2} t / 32\right)
$$

Proof. This lemma follows using standard concentration inequalities: In the first step of sampling $g$ from $\mathcal{D}_{f, t}$, where we draw $t$ i.i.d. hypotheses, it follows from Hoeffding's inequality that the hypothesis $g^{\prime}(x)=(1 / t) \sum_{i=1}^{t} h_{i}^{\prime}(x)$ satisfies

$$
\underset{g^{\prime}}{\operatorname{Pr}}\left[\left|f(x)-g^{\prime}(x)\right| \geq \mu / 2\right] \leq 2 \exp \left(-2(\mu / 2)^{2} t^{2} /(4 t)\right)=2 \exp \left(-\mu^{2} t / 8\right)
$$

In the second step, we first get by a Chernoff bound that $\operatorname{Pr}\left[t^{\prime}<t / 4\right]<\exp (-t / 16)$. Secondly, let us condition on any fixed value of $t^{\prime}$ that is at least $t / 4$. Then $h_{1}, \ldots, h_{t^{\prime}}$ is a uniform sample without replacement from $h_{1}^{\prime}, \ldots, h_{t}^{\prime}$. It follows by a Hoeffding bound without replacement that

$$
\operatorname{Pr}\left[\left|g(x)-g^{\prime}(x)\right| \geq \mu / 2\right] \leq 2 \exp \left(-2(\mu / 2)^{2}\left(t^{\prime}\right)^{2} /\left(4 t^{\prime}\right)\right)<2 \exp \left(-\mu^{2} t / 32\right)
$$

In total, we conclude that

$$
\operatorname{Pr}[|f(x)-g(x)| \geq \mu]<2 \exp \left(-\mu^{2} t / 8\right)+\exp (-t / 16)+2 \exp \left(-\mu^{2} t / 32\right)<5 \exp \left(-\mu^{2} t / 32\right)
$$

Restatement of Lemma 2. For any $x \in \mathcal{X}$, any $f \in \Delta(\mathcal{H})$ and any $\mu \geq 1 / t$ :

$$
\operatorname{Pr}_{g \sim \mathcal{D}_{f, t}}[|g(x)| \leq \mu] \leq 2 \mu \sqrt{t}
$$

Proof. Let $h_{1}^{\prime}, \ldots, h_{t}^{\prime}$ be the hypotheses sampled in the first step of drawing $g$. Define $\sigma_{i}$ to be 1 if $h_{i}^{\prime}$ is sampled in $g$ and -1 otherwise. That is, we have

$$
g(x)=\frac{1}{\left|\left\{i: \sigma_{i}=1\right\}\right|} \sum_{i: \sigma_{i}=1} h_{i}^{\prime}(x)
$$

Let $\Gamma=\sum_{i=1}^{t} h_{i}^{\prime}(x)$. Then

$$
\begin{aligned}
& \Gamma+\sum_{i=1}^{t} \sigma_{i} h_{i}^{\prime}(x) \\
= & \sum_{i: \sigma_{i}=1} h_{i}^{\prime}(x)+\sum_{i: \sigma_{i}=-1} h_{i}^{\prime}(x)+\sum_{i=1}^{t} \sigma_{i} h_{i}^{\prime}(x) \\
= & 2 \sum_{i: \sigma_{i}=1} h_{i}^{\prime}(x)=2 t^{\prime} g(x) .
\end{aligned}
$$

Therefore, $|g(x)| \leq \mu$ if and only if

$$
\left|\frac{\Gamma+\sum_{i} \sigma_{i} h_{i}^{\prime}(x)}{2 t^{\prime}}\right| \leq \mu .
$$

Since $t^{\prime} \leq t$, this implies

$$
\left|\frac{\Gamma+\sum_{i} \sigma_{i} h_{i}^{\prime}(x)}{2 t}\right| \leq \mu .
$$

Hence, we have $\operatorname{Pr}[|g(x)| \leq \mu] \leq \operatorname{Pr}\left[\sum_{i} \sigma_{i} h_{i}^{\prime}(x) \in-\Gamma \pm 2 t \mu\right]$. By Erdös' improved LittlewoodOfford lemma, as long as $2 t \mu \geq 2$, this happens with probability at most $2 t \mu\left(\begin{array}{c}t / 2\rfloor\end{array}\right) 2^{-t}$. The central binomial coefficient satisfies $\binom{t}{\lfloor t / 2\rfloor} \leq 2^{t} / \sqrt{\pi t / 2} \leq 2^{t} / \sqrt{t}$ and thus the probability is at most $2 t \mu / \sqrt{t}=2 \mu \sqrt{t}$.

Finally, we prove Lemma 3 from the main paper:
Restatement of Lemma 3. For any distribution $\mathcal{D}$ over $\mathcal{X} \times\{-1,1\}$, any $t \geq 36$ and any voting classifier $f \in \Delta(\mathcal{H})$ for a hypothesis set $\mathcal{H} \subset \mathcal{X} \rightarrow\{-1,1\}$, we have:

$$
\mathcal{L}_{\mathcal{D}}(f) \leq 3 \mathcal{L}_{\mathcal{D}}^{t}(f)
$$

For the proof, we first need the following auxiliary lemma:
Lemma C.1. For any $x \in \mathcal{X}$ and any $f \in \Delta(\mathcal{H})$, if $f(x) \neq 0$, then

$$
\operatorname{Pr}_{g \sim \mathcal{D}_{f, t}}[\operatorname{sign}(f(x))=\operatorname{sign}(g(x))] \geq 1 / 2-1 / \sqrt{t}
$$

Proof. If we condition on $t^{\prime}$, then $h_{1}, \ldots, h_{t^{\prime}}$ are i.i.d samples from $\mathcal{D}_{f}$ and thus $\operatorname{Pr}[\operatorname{sign}(g(x))=$ $\operatorname{sign}(f(x))] \geq \operatorname{Pr}[\operatorname{sign}(g(x))=-\operatorname{sign}(f(x))]$. We therefore have $\operatorname{Pr}[\operatorname{sign}(f(x))=\operatorname{sign}(g(x))] \geq$ $\operatorname{Pr}[g(x) \neq 0] / 2$, regardless of $t^{\prime}$. We thus only need to bound $\operatorname{Pr}[g(x) \neq 0]$. For this, Lemma 2 with $\mu=1 / t$ implies $\operatorname{Pr}[g(x)=0] \leq 2 / \sqrt{t}$.

Using this lemma, we can prove Lemma 3.
Proof of Lemma 3 from the main paper. Consider any example $(x, y) \in \mathcal{X} \times\{-1,1\}$ for which $\operatorname{Pr}_{g \sim \mathcal{D}_{f, t}}[y g(x) \leq 0]<1 / 2-1 / \sqrt{t}$. By Lemma C.1, it must be the case that $\operatorname{sign}(f(x))=y$. We therefore have by Markov's inequality:

$$
\begin{aligned}
\mathcal{L}_{\mathcal{D}}(f) & \leq \operatorname{Pr}_{(x, y) \sim \mathcal{D}}\left[\operatorname{Pr}_{g \sim \mathcal{D}_{f, t}}[y g(x) \leq 0] \geq 1 / 2-1 / \sqrt{t}\right] \\
& \leq \frac{\mathbb{E}_{(x, y) \sim \mathcal{D}}\left[\operatorname{Pr}_{g \sim \mathcal{D}_{f, t}}[y g(x) \leq 0]\right]}{1 / 2-1 / \sqrt{t}} \\
& =\mathcal{L}_{\mathcal{D}}^{t}(f) /(1 / 2-1 / \sqrt{t}) \\
& \leq 3 \mathcal{L}_{\mathcal{D}}^{t}(f) .
\end{aligned}
$$

## C. 2 Relating generalization error to the ghost set

In the following, we give the proof of Lemma 6 from the main paper:
Restatement of Lemma 6. For $m \geq 2400^{2}$ any $t$ and any $f$, it holds that:

$$
\operatorname{Pr}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f)\right|>\frac{1}{1200}\right] \leq 2 \cdot \operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\right]
$$

Proof. The proof uses standard techniques uneventfully. We can assume $\operatorname{Pr}_{S}\left[\sup _{f \in \Delta(\mathcal{H})} \mid \mathcal{L}_{S}^{t}(f)-\right.$ $\left.\mathcal{L}_{\mathcal{D}}^{t}(f) \mid>1 / 1200\right]>0$, otherwise we are done. We have:

$$
\begin{aligned}
& \operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\right] \\
\geq & \operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400} \wedge \sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f)\right|>\frac{1}{1200}\right] \\
= & \operatorname{Pr}\left[\sup _{S}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{D^{\prime}}^{t}(f)\right|>\frac{1}{1200}\right] \times \\
& \operatorname{Pr}_{S, S^{\prime}}\left[\left.\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\left|\sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \right\rvert\,>\frac{1}{1200}\right] .
\end{aligned}
$$

Fix a data set $S$ in the non-empty event $\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f)\right|>1 / 1200$. Let $f^{*} \in \mathcal{H}$ be any hypothesis on which $\left|\mathcal{L}_{S}^{t}\left(f^{*}\right)-\mathcal{L}_{\mathcal{D}}^{t}\left(f^{*}\right)\right|>1 / 1200$. The hypothesis $f^{*}$ does not depend on $S^{\prime}$ but only on $S$. We now condition on $S$ as well and get:

$$
\begin{aligned}
& \operatorname{Pr}_{S, S^{\prime}}\left[\left.\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\left|S ; \sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \right\rvert\,>\frac{1}{1200}\right] \\
\geq & \operatorname{Pr}_{S^{\prime}}\left[\left.\left|\mathcal{L}_{S}^{t}\left(f^{*}\right)-\mathcal{L}_{S^{\prime}}^{t}\left(f^{*}\right)\right|>\frac{1}{2400}\left|S ; \sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \right\rvert\,>\frac{1}{1200}\right] \\
\geq & \operatorname{Pr}_{S^{\prime}}\left[\left.\left|\mathcal{L}_{S^{\prime}}^{t}\left(f^{*}\right)-\mathcal{L}_{\mathcal{D}^{t}}\left(f^{*}\right)\right| \leq \frac{1}{2400}\left|S ; \sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \right\rvert\,>\frac{1}{1200}\right] .
\end{aligned}
$$

Here the last inequality follows because the events $\left|\mathcal{L}_{S^{\prime}}^{t}\left(f^{*}\right)-\mathcal{L}_{\mathcal{D}}^{t}\left(f^{*}\right)\right| \leq 1 / 2400$ and $\mid \mathcal{L}_{S}^{t}\left(f^{*}\right)-$ $\mathcal{L}_{\mathcal{D}}^{t}\left(f^{*}\right) \mid>1 / 1200$ (which holds by definition of $f^{*}$ ) implies $\left|\mathcal{L}_{S}^{t}\left(f^{*}\right)-\mathcal{L}_{S^{\prime}}^{t}\left(f^{*}\right)\right|>1 / 2400$. Since $f^{*}$ is fixed and independent of $S^{\prime}$, we may now use Hoeffding's inequality to conclude

$$
\operatorname{Pr}_{S^{\prime}}\left[\left|\mathcal{L}_{S^{\prime}}^{t}\left(f^{*}\right)-\mathcal{L}_{\mathcal{D}}^{t}\left(f^{*}\right)\right| \leq 1 / 2400\left|S ; \sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \mid>1 / 1200\right] \geq 1-2 e^{-2(1 / 2400)^{2} m}
$$

For $m \geq 2400^{2}$, this is at least $1-2 e^{-2} \geq 1 / 2$.
Multiplying with $\operatorname{Pr}\left[S\left|\sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \mid>1 / 1200\right]$ and integrating over $S$, we get

$$
\begin{aligned}
& \int_{S}\left(\operatorname{Pr}_{S^{\prime}}\left[\left.\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\left|S ; \sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \right\rvert\,>\frac{1}{1200}\right]\right. \\
& \left.\times \operatorname{Pr}\left[S\left|\sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \left\lvert\,>\frac{1}{1200}\right.\right]\right) \\
\geq & \int_{S} \frac{1}{2} \operatorname{Pr}\left[S\left|\sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \left\lvert\,>\frac{1}{1200}\right.\right] .
\end{aligned}
$$

The right hand side is simply $1 / 2$ and the left hand side is $\operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\right.$ $\left.1 / 2400\left|\sup _{f \in \Delta(\mathcal{H})}\right| \mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f) \mid>1 / 1200\right]$. We finally conclude that for $m \geq 2400^{2}$, we have:

$$
\operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\right] \geq \frac{1}{2} \operatorname{Pr}_{S}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{\mathcal{D}}^{t}(f)\right|>\frac{1}{1200}\right] .
$$

## C. 3 Relation to the growth function

Last, we prove Lemma 8 from the main paper, which is restated here for convenience:
Restatement of Lemma 8 . For any $0<\delta<1$, every $t$, and every $\mu \leq \delta /(9600 \sqrt{t})$, we have

$$
\operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\right] \leq \sup _{P} 2\left|\hat{\Delta}_{\delta}^{\mu}(P)\right| \exp \left(-2 m / 9600^{2}\right) .
$$

Proof. Let $\mu \leq \delta /(9600 \sqrt{t})$. We have that:

$$
\begin{aligned}
& \operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\right] \\
= & \int_{P} \operatorname{Pr}[P] \operatorname{Pr}_{S, S^{\prime}}\left[\left.\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400} \right\rvert\, P\right] \\
\leq & \sup _{P} \operatorname{Pr}_{S, S^{\prime}}\left[\left.\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400} \right\rvert\, P\right] \\
= & \sup _{P} \operatorname{Pr}_{S, S^{\prime}}\left[\left.\sup _{f \in \Delta(\mathcal{H})}\left|\operatorname{Pr}_{(x, y) \sim S, g \sim \mathcal{D}_{f, t}}[y g(x) \leq 0]-\operatorname{Pr}_{(x, y) \sim S^{\prime}, g \sim \mathcal{D}_{f, t}}[y g(x) \leq 0]\right|>\frac{1}{2400} \right\rvert\, P\right] \\
= & \sup _{P} \operatorname{Pr}_{S, S^{\prime}}\left[\left.\sup _{f \in \Delta(\mathcal{H})}\left|\int_{g} \operatorname{Pr}[g]\left(\operatorname{Pr}_{(x, y) \sim S}[y g(x) \leq 0]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[y g(x) \leq 0]\right)\right|>\frac{1}{2400} \right\rvert\, P\right] .
\end{aligned}
$$

We always have $\left(\operatorname{Pr}_{(x, y) \sim S}[y g(x) \leq 0]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[y g(t) \leq 0]\right) \leq 1$, and by Lemma 13, we have $\operatorname{Pr}\left[g \notin \Delta_{\delta}^{\mu}(\mathcal{H}, P)\right] \leq 1 / 4800$, hence

$$
\begin{aligned}
& \left|\int_{g} \operatorname{Pr}[g]\left(\operatorname{Pr}_{(x, y) \sim S}[y g(x) \leq 0]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[y g(x) \leq 0]\right)\right| \\
\leq & \operatorname{Pr}_{g \sim \mathcal{D}_{f, g}}\left[g \notin \Delta_{\delta}^{\mu}(\mathcal{H}, P)\right]+\sup _{g \in \Delta_{\delta}^{\mu}(\mathcal{H}, P)}\left|\operatorname{Pr}_{(x, y) \sim S}[y g(x) \leq 0]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[y g(x) \leq 0]\right| \\
\leq & \frac{1}{4800}+\sup _{g \in \Delta_{\delta}^{\mu}(\mathcal{H}, P)}\left|\operatorname{Pr}_{(x, y) \sim S}[y g(x) \leq 0]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[y g(x) \leq 0]\right| .
\end{aligned}
$$

We thus have

$$
\begin{aligned}
& \operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\right] \\
\leq & \sup _{P} \operatorname{Pr}_{S, S^{\prime}}\left[\left.\frac{1}{4800}+\sup _{g \in \Delta_{\delta}^{\mu}(\mathcal{H}, P)}\left|\operatorname{Pr}_{(x, y) \sim S}[y g(x) \leq 0]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[y g(x) \leq 0]\right|>\frac{1}{2400} \right\rvert\, P\right] \\
= & \sup _{P} \operatorname{Pr}_{S, S^{\prime}}\left[\left.\sup _{g \in \Delta_{\delta}^{\mu}(\mathcal{H}, P)}\left|\operatorname{Pr}_{(x, y) \sim S}[y g(x) \leq 0]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[y g(x) \leq 0]\right|>\frac{1}{4800} \right\rvert\, P\right] .
\end{aligned}
$$

To bound this, let $\hat{\Delta}_{\delta}^{\mu}(P)=\operatorname{sign}\left(\Delta_{\delta}^{\mu}(\mathcal{H}, P)\right)$. Then the above equals:

$$
\sup _{P} \operatorname{Pr}_{S, S^{\prime}}\left[\left.\sup _{h \in \hat{\Delta}_{\delta}^{\mu}(P)}\left|\operatorname{Pr}_{(x, y) \sim S}[h(x) \neq y]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[h(x) \neq y]\right|>\frac{1}{4800} \right\rvert\, P\right] .
$$

Since we have restricted to the fixed set $P$, the set $\hat{\Delta}_{\delta}^{\mu}(P)$ is finite. Hence we may use the union bound to bound the above by

$$
\sup _{P}\left|\hat{\Delta}_{\delta}^{\mu}(P)\right| \sup _{h \in \hat{\Delta}_{\delta}^{\mu}(P)} \operatorname{Pr}_{S, S^{\prime}}\left[\left.\left|\operatorname{Pr}_{(x, y) \sim S}[h(x) \neq y]-\operatorname{Pr}_{(x, y) \sim S^{\prime}}[h(x) \neq y]\right|>\frac{1}{4800} \right\rvert\, P\right] .
$$

For a set $P$ and hypothesis $h \in \hat{\Delta}_{\delta}^{\mu}(P)$, let $p$ denote the fraction of samples $(x, y) \in P$ for which $h(x) \neq y$. Recall that $S$ and the ghost set $S^{\prime}$ are obtained from $P$ by letting $S$ be a uniform set of $m$ samples from $P$ without replacement, and $S^{\prime}$ are the remaining $m$ samples. For shorthand, define $p_{S}=\operatorname{Pr}_{(x, y) \sim S}[h(x) \neq y \mid P]$ and $p_{S^{\prime}}$ symmetrically. Then $p=(1 / 2)\left(p_{S}+p_{S^{\prime}}\right)$. Ву Hoeffding's inequality for sampling without replacement, we have $\operatorname{Pr}_{S, S^{\prime}}\left[\left|p_{S}-p\right|>\varepsilon \mid P\right]=$ $\operatorname{Pr}_{S}\left[\left|p_{S}-p\right|>\varepsilon \mid P\right]<2 \exp \left(-2 \varepsilon^{2} m\right)$. Setting $\varepsilon=1 / 9600$, we get that for $\left|p-p_{S}\right| \leq 1 / 9600$, it must be the case that $p_{S}^{\prime}=2 p-p_{S} \in p \pm 1 / 9600$. Hence $\left|p_{S}-p_{S^{\prime}}\right| \leq 1 / 4800$ and we conclude $\operatorname{Pr}_{S, S^{\prime}}\left[\left|p_{S}-p_{S^{\prime}}\right|>1 / 4800 \mid P\right]<2 \exp \left(-2 m / 9600^{2}\right)$. Thus, we end up with the bound

$$
\operatorname{Pr}_{S, S^{\prime}}\left[\sup _{f \in \Delta(\mathcal{H})}\left|\mathcal{L}_{S}^{t}(f)-\mathcal{L}_{S^{\prime}}^{t}(f)\right|>\frac{1}{2400}\right] \leq \sup _{P} 2\left|\hat{\Delta}_{\delta}^{\mu}(P)\right| \exp \left(-2 m / 9600^{2}\right)
$$

