# Multiclass Learnability and the ERM principle 

Amit Daniely*<br>Sivan Sabato ${ }^{\dagger}$<br>Shai Ben-David ${ }^{\ddagger}$<br>Shai Shalev-Shwartz ${ }^{\S}$


#### Abstract

Multiclass learning is an area of growing practical relevance, for which the currently available theory is still far from providing satisfactory understanding. We study the learnability of multiclass prediction, and derive upper and lower bounds on the sample complexity of multiclass hypothesis classes in different learning models: batch/online, realizable/unrealizable, full information/bandit feedback. Our analysis reveals a surprising phenomenon: In the multiclass setting, in sharp contrast to binary classification, not all Empirical Risk Minimization (ERM) algorithms are equally successful. We show that there exist hypotheses classes for which some ERM learners have lower sample complexity than others. Furthermore, there are classes that are learnable by some ERM learners, while other ERM learner will fail to learn them. We propose a principle for designing good ERM learners, and use this principle to prove tight bounds on the sample complexity of learning symmetric multiclass hypothesis classes (that is, classes that are invariant under any permutation of label names). We demonstrate the relevance of the theory by analyzing the sample complexity of two widely used hypothesis classes: generalized linear multiclass models and reduction trees. We also obtain some practically relevant conclusions.


## 1 Introduction

The task of multiclass learning, that is learning to classify an object into one of many candidate classes, surfaces in many domains including document categorization, object recognition in computer vision, and web advertisement.

The centrality of the multiclass learning problem has spurred the development of various approaches for tackling the task. Many of the methods define a set of possible multiclass predictors, $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ (where $\mathcal{X}$ is the data domain and $\mathcal{Y}$ is the set of labels), called the hypothesis class, and then use the training examples to choose a predictor from $\mathcal{H}$ (for instance Crammer and Singer, 2003). In this paper we study the sample complexity of such hypothesis classes, namely, how many training examples are needed for learning an accurate predictor. This question has been extensively studied and is quite well understood for the binary case, where $|\mathcal{Y}|=2$. In contrast, the existing theory of the multiclass case, where $|\mathcal{Y}|>2$, is much less complete.

We study multiclass sample complexity in several learning models. These models vary in three aspects:

- Interaction with the data source (batch vs. online protocols): In the batch protocol, we assume that the training data is generated i.i.d. by some distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$. The goal is to find a predictor $h$ with a small probability to err, $\operatorname{Pr}_{(x, y) \sim \mathcal{D}}(h(x) \neq y)$, with a high probabilty over training samples. In the online protocol we receive examples one by one and are asked to predict the labels on the fly. Our goal is to make as few prediction mistakes as possible in the worst case (see Littlestone (1987)).
- The underlying labeling mechanism (realizable vs. agnostic): In the realizable case, we assume that the labels of the instances are determined by some $h^{\star} \in \mathcal{H}$. In the agnostic case no restrictions on the labeling rule are imposed, and our goal is to make predictions which are not much worse than the best predictor in $\mathcal{H}$.

[^0]- The type of feedback (full information vs. bandits): In the full information setting, each example is revealed to the learner along with its correct label. In the bandit setting, the learner first sees an unlabeled example, and then outputs its guess for the label. Then a binary feedback is received, indicating only whether the guess was correct or not, but not revealing the correct label in the case of a wrong guess (see for example Auer et al. (2003, 2002), Kakade et al. (2008)).

In Section 2 we consider multiclass sample complexity in the PAC model (namely, the batch protocol with full information). Natarajan (1989) provides a characterization of multiclass PAC learnability in terms of a parameter of $\mathcal{H}$ known as the Natarajan dimension and denoted $d_{N}(\mathcal{H})$ (see section 2.2 for the relevant definitions). For the realizable case we show in Section 2.3 that there are constants $C_{1}, C_{2}$ such that the sample complexity of learning $\mathcal{H}$ with error $\epsilon$ and confidence $1-\delta$ satisfies

$$
\begin{equation*}
C_{1}\left(\frac{d+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) \leq m_{\mathcal{H}}(\epsilon, \delta) \leq C_{2}\left(\frac{d\left(\ln \left(\frac{1}{\epsilon}\right)+\ln (|\mathcal{Y}|)+\ln (d)\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) \tag{1}
\end{equation*}
$$

where $d=d_{N}(\mathcal{H})$. This improves the best previously known upper bound (theorem 5), in which there is a dependence on $\ln (|\mathcal{Y}|) \cdot \ln \left(\frac{1}{\epsilon}\right)$.

The Natarajan dimension is equal to the VC dimension when $|\mathcal{Y}|=2$. However, for larger label sets $\mathcal{Y}$, the bound on the sample complexity is not as tight as the known bound for the binary case, where the gap between the lower and upper bounds is only logarithmic in $1 / \epsilon$. This invokes the challenge of tightening these sample complexity bounds for the multiclass case. A common approach to proving sample complexity bounds for PAC learning is to carefully analyze the sample complexity of ERM learners. In the case of PAC learning, all ERM learners have the same sample complexity (up to a logarithmic factor, see (Vapnik, 1995)). However, rather surprisingly, this is not the case for multiclass learning ${ }^{1}$.

In Section 2.4 we describe a family of concept classes for which there exist "good" ERM learner and "bad" ERM learner with a large gap between their sample complexities. Analyzing these examples, we deduce a rough principle on how to choose a good ERM learner. We also determine the sample complexity of the worst ERM learner for a given concept class, $\mathcal{H}$, up to a multiplicative factor of $O\left(\ln \left(\frac{1}{\epsilon}\right)\right)$. We further show that if $|\mathcal{Y}|$ is infinite, then there are hypotheses classes that are learnable by some ERM learners but not by all ERM learners. In Section 2.5 we employ the suggested principle to derive an improved sample complexity upper bound for symmetric classes $(\mathcal{H}$ is symmetric if $\phi \circ f \in \mathcal{H}$ whenever $f \in \mathcal{H}$ and $\phi$ is a permutation of $\mathcal{Y})$. Symmetric classes are useful, since they are a natural choice when there is no prior knowledge about the relations between the possible labels. Moreover, many popular hypothesis classes that are used in practice are symmetric.

We conjecture that the upper bound obtained for symmetric classes holds for the sample complexity of non-symmetric classes as well. Such a result cannot be implied by uniform convergence alone, since, by the results mentioned above, there always exist bad ERM learners whose sample complexity is higher than this conjectured upper bound. It therefore seems that a proof for our conjecture will require the derivation of new learning rules. We hope that this would lead to new insights in other statistical learning problems as well.

In Section 3 we study multiclass learnability in the online model. We describe a simple generalization of the Littlestone dimension, and derive tight lower and upper bounds on the number, in terms of that dimension, of mistakes the optimal online algorithm will make in the worst case. Section 4 is devoted to a discussion of sample complexity of multiclass learning in the Bandit settings. Finally, in Section 5 we calculate the sample complexity of some popular families of hypothesis classes, which include linear multiclass hypotheses and filter trees, and discuss some practical implications of our bounds.

## 2 Multiclass Learning in the PAC Model

### 2.1 Problem Setting and Notation

For a distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$, the error of a function $f \in \mathcal{H}$ with respect to $\mathcal{D}$ is $\operatorname{Err}(f)=$ $\operatorname{Err}_{\mathcal{D}}(f)=\operatorname{Pr}_{(x, y) \sim \mathcal{D}}(f(x) \neq y)$. A learning algorithm for a class $\mathcal{H}$ is a function, $A: \cup_{n=0}^{\infty}(\mathcal{X} \times \mathcal{Y})^{n} \rightarrow$ $\mathcal{H}$. We denote a training sequence by $S_{m}=\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)$. An ERM learner for class $\mathcal{H}$ is a learning algorithm that for any sample $S_{m}$ returns a function $f \in \mathcal{H}$ that minimizes the number of sample errors $\left|\left\{i \in[m]: f\left(x_{i}\right) \neq y_{i}\right\}\right|$. This work focuses on statistical properties of the learning algorithms and ignores computatational complexity aspects.

[^1]The (agnostic) sample complexity of an algorithm $A$ is the function $m_{A}^{a}$ defined as follows: For every $\epsilon, \delta>0, m_{A}^{a}(\epsilon, \delta)$ is the minimal integer such that for every $m \geq m_{A}^{a}(\epsilon, \delta)$ and every distribution $\mathcal{D}$ on $\mathcal{X} \times \mathcal{Y}$,

$$
\begin{equation*}
\underset{S_{m} \sim \mathcal{D}^{m}}{\operatorname{Pr}}\left(\underset{\mathcal{D}}{\operatorname{Err}}\left(A\left(S_{m}\right)\right)>\inf _{f \in \mathcal{H}} \operatorname{Err}_{\mathcal{D}}(f)+\epsilon\right) \leq \delta \tag{2}
\end{equation*}
$$

If there is no integer satisfying these requirements, define $m_{A}^{a}(\epsilon, \delta)=\infty$. The (agnostic) sample complexity of a class $\mathcal{H}$ is

$$
m_{\mathcal{H}}^{a}(\epsilon, \delta)=\inf _{A} m_{A}^{a}(\epsilon, \delta)
$$

where the infimum is taken over all learning algorithms.
We say that a distribution $\mathcal{D}$ is realizable by a hypothesis class $\mathcal{H}$ if there exists some $f \in \mathcal{H}$ such that $\operatorname{Err}_{\mathcal{D}}(f)=0$. The realizable sample complexity of an algorithm $A$ for a class $\mathcal{H}$, denoted $m_{A}^{r}$, is the minimal integer such that for every $m \geq m_{A}^{r}(\epsilon, \delta)$ and every distribution $\mathcal{D}$ on $\mathcal{X} \times \mathcal{Y}$ which is realizable by $\mathcal{H}$, Equation. (2) holds. The realizable sample complexity of a class $\mathcal{H}$ is $m_{\mathcal{H}}^{r}(\epsilon, \delta)=\inf _{A} m_{A}^{r}(\epsilon, \delta)$ where the infimum is taken over all learning algorithms.

### 2.2 Known Sample Complexity Results

We first survey some known results regarding the sample complexity of multiclass learning. We start with the realizable case and then discuss the agnostic case. Given a subset $S \subseteq \mathcal{X}$, we denote $\left.\mathcal{H}\right|_{S}=\left\{\left.f\right|_{S}: f \in \mathcal{H}\right\}$. Recall the definition of the Vapnik-Chervonenkis dimension (Vapnik, 1995):
Definition 1 (VC dimension) Let $\mathcal{H} \subseteq\{0,1\}^{\mathcal{X}}$ be a hypothesis class. A subset $S \subseteq \mathcal{X}$ is shattered by $\mathcal{H}$ if $\left.\mathcal{H}\right|_{S}=\{0,1\}^{S}$. The VC-dimension of $\mathcal{H}$, denoted $\mathrm{VC}(\mathcal{H})$, is the maximal cardinality of a subset $S \subseteq \mathcal{X}$ that is shattered by $\mathcal{H}$.

The VC-dimension is cornerstone in statistical learning theory as it characterizes the sample complexity of a binary hypothesis class. Namely

Theorem 2 (Vapnik, 1995) There are absolute constants $C_{1}, C_{2}>0$ such that the realizable sample complexity of every hypothesis class $\mathcal{H} \subseteq\{0,1\}^{\mathcal{X}}$ satisfies

$$
C_{1}\left(\frac{\mathrm{VC}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) \leq m_{\mathcal{H}}^{r}(\epsilon, \delta) \leq C_{2}\left(\frac{\mathrm{VC}(\mathcal{H}) \ln \left(\frac{1}{\epsilon}\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)
$$

Moreover, the upper bound is attained by any ERM learner.
It is natural to seek a generalization of the VC-Dimension to hypothesis classes of non-binary functions. A straightforward attempt is to redefine shattering of $S \subset \mathcal{X}$ by the property $\left.\mathcal{H}\right|_{S}=\mathcal{Y}^{S}$. However, this requirement is too strong and does not lead to tight bounds on the sample complexity. Instead, we recall two alternative generalizations, introduced by Natarajan (1989). In both definitions, shattering is redefined to require that for any partition of $S$ into $T$ and $S \backslash T$, there exists a $g \in \mathcal{H}$ whose behavior on $T$ differs from its behavior on $S \backslash T$. The two definitions differ in how "different behavior" is defined.

Definition 3 (Graph dimension and Natarajan dimension) Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a hypothesis class and let $S \subseteq \mathcal{X}$. We say that $\mathcal{H}$ G-shatters $S$ if there exists an $f: S \rightarrow \mathcal{Y}$ such that for every $T \subseteq S$ there is a $g \in \mathcal{H}$ such that

$$
\forall x \in T, g(x)=f(x), \text { and } \forall x \in S \backslash T, g(x) \neq f(x)
$$

We say that $\mathcal{H} \mathrm{N}$-shatters $S$ if there exist $f_{1}, f_{2}: S \rightarrow \mathcal{Y}$ such that $\forall y \in S, f_{1}(y) \neq f_{2}(y)$, and for every $T \subseteq S$ there is a $g \in \mathcal{H}$ such that

$$
\forall x \in T, g(x)=f_{1}(x), \text { and } \forall x \in S \backslash T, g(x)=f_{2}(x)
$$

The graph dimension of $\mathcal{H}$, denoted $d_{G}(\mathcal{H})$, is the maximal cardinality of a set that is $G$-shattered by $\mathcal{H}$. The Natarajan dimension of $\mathcal{H}$, denoted $d_{N}(\mathcal{H})$, is the maximal cardinality of a set that is $N$-shattered by $\mathcal{H}$.

Both of these dimensions coincide with the VC-dimension for $|\mathcal{Y}|=2$. Note also that we always have $d_{N} \leq d_{G}$.

By reductions from and to the binary case, it is not hard to show, similarly to Natarajan (1989) and Ben-David et al. (1995) (see Appendix A for a full proof), that

Theorem 4 For the constants $C_{1}, C_{2}$ from theorem 2, for every $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ we have

$$
C_{1}\left(\frac{d_{N}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) \leq m_{\mathcal{H}}^{r}(\epsilon, \delta) \leq C_{2}\left(\frac{d_{G}(\mathcal{H}) \ln \left(\frac{1}{\epsilon}\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)
$$

Moreover, the upper bound is attained by any ERM learner.
From this theorem it follows that the finiteness of the Natarajan dimension is a necessary condition for learnability, and the finiteness of the graph dimension is a sufficient condition for learnability. In Ben-David et al. (1995) it was proved that for every concept class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$,

$$
\begin{equation*}
d_{N}(\mathcal{H}) \leq d_{G}(\mathcal{H}) \leq 4.67 \log _{2}(|\mathcal{Y}|) d_{N}(\mathcal{H}) \tag{3}
\end{equation*}
$$

It follows that if $|\mathcal{Y}|<\infty$ then the finiteness of the Natarajan dimension is a necessary and sufficient condition for learnability. Incorporating Equation. (3) into theorem 4, it can be seen that the Natarajan dimension, as well as the graph dimension, characterize the sample complexity of $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ up to a multiplicative factor of $O\left(\log (|\mathcal{Y}|) \log \left(\frac{1}{\epsilon}\right)\right)$. Precisely,
Theorem 5 (Ben-David et al., 1995) For the constants $C_{1}, C_{2}$ from theorem 2,

$$
C_{1}\left(\frac{d_{N}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) \leq m_{\mathcal{H}}^{r}(\epsilon, \delta) \leq C_{2}\left(\frac{d_{N}(\mathcal{H}) \cdot \ln (|\mathcal{Y}|) \cdot \ln \left(\frac{1}{\epsilon}\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)
$$

Moreover, the upper bound is attained by any ERM learner.
A similar analysis can be performed for the agnostic case. For binary classification we have that for every hypothesis class $\mathcal{H} \subseteq\{0,1\}^{\mathcal{X}}$,

$$
\begin{equation*}
m_{\mathcal{H}}^{a}(\epsilon, \delta)=\Theta\left(\frac{1}{\epsilon^{2}}\left(V C(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)\right)\right) \tag{4}
\end{equation*}
$$

and this is attained by any ERM learner. Here too it is possible to obtain by reduction from and to the binary case that for every hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$,

$$
\begin{equation*}
\Omega\left(\frac{1}{\epsilon^{2}}\left(d_{N}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)\right)\right) \leq m_{\mathcal{H}}^{a}(\epsilon, \delta) \leq O\left(\frac{1}{\epsilon^{2}}\left(d_{G}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)\right)\right) \tag{5}
\end{equation*}
$$

By Equation. (3) we have

$$
\begin{equation*}
m_{\mathcal{H}}^{a}(\epsilon, \delta)=O\left(\frac{1}{\epsilon^{2}}\left(\log (|\mathcal{Y}|) \cdot d_{N}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)\right)\right) \tag{6}
\end{equation*}
$$

Thus in the agnostic case as well, the Natarajan dimension characterizes the agnostic sample complexity up to a multiplicative factor of $O(\log (|\mathcal{Y}|))$. Here too, all of these bounds are attained by any ERM learner.

### 2.3 An Improved Result for the Realizable Case

The following theorem provides a sample complexity upper bound which can be better than Theorem 5 when $\ln \left(d_{N}(\mathcal{H})\right) \ll \ln (|\mathcal{Y}|) \cdot \ln \left(\frac{1}{\epsilon}\right)$. The proof of the theorem is given in Appendix A. While the proof is a simple adaptation of previous results, we find it valuable to present this result here, as we could not find it in the research literature.

Theorem 6 For every concept class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$,

$$
m_{\mathcal{H}}^{r}(\epsilon, \delta)=O\left(\frac{d_{N}(\mathcal{H})\left(\ln \left(\frac{1}{\epsilon}\right)+\ln (|\mathcal{Y}|)+\ln \left(d_{N}(\mathcal{H})\right)\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) .
$$

Moreover, the bound is attained by any ERM learner.
Theorem 6 is the departure point of our research. As indicated above, one of our objectives is to prove sample complexity bounds for the multiclass case with a ratio of $O\left(\ln \left(\frac{1}{\epsilon}\right)\right)$ between the upper bound and the lower bound, as in the binary case. In the next section we show that such an improvement cannot be attained by uniform convergence analysis, since the ratio between the sample complexity of the worst ERM learner and the best ERM learner of a given hypothesis class might be as large as $\ln (|\mathcal{Y}|)$.

### 2.4 The Gap between "Good ERM" and "Bad ERM"

The tight bounds in the binary case given in Theorem 2 are attained by any ERM learner. In contrast to the binary case, we now show that in the multiclass case there can be a significant sample complexity gap between different ERM learners. Moreover, in the case of classification with an infinite number of classes, there are learnable hypothesis classes that some ERM learners fail to learn. We begin with showing that the graph dimension determines the sample complexity of the worst ERM learner up to a multiplicative factor of $O\left(\ln \left(\frac{1}{\epsilon}\right)\right)$.

Theorem 7 There are absolute constants $C_{1}, C_{2}>0$ such that for every hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ and every ERM learner A,

$$
m_{A}^{r}(\epsilon, \delta) \leq C_{2}\left(\frac{d_{G}(\mathcal{H}) \ln \left(\frac{1}{\epsilon}\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)
$$

Moreover, there is an ERM learner $A_{\mathrm{bad}}$ such that

$$
\begin{equation*}
m_{A_{\mathrm{bad}}}^{r}(\epsilon, \delta) \geq C_{1}\left(\frac{d_{G}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) \tag{7}
\end{equation*}
$$

Proof: The upper bound on $m_{A}^{r}$ is just a restatement of theorem 4. It remains to prove that there exists an ERM learner, $A_{\text {bad }}$, satisfying (7). We shall first consider the case where $d=d_{G}(\mathcal{H})<\infty$.

Let $S=\left\{x_{0}, \ldots, x_{d-1}\right\} \subseteq \mathcal{X}$ be a set which is $G$-Shattered by $\mathcal{H}$ using the function $f_{0}$. Let $A_{\text {bad }}$ be an ERM learner with the property that upon seeing a sample whose instances are in $T \subseteq S$, and whose labels are determined by $f_{0}$, it returns $f \in \mathcal{H}$ such that $f$ equals to $f_{0}$ on $T$ and $f$ is different from $f_{0}$ on $S \backslash T$. The existence of such an $f$ follows form the assumption that $S$ is G-shattered using $f_{0}$.

Fix $\delta<e^{-1 / 6}$ and let $\epsilon$ small enough such that $1-2 \epsilon \geq e^{-4 \epsilon}$. Define a distribution on $\mathcal{X}$ by setting $\operatorname{Pr}\left(x_{0}\right)=1-2 \epsilon$ and for all $1 \leq i \leq d-1, \operatorname{Pr}\left(x_{i}\right)=\frac{2 \epsilon}{d-1}$. Suppose that the correct hypothesis is $f_{0}$ and let the sample size be $m$. Clearly, the hypothesis returned by $A_{\mathrm{bad}}$ will err on all the examples from $S$ which are not in the sample. By Chernoff's bound, if $m \leq \frac{d-1}{6 \epsilon}$, then with probability $\geq e^{-\frac{1}{6}} \geq \delta$, the sample will include no more than $\frac{d-1}{2}$ examples from $S$. Thus the returned hypothesis will have error $\geq \epsilon$. Moreover, the probability that the sample includes only $x_{0}$ (and thus $A_{\mathrm{bad}}$ will return a hypothesis with error $2 \epsilon$ ) is $(1-2 \epsilon)^{m} \geq e^{-4 \epsilon m}$, which is more than $\delta$ if $m \leq \frac{1}{4 \epsilon} \ln \left(\frac{1}{\delta}\right)$. We therefore obtain that

$$
m_{A_{\mathrm{bad}}}^{r}(\epsilon, \delta) \geq \max \left\{\frac{d-1}{6 \epsilon}, \frac{1}{2 \epsilon} \ln (1 / \delta)\right\} \geq \frac{d-1}{12 \epsilon}+\frac{1}{4 \epsilon} \ln (1 / \delta)
$$

as required. If $d_{G}(\mathcal{H})=\infty$ then the argument above can be repeated for a sequence of pairwise disjoint G-shattered sets $S_{n}, n=1,2, \ldots$ with $\left|S_{n}\right|=n$.

The following example shows that in some cases there are learning algorithms that are much better than the worst ERM:

Example 8 (A Large Gap Between ERM Learners) Let $\mathcal{X}_{0}$ be any finite or countable domain set and let $\mathcal{X}$ be some subset of $\mathcal{X}_{0}$. Let $\mathcal{P}_{f}(\mathcal{X})$ denote the collection of finite and co-finite subsets $A \subseteq \mathcal{X}$. For every $A \in \mathcal{P}_{f}(\mathcal{X})$, define $f_{A}: \mathcal{X}_{0} \rightarrow \mathcal{P}_{f}(\mathcal{X}) \cup\{*\}$ by

$$
f_{A}(x)= \begin{cases}A & \text { if } x \in A \\ * & \text { otherwise }\end{cases}
$$

and consider the concept family $\mathcal{H}_{\mathcal{X}}=\left\{f_{A}: A \in \mathcal{P}_{f}(\mathcal{X})\right\}$. We first note that any ERM learner that sees an example of the form $(x, A)$ for some $A \subseteq \mathcal{X}$ must return the hypothesis $f_{A}$, thus to define an ERM learner we only have to specify the hypothesis it returns upon seeing a sample of the form $S_{m}=\left\{\left(x_{1}, *\right), \ldots,\left(x_{m}, *\right)\right\}$. Note also that $\mathcal{X}$ is $G$-shattered using the function $f_{\emptyset}$, and therefore $d_{G}\left(\mathcal{H}_{\mathcal{X}}\right) \geq|\mathcal{X}|$ (it is easy to see that, in fact $\left.d_{G}\left(\mathcal{H}_{\mathcal{X}}\right)=|\mathcal{X}|\right)$.

We consider two ERM learners - $A_{\text {good }}$, which on a sample of the form $S_{m}$ returns the hypothesis $f_{\emptyset}$, and $A_{\text {bad }}$, which, upon seeing $S_{m}$, returns $f_{\left\{x_{1}, \ldots, x_{m}\right\}^{c}}$, thus satisfying the specification of a bad ERM algorithm from the proof of Theorem 7. It follows that the sample complexity of $A_{\mathrm{b}}$ is $\Omega\left(\frac{|\mathcal{X}|}{\epsilon}+\frac{1}{\epsilon} \ln \left(\frac{1}{\delta}\right)\right)$. On the other hand,

Claim 9 The sample complexity of $A_{\text {good }}$ is at most $\frac{1}{\epsilon} \ln \frac{1}{\delta}$.

Proof: Let $\mathcal{D}$ be a distribution over $\mathcal{X}_{0}$ and suppose that the correct labeling is $f_{A}$. Let $m$ be the size of the sample. For any sample, $A_{\text {good }}$ returns either $f_{\emptyset}$ or $f_{A}$. If it returns $f_{A}$ then its generalization error is zero. Thus, it returns a hypothesis with error $\geq \epsilon$ only if $\operatorname{Pr}_{\mathcal{D}}(A) \geq \epsilon$ and all the $m$ examples in the sample are from $A^{c}$. Assume $m \geq \frac{1}{\epsilon} \ln \left(\frac{1}{\delta}\right)$, then probability of the latter event is no more than $(1-\epsilon)^{m} \leq e^{-\epsilon m} \leq \delta$.

Since $\mathcal{X}$ can be infinite in the above example we conclude that
Corollary 10 There exist sets $\mathcal{X}, \mathcal{Y}$ and a hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, such that $\mathcal{H}$ is learnable by some ERM learner but is not learnable by some other ERM learner.

What is the crucial feature that makes $A_{\text {good }}$ better than $A_{\mathrm{bad}}$ ? If the correct labeling is $f_{A} \in \mathcal{H}_{\mathcal{X}}$, then for any sample, $A_{\text {good }}$ might return at most one of two functions - namely $f_{A}$ or $f_{\emptyset}$. On the other hand, if the sample is labeled by the function $f_{\emptyset}, A_{\text {bad }}$ might return every function in $\mathcal{H}_{\mathcal{X}}$. Thus, to return a hypothesis with error $\leq \epsilon, A_{\text {good }}$ needs to reject only one hypothesis while $A_{\text {bad }}$ needs to reject many more. We conclude the following (rough) principle: A good ERM is an ERM that, for every target hypothesis, consider a small number of hypotheses.

Next, we formalize the above intuition by proving a general theorem that enables us to derive sample complexity bounds for ERM learners that are designed using the above principle. Fix a hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$. We view an ERM learner as an operator that for any $f \in \mathcal{H}, S \subseteq \mathcal{X}$ takes the partial function $\left.f\right|_{S}$ as input and extends it to a function $g=A\left(\left.f\right|_{S}\right) \in \mathcal{H}$ such that $\left.g\right|_{S}=\left.f\right|_{S}$. For every $f \in \mathcal{H}$, denote by $F_{A}(f)$ the set of all the functions that the algorithm $A$ might return upon seeing a sample of the form $\left\{\left(x_{i}, f\left(x_{i}\right)\right)\right\}_{i=1}^{m}$ for some $m \geq 0$. Namely,

$$
F_{A}(f)=\left\{A\left(\left.f\right|_{S}\right): S \subseteq \mathcal{X},|S|<\infty\right\}
$$

To provide an upper bound on $m_{A}^{r}(\epsilon, \delta)$, it suffices to show that for every $f \in \mathcal{H}$, with probability at least $1-\delta$, all the functions with error at least $\epsilon$ in $F_{A}(f)$ will be rejected after seeing $m$ examples. This is formalized in the following theorem.

Theorem 11 Let $A$ be an ERM learner for a hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$. Define the growth function of $A$ by $\Pi_{A}(m)=\sup _{f \in \mathcal{H}} \Pi_{F_{A}(f)}(m)$, where for $F \subseteq \mathcal{Y}^{\mathcal{X}}, \Pi_{F}(m)=\sup \left\{|F|_{S}|: S \subseteq \mathcal{X},|S| \leq m\}\right.$. Then

$$
m_{A}^{r}(\epsilon, \delta) \leq \min \left\{m: \Pi_{A}(2 m) 2^{1-\frac{\epsilon m}{2}}<\delta\right\}
$$

The theorem immediately follows from the following lemma.
Lemma 12 (The Double Sampling Lemma) Let $A$ be an ERM learner. Fix a distribution $\mathcal{D}$ over $\mathcal{X}$ and a function $f_{0} \in \mathcal{H}$. Denote by $A_{m}$ the event that, after seeing $m$ i.i.d. examples drawn from $\mathcal{D}$ and labeled by $f_{0}$, A returns a hypothesis with error at least $\epsilon$. Then $\operatorname{Pr}\left(A_{m}\right) \leq$ $2 \cdot \Pi_{A}(2 m) 2^{-\frac{\epsilon m}{2}}$.

Proof: Let $S_{1}$ and $S_{2}$ be two samples of $m$ i.i.d. examples labeled by $f_{0}$. Let $B_{m}$ be the event that there exists a function $f \in \mathcal{H}$ with error at least $\epsilon$, such that (1) $f$ is not rejected by $S_{1}$ (i.e. $f_{0}(x)=$ $f(x)$ for all examples $x$ in $S_{1}$ ), and (2) there exist at least $\frac{\epsilon m}{2}$ examples $\left(x, f_{0}(x)\right)$ in $S_{2}$ for which $f(x) \neq f_{0}(x)$. By Chernoff's bound, for $m=\Omega\left(\frac{1}{\epsilon}\right), \operatorname{Pr}\left(B_{m}\right)=\operatorname{Pr}\left(B_{m} \mid A_{m}\right) \operatorname{Pr}\left(A_{m}\right) \geq \frac{1}{2} \operatorname{Pr}\left(A_{m}\right)$. W.l.o.g., we can assume that $S_{1}, S_{2}$ are generated as follows: First, $2 m$ examples are drawn to create a sample $U$. Then $S_{1}$ and $S_{2}$ are generated by selecting a random partition of $U$ into two samples of size $m$. Now, $\operatorname{Pr}\left(B_{m}\right)$ is bounded by the probability that there is an $\left.f \in \mathcal{H}\right|_{U}$ such that (1) there are at least $\frac{\epsilon m}{2}$ examples in $U$ such that $f$ disagrees with $f_{0}$ on these examples and (2) all of these examples are located in $S_{2}$. For a single $\left.f \in \mathcal{H}\right|_{U}$ that disagrees with $f_{0}$ on $l \geq \frac{\epsilon m}{2}$ samples, the probability that all these examples are located in $S_{2}$ is $\binom{m}{l} /\binom{2 m}{l} \leq 2^{-l} \leq 2^{-\frac{c m}{2}}$. Thus, using the union bound we obtain that $\operatorname{Pr}\left(B_{m}\right) \leq|\mathcal{H}|_{U} \left\lvert\, 2^{-\frac{\epsilon m}{2}} \leq \Pi(2 m) 2^{-\frac{\epsilon m}{2}}\right.$.

The bound in theorem 6 is based on the (trivial) inequality $\Pi_{A} \leq \Pi_{\mathcal{H}}$. However, as Example 8 shows, $\Pi_{A}$ can be much smaller than $\Pi_{\mathcal{H}}$. As we shall see in the sequel, we can apply the double sampling lemma to get better sample complexity bounds for "good" ERM learners. The key tool for these sample complexity bounds is Lemma 14, that is, in turn, based on the following combinatorial result:

Lemma 13 (Natarajan, 1989) For every hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}},|\mathcal{H}| \leq|\mathcal{X}|^{d_{N}(\mathcal{H})}|\mathcal{Y}|^{2 d_{N}(\mathcal{H})}$.

Lemma 14 Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a class of functions. Assume that for some number $r$, for every $h \in \mathcal{H}$, the size of the range of $h$ is at mostr. Let $A$ be an algorithm such that, for some set of values $Y^{\prime} \subseteq \mathcal{Y}$, for every $f \in \mathcal{H}$, and every sample $S_{m}=\left(\left(x_{1}, f\left(x_{1}\right)\right), \ldots\left(x_{m}, f\left(x_{m}\right)\right)\right)$, the function returned by $A$ on input $S_{m}$ is consistent with $S_{m}$ and has its values in the set $\left\{f\left(x_{1}\right), \ldots, f\left(x_{m}\right)\right\} \cup Y^{\prime}$. Then,

$$
m_{A}^{r}(\epsilon, \delta)=O\left(\frac{d_{N}(\mathcal{H})\left(\ln \left(\frac{1}{\epsilon}\right)+\ln \left(\max \left\{r,\left|Y^{\prime}\right|\right\}\right)\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)
$$

Proof: The assumptions of the lemma imply that, for every $f \in \mathcal{H}$, the range of the functions in $F_{A}(f)$ is contained in the union of $Y^{\prime}$ and the range of $f$. Therefore, using Lemma 13 we obtain that $\Pi_{A}(2 m) \leq(2 m)^{d_{N}(\mathcal{H})}\left(\left|Y^{\prime}\right|+r\right)^{2 d_{N}(\mathcal{H})}$, and the bound follows from Theorem 11 .

Note that classes in which each function $h \in H$ uses at most $r$ values, for some $r<d_{N}(H) \log (|\mathcal{Y}|)$, can have a large range $\mathcal{Y}$ and a graph dimension that is significantly larger than their Natarajan dimension. In such cases, we may be able to show a gap between the sample complexity of bad and good ERM learners, by applying the lower bound from Theorem 7. In particular, we get such a result for the following family of hypotheses classes, which generalizes Example 8.

Corollary 15 Let $\mathcal{H}$ be a class of functions from $\mathcal{X}$ to some range set $\mathcal{Y}$, such that, for some value $y_{0} \in \mathcal{Y}$, for every $h \in H$, the range of $h$ contains at most one value besides $y_{0}$. Assume also that $\mathcal{H}$ contains the constant $y_{0}$ function. Let d denote the Natarajan dimension of $\mathcal{H}$. Then there exists an ERM learning algorithm $A$ for $H$ such that the $(\epsilon, \delta)$ sample complexity of $A$ is

$$
O\left(\frac{d \cdot \ln (1 / \epsilon)+\ln (1 / \delta)}{\epsilon}\right)
$$

Every class in that family that has a large graph dimension will therefore realize a gap between the sample complexities of different ERM learners.

Example 16 Consider the set of all balls in $\mathbb{R}^{n}$ and, for each such ball, $B=B(z, r)$ with center $z$ and radius $r$, let $h_{B}$ be the function defined by $h_{B}(x)=z$ if $x \in B$ and $h_{B}(x)=\star$ otherwise. Let $\mathcal{H}_{\mathcal{B}^{n}}=\left\{h_{B}: B=B(z, r)\right.$ for some $\left.z \in \mathbb{R}^{n}, r \in \mathbb{R}\right\} \cup\left\{h_{\star}\right\}$ (where $h_{\star}$ is the constant $\star$ function). It is not hard to see that $d_{N}\left(\mathcal{H}_{\mathcal{B}^{n}}\right)=1$ and $d_{G}\left(\mathcal{H}_{\mathcal{B}^{n}}\right)=n+1$. Furthermore, let $A_{\text {good }}$ be the ERM learner that for every sample $S=\left(x_{1}, f\left(x_{1}\right)\right), \ldots\left(x_{m}, f\left(x_{m}\right)\right)$, returns $h_{B_{S}}$, where $B_{S}$ is the minimal ball that is consistent with the sample. Note that this algorithm uses, for every $f \in \mathcal{H}_{\mathcal{B}^{n}}$ and every sample $S$ labeled by such $f$, at most one value (the value $\star$ ) on top of the values $\left\{f\left(x_{1}\right), \ldots, f\left(x_{m}\right)\right\}$.

In this case, Theorem 7 implies that for some constant $C_{1}$, there exists a bad ERM learner, $A_{\mathrm{bad}}$ such that

$$
m_{A_{\mathrm{bad}}}^{r}(\epsilon, \delta) \geq C_{1}\left(\frac{n+\ln (1 / \delta)}{\epsilon}\right)
$$

On the other hand, Lemma 14 implies that there is a good ERM learner, $A_{\text {good }}$ and a constant $C_{2}$ for which

$$
m_{A_{\text {good }}}^{r}(\epsilon, \delta) \leq C_{2}\left(\frac{\ln (1 / \epsilon)+\ln (1 / \delta)}{\epsilon}\right)
$$

Note that, if one restricts the hypothesis class to allow only balls that have their centers in some finite set of grid points, the class uses only a finite range of labels. However, if such a grid is sufficiently dense, the sample complexities of both algorithms, $A_{\text {bad }}$ and $A_{\text {good }}$, would not change.

### 2.5 Symmetric Classes

The principle for choosing a good ERM leads to tight bounds on the sample complexity of symmetric classes. Recall that a class $\mathcal{H}$ is called symmetric if for any $f \in \mathcal{H}$ and any permutation $\phi$ on labels, we have that $\phi \circ f \in \mathcal{H}$ as well.

Theorem 17 There are absolute constants $C_{1}, C_{2}$ such that for every symmetric hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$

$$
C_{1}\left(\frac{d_{N}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) \leq m_{\mathcal{H}}^{r}(\epsilon, \delta) \leq C_{2}\left(\frac{d_{N}(\mathcal{H})\left(\ln \left(\frac{1}{\epsilon}\right)+\ln \left(d_{N}(\mathcal{H})\right)\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)
$$

A key observation that enables us to employ our principle in this case is:

Lemma 18 Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a symmetric hypothesis class of Natarajan dimension $d$. Then, the range of any $f \in \mathcal{H}$ is of size at most $2 d+1$.
Proof: If $|\mathcal{Y}| \leq 2 d+1$ we are done. Thus assume that there are $2 d+2$ distinct elements $y_{1}, \ldots, y_{2 d+2} \in \mathcal{Y}$. Assume to the contrary that there is a hypothesis $f \in \mathcal{H}$ with a range of more than $d$ values. Thus there is a set $S=\left\{x_{1}, \ldots, x_{d+1}\right\} \subseteq \mathcal{X}$ such that $\left.f\right|_{S}$ has $d+1$ values in its range. It follows that $\mathcal{H} \mathrm{N}$-shatters $S$, thus reaching a contradiction. Indeed, since $\mathcal{H}$ is symmetric, there are functions $f_{0}, f_{1} \in \mathcal{H}$ such that $f_{j}\left(x_{i}\right)=y_{j(d+1)+i}$. Similarly, for every $T \subseteq S$, there is a $g \in \mathcal{H}$ such that $g(x)=f_{0}(x)$ for every $x \in T$ and $g(x)=f_{1}(x)$ for every $x \in S \backslash T$.

We are now ready to prove Theorem 17.
Proof: (of Theorem 17) The lower bound is a restatement of Theorem 4. For the upper bound, we define an algorithm $A$ that conforms to the conditions in Lemma 14: Fix a set $\mathcal{Y}^{\prime} \subseteq \mathcal{Y}$ of size $\left|\mathcal{Y}^{\prime}\right|=\min \left\{|\mathcal{Y}|, 2 d_{N}(\mathcal{H})+1\right\}$. Given a sample $\left(x_{1}, f\left(x_{1}\right)\right), \ldots,\left(x_{m}, f\left(x_{m}\right)\right)$, $A$ returns a hypothesis that is consistent with the sample and that attains only values in $\left\{f\left(x_{1}\right), \ldots, f\left(x_{m}\right)\right\} \cup \mathcal{Y}^{\prime}$. It is possible due to symmetry and Lemma 18.

A similar analysis can be performed for the agnostic case. Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a symmetric hypothesis class. Let $\mathcal{Y}^{\prime} \subseteq \mathcal{Y}$ be an arbitrary set of size $\min \left\{|\mathcal{Y}|, 4 d_{N}(G)+2\right\}$. Denote $\mathcal{H}^{\prime}=\{f \in \mathcal{H}: f(\mathcal{X}) \subseteq$ $\left.\mathcal{Y}^{\prime}\right\}$. Using lemma 18 and symmetry, it is easy to see that $d_{G}(\mathcal{H})=d_{G}\left(\mathcal{H}^{\prime}\right)$ and $d_{N}(\mathcal{H})=d_{N}\left(\mathcal{H}^{\prime}\right)$. By equation 3, we conclude that $d_{G}(\mathcal{H})=O\left(\log \left(d_{N}(\mathcal{H})\right) \cdot d_{N}(\mathcal{H})\right)$. Using equation 5 we obtain a sample complexity bound of

$$
m_{\mathcal{H}}^{a}(\epsilon, \delta)=O\left(\frac{1}{\epsilon^{2}}\left(\log \left(\min \left\{d_{N}(\mathcal{H}),|\mathcal{Y}|\right\}\right) \cdot d_{N}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)\right)\right)
$$

which is better than Equation. (6). Moreover, the ratio between this bound and the lower bound (Equation. (5)) is $O\left(\log \left(d_{N}(\mathcal{H})\right)\right)$ regardless of $|\mathcal{Y}|$. Note that this bound is attained by any ERM. We present the following open question:
Open question 19 Examples 8 and 16 show that there are (non-symmetric) hypothesis classes with a ratio of $\Omega(\ln (|\mathcal{Y}|))$ between the sample complexities of the worst ERM learner and the best ERM learner. How large can this gap be for symmetric hypothesis classes?

## 3 Multiclass Learning in the Online Model

Learning in the online model is conducted in a sequence of consecutive rounds. On each round $t=1,2, \ldots$, the environment presents a sample $x_{t} \in \mathcal{X}$, the algorithm should predict a value $\hat{y_{t}} \in \mathcal{Y}$, and then the environment reveals the correct value $y_{t} \in \mathcal{Y}$. The prediction at time $t$ can be based only on the examples $x_{1}, \ldots, x_{t}$ and the previous outcomes $y_{1}, \ldots, y_{t-1}$. We start with the realizable case, in which we assume that for some function $f \in \mathcal{H}$, all the outcomes are evaluations of $f$, namely, $y_{t}=f\left(x_{t}\right)$. Given an online learning algorithm, $A$, define its (realizable) sample complexity, $\mathcal{M}(A)$, to be the maximal number of wrong predictions that it might make on a legal sequence of any length.

The sample complexity of online learning has been studied by Littlestone (1987), who showed that a combinatorial measure, called the Littlestone dimension, characterizes the sample complexity of online learning. We now propose a generalization of the Littlestone dimension to classes of nonbinary functions.

Consider a rooted tree $T$ whose internal nodes are labeled by $\mathcal{X}$ and whose edges are labeled by $\mathcal{Y}$, such that the labels on edges from a parent to its child nodes are all different from each other. The tree $T$ is shattered by $\mathcal{H}$ if, for every path from root to leaf $x_{1}, \ldots, x_{k}$, there is a function $f \in \mathcal{H}$ such that $f\left(x_{i}\right)$ equals the label of $\left(x_{i}, x_{i+1}\right)$. The Littlestone dimension, L- $\operatorname{dim}(\mathcal{H})$, of $\mathcal{H}$ is the maximal depth of a complete binary tree that is shattered by $\mathcal{H}$.

It is not hard to see that, given a shattered tree of depth $l$, the environment can force any online learning algorithm to make $l$ mistakes. Thus, for any algorithm $A, \mathcal{M}(A) \geq \mathrm{L}-\operatorname{Dim}(\mathcal{H})$. We shall now present an algorithm whose sample complexity is upper bounded by $\mathrm{L}-\operatorname{Dim}(\mathcal{H})$.
Algorithm: Standard Optimal Algorithm (SOA)
Initialization: $V_{0}=\mathcal{H}$.
For $t=1,2 \ldots$,
receive $x_{t}$
for $y \in \mathcal{Y}$, let $V_{t}^{(y)}=\left\{f \in V_{t-1}: f\left(x_{t}\right)=y\right\}$
predict $\hat{y}_{t} \in \arg \max _{y} \operatorname{L}-\operatorname{Dim}\left(V_{t}^{(y)}\right)$
receive true answer $y_{t}$
update $V_{t}=V_{t}^{\left(t_{t}\right)}$

Theorem $20 \mathcal{M}(S O A)=\operatorname{L}-\operatorname{Dim}(\mathcal{H})$.
The proof is a simple adaptation of the proof of the binary case (see Littlestone, 1987). The idea is to note that for each $t$ there is at most one $y \in \mathcal{Y}$ with $\mathrm{L}-\operatorname{Dim}\left(V_{t}^{(y)}\right)=\mathrm{L}-\operatorname{Dim}\left(V_{t}\right)$, and for the rest of the labels we have $\operatorname{L}-\operatorname{Dim}\left(V_{t}^{(y)}\right)<\mathrm{L}-\operatorname{Dim}\left(V_{t}\right)$. Thus, whenever the algorithm errs, the Littlestone dimension of $V_{t}$ decreases by at least 1 , so $\operatorname{after} \operatorname{L}-\operatorname{Dim}(\mathcal{H})$ mistakes, $V_{t}$ is composed of a single function.

Note that we only considered deterministic algorithms. However, allowing the algorithm to make randomized predictions does not substantially improve its sample complexity. It is easy to see that given a shattered tree of depth $l$, the environment can enforce any randomized online learning algorithm to make at least $l / 2$ mistakes on average.

In the agnostic case, the sequence of outcomes, $y_{1}, \ldots, y_{m}$, is not necessarily realizable by some target function $f \in \mathcal{H}$. In that case, our goal is to have a regret of at most $\epsilon$, where the regret is defined as

$$
\frac{1}{m}\left|\left\{t \in[m]: \hat{y}_{t} \neq y_{t}\right\}\right|-\min _{f \in \mathcal{H}} \frac{1}{m}\left|\left\{t \in[m]: f\left(x_{t}\right) \neq y_{t}\right\}\right| .
$$

We denote by $m_{A}^{a}(\epsilon)$ the number of examples required so that the regret of an algorithm $A$ will be at most $\epsilon$ and by $m^{a}(\epsilon)$ the infimum, over all algorithms $A$, of $m_{A}^{a}(\epsilon)$.

Online learnability in the agnostic case, for classes of binary-output functions, has been studied in Ben-David et al. (2009), who showed that the Littlestone dimension characterizes the sample complexity in the agnostic case as well. The basic idea is to construct a set of experts by running the SOA algorithm on all sub-sequences of the examples whose length is at most L-Dim $(\mathcal{H})$, and then to run an online algorithm for learning with experts. This idea can be generalized to the multiclass case, but we leave this generalization to a longer version of this manuscript.

## 4 The Bandit Setting

So far we have assumed that each learning example is comprised of an instance and its corresponding label. In this section we deal with the so-called bandit setting. In the bandit model, the learner does not get to see the correct label of a training example. Instead, the learner first receives an instance $x \in \mathcal{X}$, and should guess a label, $\hat{y}$. The learner then receives a binary feedback, indicating whether its guess is correct or not.

### 4.1 Bandit vs Full Information in the Batch Model

Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a hypothesis class. Our goal is to analyze the realizable bandit sample complexity of $\mathcal{H}$, which we denote by $m_{\mathcal{H}}^{r, b}(\epsilon, \delta)$, and the agnostic bandit sample complexity of $\mathcal{H}$, which we denote by $m_{\mathcal{H}}^{a, b}(\epsilon, \delta)$. The following theorem provides upper bounds on the sample complexity.

Theorem 21 Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a hypothesis class. Then,

$$
m_{\mathcal{H}}^{r, b}(\epsilon, \delta)=O\left(|\mathcal{Y}| \cdot \frac{d_{G}(\mathcal{H}) \cdot \ln \left(\frac{1}{\epsilon}\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right) \text { and } m_{\mathcal{H}}^{a, b}(\epsilon, \delta)=O\left(|\mathcal{Y}| \cdot \frac{d_{G}(\mathcal{H})+\ln \left(\frac{1}{\delta}\right)}{\epsilon^{2}}\right)
$$

Proof: Since the claim is trivial if $|\mathcal{Y}|=\infty$, we can assume that $k:=|\mathcal{Y}|<\infty$. Let $A_{\text {full }}$ be a (full information) ERM learner for $\mathcal{H}$. Consider the following algorithm for the bandit setting: Given a sample $\left(x_{i}, y_{i}\right)_{i=1}^{m}$, for each $i$ the algorithm guesses a label $\hat{y}_{i} \in \mathcal{Y}$ drawn uniformly at random. Then the algorithm returns the hypothesis returned by $A_{\text {full }}$ with the input sample which consists of the pairs $\left(x_{i}, y_{i}\right)$ for which $\hat{y}_{i}=y_{i}$. We claim that $m_{A_{\text {bandit }}}(\epsilon, \delta) \leq 3 k \cdot m_{A_{\text {full }}}\left(\epsilon, \frac{\delta}{2}\right)$ (for both the agnostic and the realizable case), so the theorem is implied by the bounds in the full information setting (theorem 7 and equation 5). Indeed, suppose that $m$ examples suffice for $A_{\text {full }}$ to return, with probability at least $1-\frac{\delta}{2}$ a hypothesis with regret at most $\epsilon$. Let $\left(x_{i}, y_{i}\right)_{i=1}^{3 k m}$ be a sample for the bandit algorithm. By Chernoff bound, with probability at least $1-\frac{\delta}{2}$, the sample $A_{\text {bandit }}$ transfers to $A_{\text {full }}$ consist of at least $m$ examples. Note that the sample that $A_{\text {full }}$ receives is an i.i.d. sample according to the same distribution from which the original sample was sampled. Thus, with probability at least $1-\frac{\delta}{2}, A_{\text {full }}$ (and, consequently, $A_{\text {bandit }}$ ) returns a hypothesis with regret at most $\epsilon$.

The price of bandit information in the batch model: Let $\mathcal{H}$ be a hypotheses class. Define $P B I_{\mathcal{H}}(\epsilon, \delta)=\frac{m_{\mathcal{H}}^{r, b}(\epsilon, \delta)}{m_{\mathcal{H}}^{r}(\epsilon, \delta)}$. By Theorems 21,4 and Equation 3 we see that, $P B I_{\mathcal{H}}(\epsilon, \delta)=$
$O\left(\ln (|\mathcal{Y}|) \cdot \ln \left(\frac{1}{\epsilon}\right) \cdot|\mathcal{Y}|\right)$. This is essentially tight since it is not hard to see that if both $\mathcal{X}, \mathcal{Y}$ are finite and we let $\mathcal{H}=\mathcal{Y}^{\mathcal{X}}$, then $P B I_{\mathcal{H}}=\Omega(|\mathcal{Y}|)$.

Using Theorems 21,4 and Equations 5,3 we see that, as in the full information case, the finiteness of the Natarajan dimension is necessary and sufficient for learnability in the bandit setting as well. However, the ratio between the upper and the lower bounds is $\Omega(\ln (|\mathcal{Y}|) \cdot|\mathcal{Y}|)$. It would be interesting to find a more tight characterization of the sample complexity in the bandit setting. The Natarajan dimension (as well as the graph dimension and other known notions of dimension defined in (Ben-David et al., 1995), as they are all closely related to the Natarajan dimension) is deemed to fail for the following reason: For every $k, d$, there are classes $\mathcal{H} \subseteq[k]^{[d]}$ of Natarajan dimension $d$ where the realizable bandit sample complexity is $O\left(\frac{d}{\epsilon}+\frac{\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)$ (e.g. every class $\mathcal{H}$ such that $d_{N}(\mathcal{H})=d$ and for every $\left.x \in[d], \#\{f(x): f \in \mathcal{H}\}=2\right)$. On the other hand, the realizable bandit sample complexity of $[k]^{[d]}$ is $\Omega\left(k \cdot\left(\frac{d}{\epsilon}+\frac{\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)\right)$.

### 4.2 Bandit vs Full Information in the Online Model

We now consider Bandits in the online learning model. We focus on the realizable case, in which the feedback provided to the learner is consistent with some function $f_{0} \in \mathcal{H}$. We define a new notion of dimension of a class, that determines the sample complexity in this setting. Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a hypothesis class and denote $k=|\mathcal{Y}|$. Consider a rooted tree $T$ whose internal nodes are labeled by $\mathcal{X}$ and such that the labels on edges from a parent to its child nodes are all different from each other. The tree $T$ is $B L$-shattered by $\mathcal{H}$ if, for every path from root to leaf $x_{1}, \ldots, x_{k}$, there is a function $f \in \mathcal{H}$ such that for every $i, f\left(x_{i}\right)$ is different from the label of $\left(x_{i}, x_{i+1}\right)$. The bandit Littlestone dimension of $\mathcal{H}$, denoted $\operatorname{BL}-\operatorname{dim}(\mathcal{H})$, is the maximal depth of a complete $k$-ary tree that is BL-shattered by $\mathcal{H}$.

Theorem 22 Let $\mathcal{H}$ be a hypothesis class with $L=\operatorname{BL}-\operatorname{Dim}(\mathcal{H})$. The sample complexity of every deterministic online learning algorithm for $\mathcal{H}$ is at least $L$. Moreover, there is an online learning algorithm whose sample complexity is exactly $L$.

Proof: First, let $T$ be a BL-shattered tree of depth $L$. We first show that for every deterministic learning algorithm there is a sequence $x_{1}, \ldots, x_{L}$ and a labeling function $f_{0} \in \mathcal{H}$ such that the algorithm makes $L$ mistakes on this sequence. The sequence consists of the instances attached to nodes of $T$, when traversing the tree from the root to one of its leaves, such that the label of each edge $\left(x_{i}, x_{i+1}\right)$ is equal to the algorithm's prediction $\hat{y}_{i}$. The labeling function $f_{0} \in \mathcal{H}$ is one such that for all $i, f_{0}\left(x_{i}\right)$ is different from the label of edge $\left(x_{i}, x_{i+1}\right)$. Such a function exists since $T$ is BL-shattered.

Second, the following online learning algorithm makes at most $L$ mistakes.

```
Algorithm: Bandit Standard Optimal Algorithm (BSOA)
Initialization: }\mp@subsup{V}{0}{}=\mathcal{H}\mathrm{ .
For t=1,2\ldots,
    receive }\mp@subsup{x}{t}{
    for }y\in\mathcal{Y}\mathrm{ , let }\mp@subsup{V}{t}{(y)}={f\in\mp@subsup{V}{t-1}{}:f(\mp@subsup{x}{t}{})\not=y
    predict }\mp@subsup{\hat{y}}{t}{}\in\operatorname{arg}\mp@subsup{\operatorname{min}}{y}{}\operatorname{BL}-\operatorname{Dim}(\mp@subsup{V}{t}{(y)}
    receive an indication whether }\mp@subsup{\hat{y}}{t}{}=f(\mp@subsup{x}{t}{}
    if the prediction is wrong, update }\mp@subsup{V}{t}{}=\mp@subsup{V}{t}{(\mp@subsup{\hat{y}}{t}{})
```

To see that $\mathcal{M}(B S O A) \leq L$, note that at each time $t$, there is at least one $V_{t}^{(y)}$ with $\operatorname{BL}-\operatorname{Dim}\left(V_{t}^{(y)}\right)<$ $\operatorname{BL}-\operatorname{Dim}\left(V_{t-1}\right)$. Thus, whenever the algorithm errs, the dimension of $V_{t}$ decreases by one. Thus, after $L$ mistakes, the dimension is 0 , which means that there is a single function that is consistent with the sample, so no more mistakes can occur.

We conclude with an open question on the price of bandit information in the online model:
Open question 23 Let $P B I(\mathcal{H})=\frac{\operatorname{BL-Dim}(\mathcal{H})}{\operatorname{L-Dim}(\mathcal{H})}$ and fix $k \geq 2$. How large can $P B I(\mathcal{H})$ be when $\mathcal{H}$ is a class of functions from a domain $\mathcal{X}$ to a range $\mathcal{Y}$ of cardinality $k$ ?

## 5 The Sample Complexity of Known Multiclass Hypothesis Classes

In this section we analyze the sample complexity of two families of hypothesis classes for multiclass classification: the generalized linear construction (Duda and Hart, 1973, Vapnik, 1998, Hastie and Tibshirani,

1995, Freund and Schapire, 1997, Schapire and Singer, 1999, Collins, 2002, Taskar et al., 2003), and multiclass reduction trees (Beygelzimer et al., 2007, 2009, Fox, 1997). In particular, a special case of the generalized linear construction is the multi-vector construction (e.g. Crammer and Singer, 2003, Fink et al., 2006). We show that the sample complexity of the multi-vector construction and the reduction trees construction is similar and depends approximately linearly on the number of class labels. Due to the lack of space, proofs are omitted and can be found in the appendix.

### 5.1 The Generalized Linear Multiclass Construction

A generalized linear multiclass hypothesis class is defined with respect to a class specific feature mapping $\phi: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^{t}$, for some integer $t$. For any such $\phi$ define the hypothesis class $\mathcal{M}_{\phi}^{t}=$ $\left\{h[w] \mid w \in \mathbb{R}^{t}\right\}$, where

$$
h[w](x)=\underset{y \in \mathcal{Y}}{\operatorname{argmax}}\langle w, \phi(x, y)\rangle,
$$

where we ignore tie-breaking issues w.l.o.g. . A popular special case is the linear construction used in multiclass SVM (Crammer and Singer, 2003) where $\mathcal{X}=\mathbb{R}^{d}, \mathcal{Y}=[k], t=d k$, and $\phi=\psi_{d, k}$, defined by

$$
\psi_{d, k}(x, i) \triangleq(0, \ldots, 0, x[1], \ldots, x[d], 0, \ldots, 0)
$$

where $x[1]$ is in coordinate $d(i-1)+1$. We abbreviate $\mathcal{L}_{d}^{k} \triangleq \mathcal{M}_{\psi_{d, k}}^{d k}$. We first consider a general $\phi$ and show that the sample complexity for any $\phi$ is upper-bounded by a function of $t$.

Theorem 24 Let $d_{N}$ be the Natarajan-dimension of $\mathcal{M}_{\phi}^{t}$. Then $d_{N} \leq O(t \log (t))$.
For the linear construction a matching lower bound on the Natarajan dimension is shown in the following theorem. Thus, as one might expect, the sample complexity of learning with $\mathcal{L}_{k}^{d}$ is of the order of $d k$.

Theorem 25 For $d \geq 0$ and $k \geq 2$, let $d_{N}$ be the Natarajan-dimension of $\mathcal{L}_{k}^{d}$. Then

$$
\Omega(d k) \leq d_{N} \leq O(d k \log (d k))
$$

### 5.2 Reduction trees

Reduction trees provide a way of constructing multiclass hypotheses from binary classifiers. A reduction tree consists of a tree structure, where each internal node is mapped to a binary classifier and each leaf is mapped to one of the multiclass labels. Classification of an example is done by traversing the tree, starting from the root and ending in one of the leaves, where in each node the result of the binary classifier determines whether to go left or right.

It has been shown that by using appropriate learning algorithms, one can guarantee a multiclass classification error of no more than $\log _{2}(k) \epsilon$, where $k$ is the number of classes, and $\epsilon$ is the average error of the binary classifiers (Fox, 1997, Beygelzimer et al., 2009). However, this result does not directly provide sample complexity guarantees for these algorithms, since the value of $\epsilon$ itself depends on the sample and on the learning algorithm.

In the following we analyze the sample complexity of any fixed reduction tree, under the assumption that the binary classifiers all belong to some fixed hypothesis class with a finite VC-dimension d. We provide bounds on the Natarajan dimension of the resulting multiclass hypothesis class, and show that it can be as large as $\Omega(d k)$ for some hypothesis classes. We further analyze the special case where the binary hypothesis class is the class of linear separators in $\mathbb{R}^{d}$, and show that a similar result, though slightly weaker, holds for this class as well.

We now formally define a reduction tree and the hypothesis class related to it (see Figure 1 in the appendix for illustration). Let $\mathcal{X}$ be the domain of examples and let $[k]$ be the set of possible labels. A reduction tree is a full binary tree $T$. Denote the head node of $T$ by $H(T)$. The sub-tree which is the left child of $H(T)$ is denoted by $L(T)$ and the sub-tree which is the right child of $H(T)$ is denoted by $R(T)$. The set of internal nodes of $T$ is denoted by $N(T)$, and the set of leaf nodes of $T$ is denoted by leaf $(T)$. A multiclass classifier is a triplet $[T, \lambda, C]$ where $T$ is a reduction tree, $\lambda$ is a one-to-one mapping $\lambda[\cdot]: \operatorname{leaf}(T) \rightarrow[k]$, and $C[\cdot]: N(T) \rightarrow\{0,1\}^{\mathcal{X}}$ is a mapping from the internal nodes of $T$ to binary classifiers on the domain $\mathcal{X} .[T, \lambda, C]: \mathcal{X} \rightarrow[k]$ is defined recursively as follows:

$$
[T, \lambda, C](x)= \begin{cases}{[L(T), \lambda, C](x)} & H(T) \notin \operatorname{leaf}(T) \text { and } C[H(T)](x)=0 \\ {[R(T), \lambda, C](x)} & H(T) \notin \operatorname{leaf}(T) \text { and } C[H(T)](x)=1 \\ \lambda[H(T)](x) & H(T) \in \operatorname{leaf}(T)\end{cases}
$$

Unless otherwise mentioned, we assume a fixed $\lambda$, and identify $T$ with the pair $(T, \lambda)$. Accordingly, $[T, \lambda, C]$ is abbreviated to $[T, C]$. Let $\mathcal{H} \subseteq\{0,1\}^{\mathcal{X}}$ be a hypothesis class of binary classifiers on $\mathcal{X}$. The hypothesis class induced by $\mathcal{H}$ on the tree $T$ with label mapping $\lambda$, denoted by $\mathcal{H}_{(T, \lambda)}$, is the set of multiclass classifiers which can be generated on $T$ using binary classifiers from $\mathcal{H}$. Formally,

$$
\mathcal{H}_{(T, \lambda)}=\{[T, \lambda, C] \mid \forall n \in N(T), C[n] \in \mathcal{H}\}
$$

We abbreviate $\mathcal{H}_{(T, \lambda)}$ to $\mathcal{H}_{T}$ when the labeling $\lambda$ is fixed.
Suppose that the VC-dimension of $\mathcal{H}$ is $d$. What can be said about the sample complexity of $\mathcal{H}_{T}$ for a given tree $T$ ? First, a simple counting argument provides an upper bound on the graphdimension and the Natarajan-dimension of $\mathcal{H}_{T}$ : Any hypothesis in $\mathcal{H}_{T}$ is a function of the values of $|N(T)|=k-1$ binary hypotheses from $\mathcal{H}$. Therefore, the number of possible labelings of $A$ by $\mathcal{H}_{T}$ for any $A \subseteq \mathcal{X}$ is bounded by $\left.|\mathcal{H}|_{A}\right|^{k-1}$. By Sauer's lemma, $\left.|\mathcal{H}|_{A}\left|\leq|A|^{d}\right.$. Thus $| \mathcal{H}_{T}\right|_{A}\left|\leq|A|^{d(k-1)}\right.$. If $A$ is G-shattered or N-shattered by $\mathcal{H}_{T}$, then $\left|\mathcal{H}_{T}\right|_{A} \mid \geq 2^{|A|}$. Thus $2^{|A|} \leq|A|^{d(k-1)}$. It follows that $|A| \leq O(d k \log (d k))$, thus the same upper bound holds for the graph-dimension and the Natarajandimension. A closely matching lower bound is provided in the following theorem.

Theorem 26 Let $k \geq 2$ and $d \geq 2$ be integers. For any reduction tree $T$ with $k \geq 2$ leafs, there exists a binary hypothesis class $\mathcal{H}$ with $V C$-dimension $d$ such that $\mathcal{H}_{T}$ has Natarajan dimension $d(k-1)$.

Theorem 26 shows that for every tree there exists a binary hypothesis class which induces a high sample complexity on the resulting multiclass hypothesis class. The following theorem shows that moreover, the popular hypothesis class of linear separators in $\mathbb{R}^{d}$ induces reduction trees with a sample complexity which is almost as large, up to a logarithmic factor.

Let $\mathcal{W}^{d}$ be the class of non-homogeneous linear separators in $\mathbb{R}^{d}$, that is $\mathcal{W}^{d}=\{x \rightarrow \operatorname{sign}(\langle x, w\rangle+$ b) $\left.\mid w \in \mathbb{R}^{d}, b \in \mathbb{R}\right\}$. For a full binary tree $T$ with $k$ leaves, denote by $n_{1}(T)$ the number of internal nodes with one leaf child and one non-leaf child, and by $n_{2}(T)$ the number of internal nodes with two leaf children.

Theorem 27 For any multiclass-to-binary tree $T$ with $k$ leaves, the graph dimension of $\mathcal{W}_{T}^{d}$ is at least $(d+1) \cdot n_{2}(T)+d \cdot n_{1}(T) \geq d k / 2$. Consequently the Natarajan dimension is $\Omega(d k / \log (k))$.

We conclude that the sample complexity of different reduction trees is similar, and that this sample complexity is also similar to that of the multi-vector construction. This implies that when choosing between the different hypothesis classes, considerations other than the sample complexity should determine the choice. One such important consideration is the approximation error. Since sample complexity analysis bounds only the estimation error, one wishes to have the approximation error as low as possible. Thus if there is some prior knowledge on the match between the hypothesis class and the source distribution, this might guide the choice of the hypothesis class. The following theorem shows, however, that for fairly balanced reduction trees this match is highly dependent on the assignment of labels to leaf nodes. For any reduction tree $T$ denote by $\Lambda$ the set of one-to-one mappings from the $\operatorname{leaf}(T)$ to $[k]$, and let $U$ be the uniform distribution over $\Lambda$.

Theorem 28 Let $T$ be a full binary tree with $k$ leaves, and let $n$ be the number of leaves on the left sub-tree. For any hypothesis class $\mathcal{H}$ with VC-dimension $d$, and for any distribution $D$ over $\mathcal{X} \times[k]$ which assigns non-zero probability to each label in $[k]$,

$$
\operatorname{Pr}_{\lambda \sim U}\left[\mathcal{H}_{(T, \lambda)} \text { separates } D\right] \leq\left(\frac{e k}{d}\right)^{d}\binom{k}{n}^{-1}
$$

Thus if $k \gg d$ and $n$ is a constant fraction of $k$, this probability decreases exponentially with $k$.

## 6 Conclusions and Open Problems

In this paper we have studied several new aspects of multiclass sample complexity. Many interesting questions arise and some are listed below.

Consider the two example classes from section 2.4. It is interesting to note that, in both cases, $d_{N}(\mathcal{H})=1$, and $m_{\mathcal{H}}^{r}(\epsilon, \delta)=\Theta\left(\frac{1}{\epsilon} \ln \left(\frac{1}{\delta}\right)\right)$. It seems like the Natarajan dimension is the parameter that controls the sample complexity for those examples. That is also the case for symmetric classes as well as some other classes that we have examined but did not include in this paper. We therefore raise:

Conjecture 29 There exists a constant $C$ such that, for every hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$,

$$
m_{\mathcal{H}}^{r}(\epsilon, \delta) \leq C\left(\frac{d_{N}(\mathcal{H}) \ln \left(\frac{1}{\epsilon}\right)+\ln \left(\frac{1}{\delta}\right)}{\epsilon}\right)
$$

In light of theorem 7 and the fact that there are cases where $d_{G} \geq \log _{2}(|\mathcal{Y}|-1) d_{N}$, in order to prove the conjecture we will have to find a learning algorithm that is not just an arbitrary ERM learner. So far, all the general upper bounds that we are aware of are valid for any ERM learner. Understanding how to select among ERM learners is fundamental as it teaches us what is the correct way to learn. We suspect that such an understanding might lead to improved bounds in the binary case as well. We hope that our examples from section 2.4 and our result for symmetric classes will prove to be the first steps in the search for the best ERM.

Another direction is the study of learnability conditions for additional hypotheses classes. Section 5 shows that some well known multiclass constructions have surprisingly similar sample complexity properties. It is of practical significance and theoretical interest to study learnability conditions for other constructions, and especially to develop a fuller understanding of the relationship between different constructions, in a manner that could guide an informed choice of a hypothesis class.
Acknowledgements: Sivan Sabato is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

## References

P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. Machine learning, 47(2):235-256, 2002.
P. Auer, N. Cesa-Bianchi, Y. Freund, and R.E. Schapire. The nonstochastic multiarmed bandit problem. SICOMP: SIAM Journal on Computing, 32, 2003.
S. Ben-David, N. Cesa-Bianchi, D. Haussler, and P. Long. Characterizations of learnability for classes of $\{0, \ldots, n\}$-valued functions. Journal of Computer and System Sciences, 50:74-86, 1995.
S. Ben-David, D. Pal, , and S. Shalev-Shwartz. Agnostic online learning. In COLT, 2009.
A. Beygelzimer, J. Langford, and P. Ravikumar. Multiclass classification with filter trees. Preprint, June, 2007.

Alina Beygelzimer, John Langford, and Pradeep Ravikumar. Error-correcting tournaments. CoRR, 2009.
M. Collins. Discriminative training methods for hidden Markov models: Theory and experiments with perceptron algorithms. In Conference on Empirical Methods in Natural Language Processing, 2002.
K. Crammer and Y. Singer. Ultraconservative online algorithms for multiclass problems. Journal of Machine Learning Research, 3:951-991, 2003.
R. O. Duda and P. E. Hart. Pattern Classification and Scene Analysis. Wiley, 1973.

Michael Fink, Shai Shalev-Shwartz, Yoram Singer, and Shimon Ullman. Online multiclass learning by interclass hypothesis sharing. In International Conference on Machine Learning, 2006.
J. Fox. Applied Regression Analysis, Linear Models, and Related Methods. SAGE Publications, 1997.
Y. Freund and R.E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. Journal of Computer and System Sciences, 55(1):119-139, August 1997.
T.J. Hastie and R.J. Tibshirani. Generalized additive models. Chapman \& Hall, 1995.
S.M. Kakade, S. Shalev-Shwartz, and A. Tewari. Efficient bandit algorithms for online multiclass prediction. In International Conference on Machine Learning, 2008.
N. Littlestone. Learning when irrelevant attributes abound. In FOCS, pages 68-77, October 1987.
B. K. Natarajan. On learning sets and functions. Mach. Learn., 4:67-97, 1989.
R. E. Schapire and Y. Singer. Improved boosting algorithms using confidence-rated predictions. Machine Learning, 37(3):1-40, 1999.
S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability, stability and uniform convergence. The Journal of Machine Learning Research, 9999:2635-2670, 2010.
B. Taskar, C. Guestrin, and D. Koller. Max-margin markov networks. In NIPS, 2003.
V. N. Vapnik. Statistical Learning Theory. Wiley, 1998.
V.N. Vapnik. The Nature of Statistical Learning Theory. Springer, 1995.

## A Proofs Omitted from the Text

Proof: (of theorem 4)
The lower bound: Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a hypothesis class of Natarajan dimension $d$ and Let $\mathcal{H}_{d}:=$ $\{0,1\}^{[d]}$. We claim that $m_{\mathcal{H}_{d}} \leq m_{\mathcal{H}}$, so the lower bound is obtained by theorem 2. Let $A$ be a learning algorithm for $\mathcal{H}$. Consider the learning algorithm, $\bar{A}$, for $\mathcal{H}_{d}$ defined as follows. Let $S=\left\{s_{1}, \ldots, s_{d}\right\} \subseteq X, f_{0}, f_{1}$ be a set and functions that indicate that $d_{N}(\mathcal{H})=d$. Given a sample $\left(x_{i}, y_{i}\right) \in[d] \times\{0,1\}, i=1, \ldots, m$, let $g=A\left(\left(s_{x_{i}}, f_{y_{i}}\left(s_{x_{i}}\right)\right)_{i=1}^{m}\right)$. Define $f=\bar{A}\left(\left(x_{i}, y_{i}\right)_{i=1}^{m}\right)$ by setting $f(i)=1$ if and only if $g\left(s_{i}\right)=f_{1}\left(s_{i}\right)$. It is not hard to see that $m_{\bar{A}} \leq m_{A}^{r}$, thus, $m_{\mathcal{H}_{d}} \leq m_{\mathcal{H}}$.
The upper bound: Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a hypothesis class of graph dimension $d$. For every $f \in \mathcal{H}$ define $\bar{f}: \mathcal{X} \times \mathcal{Y} \rightarrow\{0,1\}$ by setting $\bar{f}(x, y)=1$ if and only if $f(x)=y$ and let $\overline{\mathcal{H}}=\{\bar{f}: f \in \mathcal{H}\}$. It is not hard to see that $V C(\overline{\mathcal{H}})=d_{G}(\mathcal{H})$.

Suppose that $f \in \mathcal{H}$ is consistent with a sample $\left(x_{i}, f_{0}\left(x_{i}\right)\right)_{i=1}^{m}$ of $m=\Omega\left(\frac{d}{\epsilon} \ln \left(\frac{1}{\epsilon}\right)+\frac{1}{\epsilon} \ln \left(\frac{1}{\delta}\right)\right)$ examples, drawn i.i.d. according to some distribution $\mathcal{D}$ on $\mathcal{X}$. We must show that, with probability $\geq 1-\delta, \operatorname{Err}_{\mathcal{D}, f_{0}}(f) \leq \epsilon$. However, by theorem 2,

$$
\operatorname{Err}_{\mathcal{D}, f_{0}}(f)=\operatorname{Pr}_{x \sim \mathcal{D}}\left(\bar{f}\left(x, f_{0}(x)\right) \neq 1\right) \leq \epsilon
$$

With probability $\geq 1-\delta$.
Proof: (of Theorem 6) Let $A$ be an ERM learner. Since $F_{A}(f) \subseteq \mathcal{H}$ for every $f$, it follows that $\Pi_{A} \leq \Pi_{\mathcal{H}}$. By lemma 13, $\Pi_{\mathcal{H}}(m) \leq m^{d_{N}(\mathcal{H})}|\mathcal{Y}|^{2 d_{N}(\mathcal{H})}$. Incorporating it into Theorem 11 we get the desired bound.
Proof: (of Theorem 24) Let $S=\left\{x_{1}, \ldots, x_{d_{N}}\right\} \subseteq \mathbb{R}^{d}$ be a set which is $N$-shattered by $\mathcal{M}_{\phi}^{t}$, and let $f_{1}, f_{2}: S \rightarrow \mathcal{Y}$ be the functions that witness the shattering. For every $i \in\left[d_{N}\right]$ let $z_{i}=\phi\left(x_{i}, f_{1}\left(x_{i}\right)\right)-$ $\phi\left(x_{i}, f_{2}\left(x_{i}\right)\right) \in \mathbb{R}^{t}$. Denote $Z=\left\{z_{i}\right\}_{i \in\left[d_{N}\right]}$. Consider the hypothesis class of homogeneous linear separators in $\mathbb{R}^{t}$, defined by $\left\{z \rightarrow \operatorname{sign}(\langle w, z\rangle) \mid w \in \mathbb{R}^{t}\right\}$. Since the VC-dimension of this class is $t$, by Sauer's lemma the number of possible labelings of $Z$ with this class is upper-bounded by $\left(d_{N}\right)^{t}$. We now show that there is a one-to-one mapping from subsets $T \subseteq S$ to labelings of $Z$ : For any $T \subseteq S$, let $w \in \mathbb{R}^{t}$ such that

$$
\left\{x \in S \mid h[w](x)=f_{1}(x)\right\}=T, \text { and }\left\{x \in S \mid h[w](x)=f_{2}(x)\right\}=S \backslash T .
$$

Then $T=\left\{x \in S \mid\left\langle w, \phi\left(x, f_{1}(x)\right)\right\rangle \geq\left\langle w, \phi\left(x, f_{2}(x)\right)\right\rangle\right\}=\left\{x_{i} \mid\left\langle w, z_{i}\right\rangle \geq 0\right\}$. Thus every $T$ induces a different labeling of $Z$. It follows that the number of subsets of $S$ is bounded by the number of labelings of $Z$, thus $2^{d_{N}} \leq\left(d_{N}\right)^{t}$. It follows that $d_{N} \leq O(t \log (t))$.
Proof: (of Theorem 25) The upper bound is a direct consequence of Theorem 24. For the lower bound, we show that there exists an N-shattered set of size $\lfloor d / 2\rfloor \cdot\lfloor k / 2\rfloor$. Let $b=\lfloor k / 2\rfloor$. Let $x_{1}, \ldots, x_{b} \in \mathbb{R}^{2}$ be $b$ different vectors such that $\forall i \in[b],\left\|x_{i}\right\|=1$. Let $S=\left\{y_{i, j}\right\}_{i \in[b], j \in[[d / 2]]} \subseteq \mathbb{R}^{d}$, where for $s \in[d]$ :

$$
y_{i, j}[s]= \begin{cases}x_{i}[1] & s=2 j-1 \\ x_{i}[2] & s=2 j \\ 0 & \text { otherwise } .\end{cases}
$$

We show that $S$ is $N$-shattered, thus $d_{N} \geq|S|=\lfloor k / 2\rfloor \cdot\lfloor d / 2\rfloor$. Define functions $f_{1}, f_{2}: S \rightarrow[k]$ such that for $y_{i, j} \in S, f_{1}\left(y_{i, j}\right)=i$ and $f_{2}\left(y_{i, j}\right)=b+i$. For a subset $T \subseteq S$, let $w \in \mathbb{R}^{d k}$ such that for $i \in[b], s \in[d]$

$$
w[d(i-1)+s]= \begin{cases}x_{i}[1] & y_{i, j} \in Z \text { and } s=2 j-1, \\ x_{i}[2] & y_{i, j} \in Z \text { and } s=2 j, \\ 0 & \text { otherwise }\end{cases}
$$

and for $i \in\{b+1, \ldots, 2 b\}, s \in[d]$,

$$
w[d(i-1)+s]= \begin{cases}x_{i}[1] & y_{i-b, j} \notin Z \text { and } s=2 j-1, \\ x_{i}[2] & y_{i-b, j} \notin Z \text { and } s=2 j, \\ 0 & \text { otherwise. }\end{cases}
$$

Then $h[w]=f_{1}(y)$ for $y \in T$ and $h[w]=f_{2}(y)$ for $y \in S \backslash T$. Thus $S$ is N -shattered.
Proof: (of Theorem 26) Let $\mathcal{H}(T)$ be a binary hypothesis class for tree $T$. We construct $\mathcal{H}(T)$ inductively on the structure of the tree. For every tree $T$, the domain of the binary hypotheses in $\mathcal{H}(T)$ will be $[d] \times N(T)$.


Figure 1: Illustration of a reduction tree

Induction basis: Assume that both $L(T)$ and $R(T)$ are leafs, thus $k=2$ and $|N(T)|=1$. Define $\mathcal{H}(T)=\{h \mid h:[d] \times\{H(T)\} \rightarrow\{0,1\}\}$.

Inductive step: Assume $T$ has two children $L(T)$ and $R(T)$, and at least one of them is not a leaf. By the induction hypothesis, if $L(T)$ is a non-leaf then $\mathcal{H}(L(T))$ is a set of binary hypotheses with domain $[d] \times N(L(T))$. $\mathcal{H}(L(T))$ has VC-dimension $d$, and the Natarajan dimension of $\mathcal{H}(L(T))_{L(T)}$ is $d \cdot|N(L(T))|$. The same holds for $R(T)$. Define $\mathcal{H}(T)=\left\{h_{0}, h_{1}\right\} \cup \mathcal{H}_{L} \cup \mathcal{H}_{R} \cup \mathcal{H}_{H}$, where:

- $h_{0}(x)=0$ and $h_{1}(x)=1$ for all $x \in[d] \times N(T)$,
- If $L(T)$ is a leaf, $\mathcal{H}_{L}=\emptyset$. Otherwise,

$$
\begin{aligned}
& \mathcal{H}_{L}=\left\{h:[d] \times N(T) \rightarrow\{0,1\} \mid \exists h_{L} \in \mathcal{H}(L(T)),\right. \forall x \in[d] \times N(T) \\
& h(x)= \begin{cases}h_{L}(x) & x \in[d] \times N(L(T)) \\
0 & \text { otherwise }\end{cases}
\end{aligned}
$$

- $\mathcal{H}_{R}$ is defined similarly, for $R(T)$ instead of $L(T)$.
- $\mathcal{H}_{H}$ is defined as follows:

$$
\begin{aligned}
\mathcal{H}_{H}=\{h:[d] \times N(T) \rightarrow\{0,1\} & \mid \forall x \in[d] \times N(L(T)), h(x)=0 \\
& \forall x \in[d] \times N(R(T)), h(x)=1\}
\end{aligned}
$$

We now prove by induction that for every tree $T$ the following claims hold:

- $\mathcal{H}(T)$ has VC-dimension $d$,
- $\mathcal{H}(T)_{T}$ has Natarajan dimension $d \cdot|N(T)|$.
- An auxiliary claim: $\mathcal{H}(T)$ includes the hypotheses $h_{0}$ and $h_{1}$.

Induction Basis: If both $L(T)$ and $R(T)$ are leafs, then the VC-dimension of $\mathcal{H}(T)$ is clearly d. The induced multiclass hypothesis class $\mathcal{H}(T)_{T}$ is in fact a set of binary hypotheses which is isomorphic to $\mathcal{H}(T)$, thus its Natarajan dimension is also $d=d(k-1)$. The zero hypothesis is clearly in $\mathcal{H}(T)$ by construction.

Induction Step: Assume $T$ has two children $L(T)$ and $R(T)$, and at least one of them is not a leaf. By the construction of $\mathcal{H}(T)$, the auxiliary claim clearly holds. The following lemmas, whose proofs follows, prove the two other claims:

Lemma $30 \mathcal{H}(T)$ has VC-dimension d.
Lemma $31 \mathcal{H}(T)_{T}$ has Natarajan dimension $d|N(T)|$.
Thus the induction hypothesis holds.
Proof: (of Lemma 30) The VC-dimension of $\mathcal{H}(T)$ is at least $d$, since the VC-dimension of at least one of $\mathcal{H}(L(T))$ and $\mathcal{H}(R(T))$ is $d$. Assume to the contrary that it is larger than $d$, then there exists a set $A=\left\{x_{1}, \ldots, x_{d+1}\right\} \subseteq[d] \times N(T)$ which is shattered by $\mathcal{H}(T)$. Denote for brevity $S_{L}=[d] \times N(L(T)), S_{R}=[d] \times N(R(T))$ and $S_{H}=[d] \times H(T)$. By the construction of $\mathcal{H}(T)$ and the auxiliary claim, $\left.\mathcal{H}(T)\right|_{S_{L}}=\mathcal{H}(L(T))$ and $\left.\mathcal{H}(T)\right|_{S_{R}}=\mathcal{H}(R(T))$ whenever $L(T)$ and $R(T)$ are not leaves respectively. In addition, since $\left|S_{H}\right|=d, A \nsubseteq S_{H}$. Since $|A| \geq 3$, there exist three different elements in $x, y, z \in A$ such that at least two of them are in different sets out of $S_{L}, S_{H}, S_{R}$. We consider the different cases (where names of elements are w.l.o.g.) and show for each case a labeling $l_{x}, l_{y}, l_{z}$ for $x, y, z$ that cannot be achieved with a hypothesis in $\mathcal{H}(T)$ :

- If $x \in S_{H}, y \in S_{R}$ then $l_{x}=1, l_{y}=0$ cannot be achieved.
- If $x, y \in S_{L}, z \in S_{H} \cup S_{R}$ then $l_{x}=1, l_{y}=0, l_{z}=1$ cannot be achieved.
- If $x \in S_{L}, y, z \in S_{R}$ then $l_{x}=1, l_{y}=0, l_{z}=1$ cannot be achieved.
- If $x \in S_{L}, y, z \in S_{H}$ then $l_{x}=1, l_{y}=0, l_{z}=1$ cannot be achieved.

We have reached a contradiction, therefore no such $A$ exists.
Proof: (of Lemma 31) The Natarajan dimension is upper bounded by the size of the domain, which is $d|N(T)|$. By the induction hypothesis, $\mathcal{H}(L(T))_{L(T)}$ and $\mathcal{H}(R(T))_{R(T)}$ have Natarajan dimension $d_{L}=d|L(T)|$ and $d_{R}=d|R(T)|$ respectively. Thus $[d] \times N(L(T))$ and $[d] \times N(R(T))$ are N -shattered by $\mathcal{H}(L(T))_{L(T)}$ and $\mathcal{H}(R(T))_{R(T)}$ respectively. Let $f_{1}^{L}, f_{2}^{L}$, and $f_{1}^{R}, f_{2}^{R}$ be the pairs of functions that witness the N -shattering of $\mathcal{H}(L(T))_{L(T)}$ and $\mathcal{H}(R(T))_{R(T)}$ respectively. Let $c_{L}$ be the class of the left-most child in $L(T)$, and let $c_{R}$ be the class of the left-most child in $R(T)$. define $g_{1}$ and $g_{2}$ as follows:

$$
\begin{aligned}
& g_{1}(x)= \begin{cases}f_{1}^{L}(x) & x \in[d] \times N(L(T)) \\
f_{1}^{R}(x) & x \in[d] \times N(R(T)) \\
c_{L} & x \in[d] \times\{H(T)\}\end{cases} \\
& g_{2}(x)= \begin{cases}f_{2}^{L}(x) & x \in[d] \times N(L(T)) \\
f_{2}^{R}(x) & x \in[d] \times N(R(T)) \\
c_{R} & x \in[d] \times\{H(T)\}\end{cases}
\end{aligned}
$$

It is easy to verify that $[d] \times N(T)$ is N -shattered using $g_{1}$ and $g_{2}$.
Proof: (of Theorem 27) The proof is by induction on the structure of the tree.
Induction basis: Assume that $T$ is a tree with one internal node and two leaf children. Then $\mathcal{W}_{T}^{d}$ is isomorphic up to label names to $\mathcal{W}^{d}$. Thus the graph dimension of $\mathcal{W}_{T}^{d}$ is equal to the VC-dimension of $\mathcal{W}^{d}$, that is $d+1=(d+1) \cdot n_{1}(T)$.

Inductive step: We consider two cases: Either both $R(T)$ and $L(T)$ are non-leaves or ons is a leaf and one is not.

Case 1: Let $T$ be a tree where both $L(T)$ and $R(T)$ are non-leaves. By the induction hypothesis, the graph dimension of $\mathcal{W}_{L(T)}^{d}$ is at least $d_{L}=(d+1) \cdot n_{2}(L(T))+d \cdot n_{1}(L(T))$ and the graph dimension of $\mathcal{W}_{R(T)}^{d}$ is at least $d_{R}=(d+1) \cdot n_{2}(R(T))+d \cdot n_{1}(R(T))$. Thus there exist sets $A_{L}=\left\{a_{1}, \ldots, a_{d_{L}}\right\}$ and $B_{R}=\left\{b_{1}, \ldots, b_{d_{R}}\right\}$ which are G-shattered by $L(T)$ and $R(T)$ respectively, using functions $f_{L}$ and $f_{R}$ respectively. Let

$$
\begin{aligned}
a_{L} & =\left(\min _{i \in\left[d_{L}\right]}\left\{a_{i}[1]\right\}+1,0, \ldots, 0\right) \in \mathbb{R}^{d} \\
b_{R} & =\left(-\max _{i \in\left[d_{R}\right]}\left\{b_{i}[1]\right\}-1,0, \ldots, 0\right) \in \mathbb{R}^{d}
\end{aligned}
$$

Let $\tilde{A}_{L}=\left\{a_{1}+a_{L}, \ldots, a_{d_{L}}+a_{L}\right\}$ and let $\tilde{B}_{R}=\left\{b_{1}+b_{R}, \ldots, b_{d_{L}}+b_{R}\right\}$. Then $\forall x \in \tilde{A}_{L}, x[1]>0$, and $\forall x \in \tilde{B}_{R}, x[1]<0$.

We show that the set $\tilde{A}_{L} \cup \tilde{B}_{R}$ is G-shattered by $\mathcal{W}_{T}^{d}$ : Define

$$
f(x)= \begin{cases}f_{R}(x) & x[1]>0 \\ f_{L}(x) & \text { otherwise }\end{cases}
$$

Let $Z \subseteq \tilde{A}_{L} \cup \tilde{B}_{R}$. We construct a mapping $C: N(T) \rightarrow \mathcal{H}$ such that

$$
\left\{x \in \tilde{A}_{L} \cup \tilde{B}_{R} \mid[T, C](x)=f(x)\right\}=Z
$$

Let $Y \subseteq A_{L} \cap B_{R}=\left\{a_{i} \mid a_{i}+a_{L} \in Z\right\} \cap\left\{b_{i} \mid b_{i}+b_{R} \in Z\right\}$. Since $A_{L}$ and $B_{R}$ are G-shattered with $f_{L}$ and $f_{R}$, there exist mappings $C_{L}: N(L(T)) \rightarrow \mathcal{W}^{d}$ and $C_{R}: N(R(T)) \rightarrow \mathcal{W}^{d}$ such that

$$
\begin{aligned}
& \left\{x \in A_{L} \mid\left[L(T), C_{L}\right](x)=f_{L}(x)\right\}=Y \cap A_{L} \\
& \left\{x \in B_{R} \mid\left[R(T), C_{R}\right](x)=f_{R}(x)\right\}=Y \cap B_{R}
\end{aligned}
$$

Define the mapping $C$ as a translation of the mappings $C_{L}$ and $C_{R}$, defined by:

$$
\begin{aligned}
& \forall n \in L(T), C_{L}[n]=(w, b) \Rightarrow C[n]=\left(w, b-\left\langle w, a_{L}\right\rangle\right) \\
& \forall n \in R(T), C_{R}[n]=(w, b) \Rightarrow C[n]=\left(w, b-\left\langle w, b_{R}\right\rangle\right)
\end{aligned}
$$

Then

$$
\begin{aligned}
& \left\{x \in \tilde{A}_{L} \mid[L(T), C](x)=f_{L}(x)\right\}=Z \cap \tilde{A}_{L} \\
& \left\{x \in \tilde{B}_{R} \mid[R(T), C](x)=f_{R}(x)\right\}=Z \cap \tilde{B}_{R}
\end{aligned}
$$

Now, set $C[H(T)](x)=\operatorname{sign}(\langle x, w\rangle+b)$ where $w=(1,0, \ldots, 0)$ and $b=0$. Then

$$
\begin{aligned}
& \forall x \in \tilde{A}_{L},[T, C](x)=[L(T), C](x)=f_{L}(x)=f(x) \\
& \forall x \in \tilde{B}_{R},[T, C](x)=[R(T), C](x)=f_{R}(x)=f(x)
\end{aligned}
$$

Thus $\tilde{A}_{L} \cup \tilde{B}_{R}$ is G-shattered by $\mathcal{W}_{T}^{d}$. It follows that the graph dimension of $\mathcal{W}_{T}^{d}$ is at least $\left|\tilde{A}_{L} \cup \tilde{B}_{R}\right|=d_{L}+d_{R}=(d+1) \cdot n_{2}(T)+d \cdot n_{1}(T)$.

Case 2: Assume w.l.o.g. that $T$ is a tree where $L(T)$ is not a leaf node and $R(T)$ is a leaf node with $\lambda[R(T)]=t$. By the induction hypothesis, the graph dimension of $\mathcal{W}_{L(T)}^{d}$ is at least $d_{L}=(d+1) \cdot n_{2}(L(T))+d \cdot n_{1}(L(T))$. Thus there exists a set $A=\left\{a_{1}, \ldots, a_{d_{L}}\right\}$ which is Gshattered by $L(T)$ using the function $f_{L}$.

Denote by $e_{i}$ the $i$ 'th unit vector in $\mathbb{R}^{d}$, and let $q>0$ be large enough such that $\left\{(0, \ldots, 0), q e_{1}, \ldots, q e_{d}\right\}$ is shattered with a margin of $2 M$, where $M=\max _{x \in A}\|x\|_{2}$. Let $B=A \cup\left\{q e_{1}, \ldots, q e_{d}\right\}$. Then we show $B$ is G-shattered using the following function $f$ :

$$
f(x)= \begin{cases}f_{L} & \|x\| \leq q \\ t & \text { otherwise }\end{cases}
$$

Let $Z \subseteq B$. We construct a mapping $C: N(T) \rightarrow \mathcal{H}$ such that

$$
\begin{equation*}
\{x \in B \mid[T, C](x)=f(x)\}=Z \tag{8}
\end{equation*}
$$

Since $A$ is G-shattered using $f_{L}$, there exists a mapping $C_{L}: N(L(T)) \rightarrow \mathcal{W}^{d}$ such that $\{x \in A \mid$ $\left.\left[L(T), C_{L}\right](x)=f_{L}(x)\right\}=Z \cap A$. Define $C$ such that $\forall n \in N(L(T)), C[n]=C_{L}[n]$. In addition, Let $C[H(T)] \in \mathcal{W}^{d}$ be a hypothesis such that $\forall i, e_{i} \in Z \Longleftrightarrow h\left(e_{i}\right)=1$, and $\forall x,\|x\|_{2} \leq M \rightarrow h(0)=0$. Then Equation. (8) holds. Thus the graph dimension of $\mathcal{W}_{T}^{d}$ is at least $|B|=d_{L}+d \geq(d+1)$. $n_{2}(T)+d \cdot n_{1}(T)$.
Proof: (of Theorem 28) If suffices to consider distributions with deterministic labeling, such that the correct label is a function $f: \mathcal{X} \rightarrow[k]$. Let $A=\left\{x_{1}, \ldots, x_{k}\right\} \in \mathcal{X}$ such that for all $i \in[k]$, $f\left(x_{i}\right)=i$. For any labeling $\lambda \in \Lambda$, let $f_{\lambda}: A \rightarrow\{0,1\}$ be the indicator function of the set of labels assigned to leaves in $L(T)$, that is $f_{\lambda}\left(x_{i}\right)=\overrightarrow{1}[\exists n \in \operatorname{leaf}(L(T)), \lambda[n]=i]$. If $D$ is separable with $\mathcal{H}_{(T, \lambda)}$ then $f_{\lambda}=\left.\left.C[H((T, \lambda))]\right|_{A} \in \mathcal{H}\right|_{A}$. By Sauer's lemma, $|\mathcal{H}|_{A} \left\lvert\, \leq\left(\frac{e k}{d}\right)^{d}\right.$. There are $\binom{k}{n}$ possible indicator functions $f_{\lambda}$ for a labeling $\lambda$, and they all have equal probability for $\lambda \sim U$. Thus $\mathbb{P}_{\lambda \sim U}\left[\left.f_{\lambda} \in \mathcal{H}\right|_{A}\right] \leq\left(\frac{e k}{d}\right)^{d} /\binom{k}{n}$.


[^0]:    *Dept. of Mathematics, The Hebrew University, Jerusalem, Israel
    ${ }^{\dagger}$ School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel
    ${ }^{\ddagger}$ David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada
    ${ }^{\S}$ School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel

[^1]:    ${ }^{1}$ Note that Shalev-Shwartz et al. (2010) established gaps between ERM learners in the general learning setting. However, here we consider multiclass learning, which seems very similar to binary classification.

