## A Omitted Details from the Main Body

## A. 1 Previous Work

Classification with multiple labels is extensively studied and, for the setting with bounded $k$, PAC learnability is well-understood [Nat89, BCHL95, SSBD14, DSBDSS11]. The works of [VC71, BEHW89] provide a fundamental dichotomy/equivalence between the finiteness of the VC dimension and binary classification (PAC learnability). [NT88] and [Nat89] extend the PAC framework to the multiclass setting by providing the notions of the Graph and the Natarajan dimension. When the number of labels $k$ is a finite constant, then these two dimensions both characterize PAC learnability since $\operatorname{Ndim}(\mathcal{H}) \leq \operatorname{Gdim}(\mathcal{H}) \leq \operatorname{Ndim}(\mathcal{H}) \cdot O(\log (k))$ [DSBDSS11]. Afterwards, [BCHL95] and [HL95] provide a combinatorial abstraction which captures as special cases e.g., the Graph and Natarajan dimensions and Pollard's pseudo-dimension. In this general setting, [BCHL95] identify the notion of $\Psi$-distinguishers that characterize PAC learnability when $k$ is bounded. More to that, uniform convergence still applies when the number of classes is bounded and, so, PAC learning provides the ERM principle for algorithm design. We remark the situation gets much more complicated when the number of labels is not bounded [DSS14, BCD ${ }^{+} 22$, RBR09, DSBDSS11]. For instance, [DSS14] show that the ERM principle does not apply in this case. In a recent breakthrough, $\left[\mathrm{BCD}^{+} 22\right]$ show that the DS dimension captures learnability and the Natarajan dimension provably fails to achieve this.

## A. 2 Assumption on Cardinality of $\mathcal{H}$

We briefly discuss why the assumption that $|\mathcal{H}|>k+2$ comes without loss of generality. If $\mathcal{H}$ contains either a single hypothesis or contains some hypotheses which have no conflict, i.e., $h(x) \neq h^{\prime}(x)$ everywhere, then $\operatorname{Pr}_{(x, y) \sim P}\left[\widehat{h}_{n}(x) \neq y\right]=0$ is trivially achievable for all $n$. Now, if $|\mathcal{H}| \leq k+2$ but does not fall in the above two cases, then $\mathcal{H}$ is learnable at an optimal exponential rate. To see this, let $\epsilon$ be the minimal error among all hypotheses $h \in \mathcal{H}$ with non-zero error. The probability that there exists a hypothesis with error $\epsilon$ that makes no mistakes in the $n$ training data is at most $|\mathcal{H}|(1-\epsilon)^{n}$ using the union bound. Thus a learning algorithm that outputs any classifier $\widehat{h}_{n} \in \mathcal{H}$ that correctly classifies the training set satisfies $\mathbf{E}\left[\operatorname{Pr}_{(x, y) \sim P}\left[\widehat{h}_{n}(x) \neq y\right]\right] \leq C \exp (-c n)$ with $C=C(\mathcal{H}, P)$ and $c=c(\mathcal{H}, P)$. Moreover, this is optimal due to Proposition 1.

## A. 3 Future Directions

We deal with the settings of universal learning and partial concept classes, two fundamental questions $\left[\mathrm{BHM}^{+} 21\right.$, AHHM22] which are witnessed in real-life applications, nevertheless the classical theory fails to explain. Our results raise various interesting questions for future work (apart from Open Question 1). First, it would be interesting to extend our results to the agnostic setting. Second, for the universal setting, we believe it is an important next step to shed light towards multiclass learning with unbounded labels (whose uniform learnability was recently characterized by $\left[\mathrm{BCD}^{+} 22\right]$ ). Moreover, for the partial concepts setting, the work of [AHHM22] leaves numerous fascinating open questions for the binary setting that can be asked in the multiclass setting too. In general, our work along with its seminal binary counterparts [ $\mathrm{BHM}^{+} 21$, AHHM22] shows that the algorithmic landscape occuring in practice is quite diverse and the ERM principle is provably insufficient. It is important to come up with principled algorithmic strategies that bring theory closer to practice.

## B Preliminaries

In this section we discuss more extensively the Preliminaries from Section 1.4.

## B. 1 Complexity Measures

We first state the classical definition of the Littlestone-dimension [Lit88] that characterizes learnability in the online setting.
Definition 5 (Littlestone dimension). Consider a complete binary tree $T$ of depth $d+1$ whose internal nodes are labeled by points in $\mathcal{X}$ and edges by $\{0,1\}$, when they connect the parent to the
right, left child, respectively. We say that $\mathcal{H} \subseteq\{0,1\}^{\mathcal{X}}$ Littlestone-shatters $T$ if for every root-to-leaf path $x_{1}, y_{1}, x_{2}, y_{2}, \ldots, x_{d}, y_{d}, x_{d+1}$ there exists some $h \in \mathcal{H}$ such that $h\left(x_{i}\right)=y_{i}, 1 \leq i \leq d$. The Littlestone dimension is denoted by $\operatorname{Ldim}(\mathcal{H})$ is defined to be the largest $d$ such that $\mathcal{H}$ Littlestoneshatters such a binary tree of depth $d+1$. If this happens for every $d \in \mathbb{N}$ we say that $\operatorname{Ldim}(\mathcal{H})=\infty$.

The above definition naturally extends to multiple labels [DSBDSS11]; we denote the multiclass Littlestone dimension by $\operatorname{Ldim}_{k}(\cdot)$. We next recall the well-known notion of VC dimension that characterizes PAC-learnability of binary concept classes [VC71].
Definition 6 (VC dimension). We say that $\mathcal{H} \subseteq\{0,1\}^{\mathcal{X}}$ VC-shatters a sequence $\left\{x_{1}, \ldots, x_{n}\right\} \in \mathcal{X}^{n}$ if $\left\{h\left(x_{1}\right), \ldots, h\left(x_{n}\right): h \in \mathcal{H}\right\}=\{0,1\}^{n}$. The VC-dimension of $S$ is denoted by $\operatorname{VCdim}(\mathcal{H})$ and is defined to be the largest $d$ such that $\mathcal{H} V C$-shatters some sequence of length $d$. If this happens for every $d \in \mathbb{N}$ we say that $\operatorname{VCdim}(\mathcal{H})=\infty$.

When moving to $\mathcal{H} \subseteq\{0, \ldots, k\}^{\mathcal{X}}, k>1$, there are many different extensions of the VC-dimension that have been considered in the literature. We recall the definition of two of them that are important in our work, the Natarajan dimension [Nat89] and the Graph dimension [NT88, Nat89].
Definition 7 (Natarajan dimension). We say that $\mathcal{H} \subseteq\{0, \ldots, k\}^{\mathcal{X}} N$-shatters a sequence $\left\{x_{1}, \ldots, x_{n}\right\} \in \mathcal{X}^{n}$ if there exist two colorings $s^{(0)}, s^{(1)}$ such that $s^{(0)}\left(x_{i}\right) \neq s^{(1)}\left(x_{i}\right), \forall i \in[n]$, and $\forall I \subseteq[n]$ there exists some $h \in \mathcal{H}$ with $h\left(x_{i}\right)=s^{(0)}\left(x_{i}\right), i \in I$ and $h\left(x_{i}\right)=s^{(1)}\left(x_{i}\right), i \in[n] \backslash I$. The Natarajan dimension is denoted by $\operatorname{Ndim}(\mathcal{H})$ and is defined to be the largest $d$ for which $\mathcal{H}$ $N$-shatters a sequence of length $d$. If this happens for every $d \in \mathbb{N}$ we say that $\operatorname{Ndim}(\mathcal{H})=\infty$.
Definition 8 (Graph dimension). We say that $\mathcal{H} \subseteq\{0, \ldots, k\}^{\mathcal{X}} G$-shatters a sequence $\left\{x_{1}, \ldots, x_{n}\right\} \in \mathcal{X}^{n}$ if there exists a coloring $s^{(0)}$ such that $\forall I \subseteq[n]$ there exists some $h \in \mathcal{H}$ with $h\left(x_{i}\right)=s^{(0)}\left(x_{i}\right), i \in I$ and $h\left(x_{i}\right) \neq s^{(0)}\left(x_{i}\right), i \in[n] \backslash I$. The Graph dimension is denoted by $\operatorname{Gdim}(\mathcal{H})$ and is defined to be the largest $d$ for which $\mathcal{H} G$-shatters a sequence of length $d$. If this happens for every $d \in \mathbb{N}$ we say that $\operatorname{Gim}(\mathcal{H})=\infty$.

Notice that if $k=2$ the definitions of the Natarajan dimension and the Graph dimension are equivalent to the definition of the VC-dimension.
Remark 2. We underline that the above definitions can be naturally extended to the partial concepts setting. For instance, we say that $\mathcal{H} \subseteq\{0,1, \star\}^{\mathcal{X}}$ VC-shatters a set $S$ if every binary pattern is realized by some $h \in \mathcal{H}$. For further details, see [AHHM22].

## B. 2 One-Inclusion Hypergraph Algorithm

We next review a fundamental result which is a crucial ingredient in the design of our algorithms, namely the one-inclusion hypergraph algorithm $\mathbb{A}_{\mathcal{H}}$ for the class $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ [HLW94, RBR09, $\mathrm{DSS} 14, \mathrm{BCD}^{+} 22$ ]. This algorithm gets as input a training set $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)$ realizable by $\mathcal{H}$ and an additional example $x$. The goal is to predict the label of $x$. In this sense, the one-inclusion graph constitutes a transductive model in machine learning. The idea is to construct the one-inclusion (hyper)graph of $\left.\mathcal{H}\right|_{x_{1}, \ldots, x_{n}, x} \subseteq[k]^{n+1}$. The nodes of this graph are the equivalence classes of $\mathcal{H}$ induced by the examples $x_{1}, \ldots, x_{n}, x$. For the binary classification case, two equivalence classes are connected with an edge if the nodes differ by exactly one element $x$ of the $n+1$ points and $x$ appears only once in $\left(x_{1}, \ldots, x_{n}, x\right)$. For the case $k>1$, the hyperedge set is is generalized accordingly. Having created the one-inclusion graph, the goal is to orient the edges; the crucial property is that good orientations of this graph yield low error learning algorithms. Here, an orientation is good if the maximum out-degree of the graph is small. Intuitively, if the maximum out-degree of any node is $M$, then this can yield a prediction for the label $x$ with $n+1$ training samples with permutation mistake bound at most $M /(n+1)$.
Lemma 1 (One-Inclusion Hypergraph Algorithm (See Lemma 17 of [ $\left.\left.\mathrm{BCD}^{+} 22\right]\right)$ ). Let $\mathcal{X}=[n]$ with $n \in \mathbb{N}$, $k$ be positive constant and $\mathcal{H} \subseteq[k]^{n}$ be a class with Natarajan dimension $\operatorname{Ndim}(\mathcal{H})<\infty$. There exists an algorithm $\mathbb{A}: 2^{[k]^{X}} \times(X \times[k])^{n-1} \times X$ such that
$\frac{1}{n!} \sum_{\sigma \sim \mathcal{U}\left(\mathbb{S}_{n}\right)}[\mathbb{A}(\mathcal{H}, \sigma(1), h(\sigma(1)), \ldots, \sigma(n-1), h(\sigma(n-1)), \sigma(n)) \neq h(\sigma(n))] \leq \frac{\mathrm{Ndim}(\mathcal{H}) \log (k)}{n}$,
for any $h \in \mathcal{H}$.

## B. 3 Ordinals

The discussion of ordinals is borrowed from $\left[\mathrm{BHM}^{+} 21\right]$. For a thorough treatment of the subject, the interested reader is referred to [HJ99, Sie58].
We consider some set $S$. A well ordering of $S$ is defined to be any linear ordering $<$ so that every non-empty subset of $S$ contains a least element. For example, the set of natural numbers $\mathbb{N}$ along with the usual ordering is well-ordered, whereas $[0,1]$ is not (take, e.g., $S=(0,1)$, then it is clear that there is not a least element in $S$.)

We say that two-well ordered sets are isomorphic if there exists a bijection between the two which preserves the ordering. There is a canonical way to construct "equivalence classes" of well-ordered sets, called ordinals, so that every well-ordered set is isomorphic to exactly one such ordinal. In that sense, ordinals uniquely encode well-ordered sets in the same way as cardinals uniquely encode sets. We denote the class of all ordinals by ORD.
An important property of ordinals is that every pair of well-ordered sets is either isomorphic, or one of them is isomorphic to an initial segment of the other. This fact induces an ordering over the ordinals. To be more precise, we say that for two ordinals $\alpha, \beta \in$ ORD we have that $\alpha<\beta$ if $\alpha$ is isomorphic to an initial segment of $\beta$. The defining property of ordinals is that every ordinal $\beta$ is isomorphic to the set of ordinals that precede it, i.e., $\{\alpha: \alpha<\beta\}$. Moreover, the ordering $<$ is itself a well-ordering over the ordinals. This is because every non-empty set of ordinals contains a least element and it has a least upper bound.
Given the above discussion, we can see that ordinals provide a set-theoretic extension of natural numbers. This is because every ordinal $\beta$ has a successor $\beta+1$, which is the smallest ordinal that is larger that $\beta$. Thus, we can create a list of all ordinals: the first elements in the least are $0,1,2, \ldots$, which are all the finite ordinals where we identify each number $k$ with the well-ordered set $\{0, \ldots, k-1\}$. We let the "smallest" infinite ordinal be $\omega$. This can be identified with the family of all natural numbers along with their natural ordering. The way to count past infinity is the following: we write $0,1,2 \ldots, \omega, \omega+1, \omega+2, \ldots, \omega+\omega, \ldots$, and we denote the smallest uncountable ordinal with $\omega_{1}$.
A concept that is defined by ordinals and is useful for proving the guarantees of our algorithms in the Gale-Stewart games is that of transfinite recursion. Roughly speaking, this principle states that if we have a "recipe", which is given sets of "objects" $O_{\alpha}$, indexed by all ordinals $\alpha<\beta$, defines a new set of "objects" $O_{\beta}$ and has access to a "base set" $\left\{O_{\alpha}: \alpha<\alpha_{0}\right\}$, then $O_{\beta}$ is uniquely defined for all ordinals $\beta$. To give an example, this concept helps us define addition $\gamma+\beta$ between two ordinals. We set $\gamma+0=\gamma$ and $\gamma+1$ to be the successor of $\gamma$. We continue inductively and define for any $\beta$ the addition $\gamma+\beta=\sup \{(\gamma+\alpha)+1: \alpha<\beta\}$. Following this principle we can define an arithmetic in the ordinals.

## B. 4 Well-founded Relations and Ranks

We continue the above discussion by extending the notion of a well-ordering to a more general type of orders, and we introduce the notion of rank, which is important for the derivation of the winning strategies in the Gale-Stewart games. We follow the presentation from $\left[\mathrm{BHM}^{+} 21\right]$. The classical reference for this topic is [Kec 12].
We define a relation $\prec$ on a set $S$ by an arbitrary subset $R_{\prec} \subseteq S \times S$ and we let $x \prec y$ if and only if $(x, y) \in R_{\prec}$. An element $x$ of $(S, \prec)$ is called minimal if there is no $x^{\prime} \in S$ with $x^{\prime} \prec x$. The relation is called well-founded if every non-empty subset of $S$ has a minimal element.
We associate every well-founded relation on $S$ with a rank function $\rho_{\prec}: S \rightarrow$ ORD, which is defined by the following transfinite recursion: we let $\rho_{\prec}(x)=0$ if $x$ is the minimal element of $S$ and we let $\rho_{\prec}(x)=\sup \left\{\rho_{\prec}(y)+1: y \prec x\right\}$. Intuitively, the rank of some element $x$ quantifies how far away it is from the minimal element. The following property of rank justifies the name "transfinite recursion".

Remark 3. Any element $x \in S$ has a well defined rank. This is because the transfinite recursion defines $\rho_{\prec}(x)$ as soon as all $\rho_{\prec}(y)$ is defined, for all $y \prec x$. Indeed, if $\rho(x)$ is undefined it means that $\rho_{\prec}\left(x^{\prime}\right)$ is also undefined, for some $x^{\prime} \prec x$. Repeating this for $x^{\prime}$ constructs an infinitely decreasing
chain of elements in $S$, which contradicts the fact that $\prec$ is well-founded, since an infinitely decreasing chain does not contain a minimal element.

We give some examples that help develop some intuition about the notion of the rank of a wellfounded relation. Even though a well-founded relation does not admit an infinitely decreasing chain, it can contain finite chains of arbitrary length. Essentially, the rank of some element $\rho_{\prec}(x)$ measures the length of a decreasing chain starting from $x$.

Example 1. As a warmup, consider the case where $\rho_{\prec}(x)=k$, where $k$ is a finite ordinal. Then, there exists some $x_{1} \prec x$ such that $\rho_{\prec}\left(x_{1}\right)=k-1$. Continuing in the same manner, we can see that we can create a decreasing chain of length $k+1$.
Example 2. Moving on, assume that $\rho_{\prec}(x)=\omega$. Recall that $\omega$ is the smallest infinite ordinal. This means that any $y \prec x$ has rank $\rho_{\prec}(y)=k$, for some finite number $k$. Hence, starting from $x$ we can create a decreasing chain of arbitrary length, but this length is determined when we fix the first element of the chain.
Example 3. Now assume that $\rho_{\prec}(x)=\omega+k$, for some finite number $k$. Then, we can create $x \succ x_{1} \succ \ldots \succ x_{k}$, with $\rho_{\prec}\left(x_{k}\right)=\omega$. Thus, the length of the chain is determined after we pick $x_{k+1}$.
Example 4. A slightly more involved example is the case where $\rho_{\prec}(x)=\omega+\omega$. As a first element in the chain, we pick some $x_{1} \prec x$ with $\rho_{\prec}\left(x_{1}\right)=\omega+k$, for some finite $k$. Then, we continue as in the previous example. So in this case, the length of the chain is determined by two choices.

## B. 5 Measurability

This short exposition is mainly from $\left[\mathrm{BHM}^{+} 21\right]$. A Polish space is a separable topological space that can be metrized by a complete metric. For example, the $n$-dimensional Euclidean space, any compact metric space, any separable Banach space are Polish spaces.
Definition 9 (See $\left[\mathrm{BHM}^{+} 21\right]$ ). A concept class $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ on a Polish space $\mathcal{X}$ is said to be measurable if there is a Polish space $\Theta$ and Borel-measurable map $h: \Theta \times \mathcal{X} \rightarrow[k]$ so that $\mathcal{H}=\{h(\theta, \cdot): \theta \in \Theta\}$.

Roughly, a subset $B$ of a Polish space $\mathcal{X}$ is universally measurable if it is measurable with respect to every complete probability measure on $\mathcal{X}$.
Definition 10 (Universally Measurable). Let $\mathcal{F}$ be the Borel $\sigma$-field on a Polish space $\mathcal{X}$. For any probability measure $\mu$, denote by $\mathcal{F}_{\mu}$ the completion of $\mathcal{F}$ with respect to $\mu$, that is, the collection of all subsets of $\mathcal{X}$ that differ from a Borel set at most on a set of zero probability. A set $B \subseteq \mathcal{X}$ is called universally measurable if $B \in \mathcal{F}_{\mu}$ for every probability measure $\mu$. Similarly, a function $f: \mathcal{X} \rightarrow \mathcal{Y}$ is called universally measurable if $f^{-1}(B)$ is universally measurable for any universally measurable set $B$.

Let $\mathcal{X}, \mathcal{Y}$ be Polish spaces, and let $f: \mathcal{X} \rightarrow \mathcal{Y}$ be a continuous function. It holds that $f$ is Borel measurable, that is, $f^{-1}(B)$ is a Borel subset of $\mathcal{X}$ for any Borel subset $B$ of $\mathcal{Y}$. A subset $B \subseteq \mathcal{X}$ of a Polish space is called analytic if it is the image of some Polish space under a continuous map. The complement of an analytic set is called coanalytic. A set is Borel if and only if it is both analytic and coanalytic. The following is a consequence of Choquet's Capacitability Theorem.
Fact 1. Every analytic (or coanalytic) set is universally measurable.
An important property of analytic sets is that they help us establish measurability of uncountable unions of measurable sets.
We now state a very important result regarding well-founded relations on Polish spaces. We let $\mathcal{X}$ be a Polish space and $\prec$ a well-found relation on $\mathcal{X}$. We say that $\prec$ is analytic if $R_{\prec} \subseteq \mathcal{X} \times \mathcal{X}$ is an analytic set. The following theorem, known as Kunen-Martin [Kec12, Del77], bounds the rank function of such a relation.

Theorem 7. Let $\prec$ be an analytic well-founded relation on a Polish space $\mathcal{X}$. Then, its rank function satisfies $\sup _{x \in \mathcal{X}} \rho_{\prec}(x) \leq \omega_{1}$.

## B. 6 Gale-Stewart Games

We add a short discussion about GS games from $\left[\mathrm{BHM}^{+} 21\right]$. We refer to their work for further details and pointers. Fix sequences of sets $\mathcal{X}_{t}, \mathcal{Y}_{t}$ for $t \geq 1$. We consider infinite games between two players: in each round $t \geq 1$, the first player $\mathrm{P}_{A}$ selects an element $x_{t} \in \mathcal{X}_{t}$, and then player $\mathrm{P}_{L}$ selects an element $y_{t} \in \mathcal{Y}_{t}$. The rules of the game are determined by specifying a set $\mathcal{W} \subseteq \prod_{t \geq 1}\left(\mathcal{X}_{t} \times \mathcal{Y}_{t}\right)$ of winning sequences for $\mathrm{P}_{L}$. That is, after an infinite sequence of consecutive plays $x_{1}, y_{1}, x_{2}, y_{2}, \ldots$, we say that $\mathrm{P}_{L}$ wins if $\left(x_{1}, y_{1}, x_{2}, y_{2}, \ldots\right) \in \mathcal{W}$; otherwise, $\mathrm{P}_{A}$ is declared the winner of the game.
A strategy is a rule used by a given player to determine the next move given the current position of the game. A strategy for $\mathrm{P}_{A}$ is a sequence of functions $f_{t}: \prod_{s<t}\left(\mathcal{X}_{s} \times \mathcal{Y}_{s}\right) \rightarrow \mathcal{X}_{t}$ for $t \geq 1$, so that $\mathrm{P}_{A}$ plays $x_{t}=f_{t}\left(x_{1}, y_{1}, \ldots, x_{t-1}, y_{t-1}\right)$ in round $t$. Similarly, a strategy for $\mathrm{P}_{L}$ is a sequence of $g_{t}: \prod_{s<t}\left(\mathcal{X}_{s} \times \mathcal{Y}_{s}\right) \times \mathcal{X}_{t} \rightarrow \mathcal{Y}_{t}$ for $t \geq 1$, so that $\mathrm{P}_{L}$ plays $y_{t}=g_{t}\left(x_{1}, y_{1}, \ldots, x_{t-1}, y_{t-1}, x_{t}\right)$ in round $t$. A strategy for $\mathrm{P}_{A}$ is called winning if playing that strategy always makes $\mathrm{P}_{A}$ win the game regardless of what $\mathrm{P}_{L}$ plays; a winning strategy for $\mathrm{P}_{L}$ is defined analogously. The crucial question follows:

> When do winning strategies exist in infinite two-player games?

The simplest assumption was introduced by [GS53]: we call $\mathcal{W}$ finitely decidable if for every sequence of plays $\left(x_{1}, y_{1}, x_{2}, y_{2}, \ldots\right) \in \mathcal{W}$, there exists $n<\infty$ so that

$$
\left(x_{1}, y_{1}, \ldots, x_{n}, y_{n}, x_{n+1}^{\prime}, y_{n+1}^{\prime}, x_{n+2}^{\prime}, y_{n+2}^{\prime}, \ldots\right) \in \mathcal{W}
$$

for all choices of $x_{n+1}^{\prime}, y_{n+1}^{\prime}, x_{n+2}^{\prime}, y_{n+2}^{\prime}, \ldots$ In words, that " $\mathcal{W}$ is finitely decidable" means that if $P_{L}$ wins, then she knows that she won after playing a finite number of rounds. Conversely, in this case $P_{A}$ wins the game precisely when $P_{L}$ does not win after any finite number of rounds.
An infinite game whose set $\mathcal{W}$ is finitely decidable is called a Gale-Stewart game. The main result behind GS games follows.
Proposition. In any Gale-Stewart game, either $P_{A}$ or $P_{L}$ has a winning strategy.
The above existential result provides no information, however, about the complexity of the winning strategies. In particular, it is completely unclear whether winning strategies can be chosen to be measurable. The next lemma addresses this concern.
Lemma 2 (Theorem B. 1 of $\left[\mathrm{BHM}^{+} 21\right]$ ). Let $\left\{X_{t}\right\}_{t \geq 1}$ be Polish spaces and $\left\{Y_{t}\right\}_{t \geq 1}$ be countable sets. Consider a Gale-Stewart game whose set $\mathcal{W} \subseteq \prod_{t \geq 1}\left(X_{t} \times Y_{t}\right)$ of winning strategies for $P_{L}$ is finitely decidable and coanalytic. Then there is a universally measurable winning strategy.

For an extensive exposition on game values which connect ordinals with the positions of the game, i.e., the sequences of choices between the two players, we refer to $\left[\mathrm{BHM}^{+} 21\right]$.

## C Universal Multiclass Learning: The Proof of Theorem 2

In this section, we prove Theorem 2 which then directly gives us Theorem 1 as well. Our text is organized as follows:

- In Appendix C.1, we analyze the exponential rates case for multiclass learning. We introduce the notion of the multiclass Littlestone tree and prove Theorem 8.
- In order to achieve this, we first analyze the problem from an adversarial online perspective in Appendix C.1.1. The main result in this section is Theorem 9.
- The above result is in the adversarial setting and hence we have to transform this online algorithm into a statistical one, i.e., we have to move from the adversarial setting to the probabilistic setting. This is done is Appendix C.1.2. We provide the analysis of our final algorithm in Theorem 10.
- In Appendix C.2, we show that a class with infinite multiclass Littlestone tree cannot be learned at a rate faster than linear.
- In Appendix C.3, we introduce the notion of a Natarajan-Littlestone tree and we prove Theorem 12.
- In Appendix C.3.1, we provide the important notion of a Natarajan-Littlestone game which guarantees the existence of an eventually correct algorithm when the class does not have an infinite NL tree.
- In Appendix C.3.2, we introduce the notion of an NL pattern in the data and using the above algorithm as a pattern avoidance function.
- The behavior of the pattern avoidance algorithm in the probabilistic setting is given in Appendix C.3.3.
- The final linear rate algorithm can be found at Appendix C.3.4.
- In Appendix C.4, the final lower bound which is related to arbitrarily slow rates (see Theorem 14).
- In Appendix C.5, we provide the notion of a Graph-Littlestone tree and give a sufficient condition for learning with a linear rate. More to that, we propose Open Question 1.


## C. 1 Exponential Rates for Multiclass Learning

For a sequence $\boldsymbol{y}=\left(y_{1}, y_{2}, \ldots\right)$, we denote $\boldsymbol{y}_{\leq k}=\left(y_{1}, \ldots, y_{k}\right)$. We may also usually identify elements of $\{0,1\}^{d}$ with strings or a prefix of a sequence of length $d$. We begin with a formal definition of a crucial combinatorial measure, namely the multiclass Littlestone tree of a class $\mathcal{H}$.
Definition 11 (Multiclass Littlestone Tree). A multiclass Littlestone tree for $\mathcal{H} \subseteq\{0,1, \ldots, k\}^{\mathcal{X}}$ is a complete binary tree of depth $d \leq \infty$ whose internal nodes are labeled by $\mathcal{X}$, and whose two edges connecting a node to its children are labeled by two different elements in $[k]$, such that every path of length at most d emanating from the root is consistent with a concept $h \in \mathcal{H}$. Typically, a multiclass Littlestone tree is a collection

$$
\bigcup_{0 \leq \ell<d}\left\{x_{u}: u \in\{0,1\}^{\ell}\right\}=\left\{x_{\emptyset}\right\} \cup\left\{x_{0}, x_{1}\right\} \cup\left\{x_{00}, x_{01}, x_{10}, x_{11}\right\} \cup \ldots
$$

such that for every path $\boldsymbol{y} \in\{0,1\}^{d}$ and finite $n<d$, there exists $h \in \mathcal{H}$ so that $h\left(x_{\boldsymbol{y}_{\leq \ell}}\right)=s_{\boldsymbol{y}_{\leq \ell+1}}$ for $0 \leq \ell \leq n$, where $s_{\boldsymbol{y}_{\leq \ell+1}}$ is the label of the edge connecting the nodes $x_{\boldsymbol{y}_{\leq \ell}}$ and $x_{\boldsymbol{y}_{\leq \ell+1}}$. We say that $\mathcal{H}$ has an infinite multiclass Littlestone tree if there is a multiclass Littlestone tree for $\mathcal{H}$ of depth $d=\infty$.

To give some intuition about this construction we state some of its properties. First, it is crucial to note that a class $\mathcal{H}$ with a finite multiclass Littlestone tree can have infinite multiclass Littlestone dimension. In fact, a class has finite multiclass Littlestone dimension if the depth of the tree admits a uniform upper bound. However, it may be the case that the class admits an unbounded tree, in the sense that for any finite depth, there exists a tree of that depth; nevertheless the class does not have an infinite tree. Second, a class with a finite multiclass Littlestone tree may contain trees with infinite paths (e.g., a class that contains the constant mapping $h=1$ can shatter the rightmost path at an infinite depth). Finally, a class with finite multiclass Littlestone tree cannot have a tree with an infinite complete binary subtree, since one could use only this subtree and obtain an infinite tree. Essentially, this tree captures "how infinite" the multiclass Littlestone dimension of $\mathcal{H}$ is: even if there is no uniform bound on the multiclass Littlstone dimension of $\mathcal{H}$, whenever $\mathcal{H}$ does not have an infinite tree we know that for any $S=\bigcup_{0 \leq \ell<d}\left\{x_{u}: u \in\{0,1\}^{\ell}\right\}$ there is some $n^{*}(S)<\infty$, which depends on the sequence $S$, so that the tree can be shattered up to level $n^{*}(S)$.
The goal of this section is to prove the next result.
Theorem 8 (Exponential rates). If $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ does not have an infinite multiclass Littlestone tree, then for any distribution $P, \mathcal{H}$ is learnable with an exponential optimal rate.

In order to prove this result, two steps are required: first, it suffices to show that no class of hypotheses can be learned in a rate faster than exponential (lower bound) and second, we have to show that not having an infinite multiclass Littlestone tree is a sufficient for learning at an exponential rate condition (upper bound). Hence, these two directions tightly characterize universal learnability at exponential rates for multiclass classification.
For the lower bound, the argument is essentially the same as in $\left[\mathrm{BHM}^{+} 21\right]$. Let us first give some intuition. If the distribution is supported on a finite number of points, e.g., two, there is an exponentially small probability that all the $n$ samples will contain the same point $(x, y)$, so it will
not help the learner distinguish between the functions $h \in \mathcal{H}$ that label this point correctly. For completeness, we present the argument formally.
Proposition 1 (Exponential Rates (Lower Bound)). Fix $\mathcal{H} \subseteq\{0,1, \ldots, k\}^{\mathcal{X}}$. For any learning algorithm $\widehat{h}_{n}$, there exists a realizable distribution P over $\mathcal{X} \times\{0,1, \ldots, k\}$ such that $\mathbf{E}\left[\operatorname{er}\left(\widehat{h}_{n}\right)\right] \geq$ $\Omega\left(2^{-n}\right)$ for infinitely many $n$. This means that $\mathcal{H}$ is not learnable at rate faster than exponential.

Proof. Recall that $|\mathcal{H}|>k+2$. Hence, there are $h_{0}, h_{0} \in \mathcal{H}$ and $x, x^{\prime} \in \mathcal{X}$ such that $h_{0}(x)=$ $h_{1}(x)=y$ and $h_{0}\left(x^{\prime}\right)=y_{0} \neq h_{1}(x)=y_{1}$. We fix some learning algorithm $\widehat{h}_{n}$ and two distributions $P_{0}, P_{1}$ where $P_{i}\{(x, y)\}=\frac{1}{2}, P_{i}\left\{\left(x, y_{i}\right)\right\}=\frac{1}{2}, i \in\{0,1\}$. We let $I \sim$ Bernoulli(1/2) and given I, we let $\left(X_{1}, Y_{1}\right),\left(X_{2}, Y_{2}\right), \ldots$ be i.i.d. samples from $P_{I}$ and $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ are the training samples for $\widehat{h}_{n}$ and $\left(X_{n+1}, Y_{n+1}\right)$ is the test point. Then, we have that

$$
\mathbf{E}\left[\operatorname{Pr}\left(\widehat{h}_{n}\left(X_{n+1}\right) \neq Y_{n+1} \mid\left\{X_{t}, Y_{y}\right\}_{t=1}^{n}, I\right)\right] \geq \frac{1}{2} \operatorname{Pr}\left(X_{1}=\ldots=X_{n}, X_{n+1}=x^{\prime}\right)=2^{-n-2}
$$

We also have that

$$
\begin{aligned}
& \mathbf{E}\left[\operatorname{Pr}\left(\widehat{h}_{n}\left(X_{n+1}\right) \neq Y_{n+1} \mid\left\{X_{t}, Y_{y}\right\}_{t=1}^{n}, I\right)\right] \\
& =\frac{1}{2} \sum_{i \in\{0,1\}} \mathbf{E}\left[\operatorname{Pr}\left(\widehat{h}_{n}\left(X_{n+1}\right) \neq Y_{n+1} \mid\left\{X_{t}, Y_{y}\right\}_{t=1}^{n}, I=i\right) \mid I=i\right]
\end{aligned}
$$

Thus, for every $n$, there exists some $i_{n}\{0,1\}$ such that for $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ i.i.d. from $P_{i_{n}}$ it holds that

$$
\mathbf{E}\left[\operatorname{er}_{P_{i_{n}}}\left(\widehat{h}_{n}\right)\right] \geq 2^{-n-2}
$$

Hence, there exists some fixed $i \in\{0,1\}$ such that $\mathbf{E}\left[\operatorname{er}_{P_{i}}\left(\widehat{h}_{n}\right)\right] \geq 2^{-n-2}$ for infinitely many $n$.
We continue with the upper bound, i.e., the design of an algorithm that learns $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ at an exponentially fast rate when its multiclass Littlestone tree is not infinite. Our approach consists of two main steps. First, we consider the classical online adversarial setting where a learner has to guess the label of a point that is presented to her by an adversary. We prove that if $\mathcal{H}$ does not have an infinite multiclass Littlestone tree, there is a strategy the learner can employ that makes a finite number of mistakes. Then, given such a strategy, we prove that there is an algorithm in the statistical setting which achieves exponential learning rates. The details are outlined in the subsequent sections.

## C.1. 1 Viewing Exponential Rates in an Online Setting

In order to design our algorithms, we have to consider the following setting. We introduce the multiclass online learning game (Figure 3) that has been studied extensively (see e.g., [DSBDSS11]). In this game, there are two players, the adversary that chooses features and reveals them to the second player, the learner whose goal is to guess a label for the given example. The learner makes a mistake

1. The adversary picks a point $x_{t} \in \mathcal{X}$.
2. The learner guesses a value $\widehat{y}_{t} \in[k]$
3. The adversary chooses the value $y_{t}$ as true label so that $y_{t}=h\left(x_{t}\right)$ for some $h \in \mathcal{H}$ that is consistent with the previous examples $\left(x_{p}, y_{p}\right)$ for any $p \leq t$.

Figure 3: Realizable Online Setting
in round $t$ whenever the guess $\widehat{y}_{t}$ differs from the true label $y_{t}$. The goal of the learner is to minimize her loss and the adversary's intention is to provoke many errors to the learner.

We say that the concept class $\mathcal{H}$ is online learnable if there exists a strategy $\widehat{y}_{t}=$ $\widehat{y}_{t}\left(x_{1}, y_{1}, \ldots, x_{t-1}, y_{t-1}, x_{t}\right)$ that makes a mistake only finitely many times, regardless of what realizable sequence is presented by the adversary. Notice that compared to the classical online learning setting that asks for a bounded number of mistakes $d$, we settle for a more modest goal. The main result in this setting is the following.

Theorem 9 (Strategies in the Adversarial Setting). For any concept class $\mathcal{H} \subseteq[k]^{\mathcal{X}}$, the following dichotomy occurs.

1. If $\mathcal{H}$ does not have an infinite multiclass Littlestone tree, then there is a strategy for the learner that makes only finitely many mistakes against any adversary.
2. If $\mathcal{H}$ has an infinite multiclass Littlestone tree, then there is a strategy for the adversary that forces any learner to make a mistake in every round.

Before the formal proof, we provide a proof sketch. The adversary's strategy is clear. Whenever the learner predicts a label, she must choose a different label that will cause the learner to make a mistake in that round; an infinite Littlestone tree is exactly the combinatorial structure the adversary is looking for. On the other hand, since there is no infinite multiclass Littlestone tree, the learner's predictions should lead her to a leaf of the Littlestone tree that is defined by her interaction with the learner. This is established by a variant of the (multiclass) Standard Optimal Algorithm (SOA) [Lit88, DSBDSS11], which works whenever $\mathcal{H}$ has a finite multiclass Littlestone dimension $d$. To gain some intuition, it is instructive to consider the first step of her algorithm. The learner is presented with a point $x_{1}$ and she picks the label $y_{1}=\operatorname{argmax}_{y \in[k]} \operatorname{Ldim}_{k}\left(\mathcal{H}_{x_{1}, y}\right)$. It is easy to see that there is at most one $y$ such that $\operatorname{Ldim}_{k}\left(\mathcal{H}_{x_{1}, y}\right)=d$; if there were two of them then $\operatorname{Ldim}_{k}(\mathcal{H})=d+1$. Thus, by picking the one that induces the subset of the hypothesis class with the largest Littlestone dimension, the learner either does not make a mistake or gets closer to a leaf of the tree.

In our universal setting, this approach does not work immediately since it may be the case that the multiclass Littlestone tree is not infinite but the associated class has infinite multiclass Littlestone dimension. To this end, we have to introduce the ordinal multiclass Littlestone dimension, which quantifies "how infinite" the multiclass Littlestone dimension is. Hence, the learner's strategy will be to play according to the ordinal Standard Optimal Algorithm. The intuition is similar as in the classical setting, but the proof becomes more involved. To establish the result, we follow the approach of $\left[\mathrm{BHM}^{+} 21\right]$ and define a more general infinite game $\mathcal{G}$ between the learner and the adversary and prove that it belongs to the family of the so-called Gale-Stewart games. Then, we leverage results from the theory of Gale-Stewart games and $\left[\mathrm{BHM}^{+} 21\right]$ in order to show that the ordinal SOA makes a finite number of mistakes whenever $\mathcal{H}$ does not have an infinite multiclass Littlestone tree. Let us continue with the proof.

Proof of Theorem 9. We first introduce a two-player game $\mathcal{G}$ that is played in discrete timesteps $t=1,2, \ldots$ between the adversary and the learner.

1. The adversary picks a point $\kappa_{t}=\left(\xi_{t}, y_{t}^{(0)}, y_{t}^{(1)}\right) \in \mathcal{X} \times[k] \times[k]$ and reveals it to the learner.
2. The learner chooses a point $\eta_{t} \in\{0,1\}$.

Figure 4: Adversarial Setting - 2-Player Game

The learning player wins in some finite round $t$ if $\mathcal{H}_{\xi_{1}, y_{1}^{\left(\eta_{1}\right)}, \ldots, \xi_{t}, y_{t}^{\left(\eta_{t}\right)}}=\emptyset$. The adversary wins if the game continues indefinitely (i.e., the class of consistent hypotheses from $\mathcal{H}$ never gets empty). Clearly, the set of winning strategies for the learning player is
$\mathcal{W}=\left\{(\boldsymbol{\kappa}, \boldsymbol{\eta}) \in(\mathcal{X} \times[k] \times[k] \times\{0,1\})^{\infty}: \exists 1 \leq t^{\star}<\infty\right.$ such that $\mathcal{H}_{\left.\xi_{1}, y_{1}^{\left(\eta_{1}\right)}, \ldots, \xi_{t^{\star}, y_{t^{\star}}^{\left(\eta_{\star}\right)}}=\emptyset\right\} . . . ~} \quad$.
We now recall an important theorem (see Appendix B.6) about Gale-Stewart games: In any GaleStewart game, exactly one of the adversary player and the learning player has a winning strategy.

Equipped with Appendix B.6, we can show that the adversary has a winning strategy if and only if $\mathcal{H}$ has an infinite multiclass Littlestone tree (provided that $\mathcal{G}$ is a Gale-Stewart game). This is summarized in the next claim.

Claim 1. The game $\mathcal{G}$ is a Gale-Stewart game and the adversary has a winning strategy in $\mathcal{G}$ if and only if the hypothesis class $\mathcal{H}$ has an infinite multiclass Littlestone tree.

Proof. It is clear from the definition of $\mathcal{W}$ that every winning strategy of the learner is finitely decidable, hence $\mathcal{G}$ is a Gale-Stewart game. For the other part of the claim, notice that if $\mathcal{H}$ has an infinite multiclass Littlestone tree, then the adversary's strategy is to present the learner at step $t$ the point of the tree at depth $t$ that is consistent with the execution of the game so far along with the labels of the edges that connect it with its children. By the definition of the tree, this strategy ensures that the game will keep going on forever. For the other direction, assume that the adversary has a winning strategy $\kappa_{\tau}\left(\eta_{1}, \ldots, \eta_{\tau-1}\right) \in \mathcal{X} \times[k] \times[k]$. Then, define the multiclass Littlestone tree $\mathcal{T}=\left\{x_{\boldsymbol{u}}: 0 \leq k<\infty, \boldsymbol{u} \in\{0,1\}^{k}\right\}$ where $x_{\eta_{1}, \ldots, \eta_{\tau-1}}=\xi_{\tau}\left(\eta_{1}, \ldots, \eta_{\tau-1}\right)$ where the labels that connect $x_{\eta_{1}, \ldots, \eta_{\tau-1}}$ with its left, right children are $y_{\tau}^{(0)}\left(\eta_{1}, \ldots, \eta_{\tau-1}\right), y_{\tau}^{(1)}\left(\eta_{1}, \ldots, \eta_{\tau-1}\right)$, respectively. We can see that $\mathcal{T}$ is infinite since this is a winning strategy for the adversary.

Having shown the above statement, we are ready to establish the desired dichotomy in the online game. Assume first that $\mathcal{H}$ has an infinite multiclass Littlestone tree $\left\{x_{u}\right\}$. The adversary's strategy is defined inductively based on the path followed so far in the game: In round $t$, set $\boldsymbol{b}_{t}=\left(b_{1}, \ldots, b_{t-1}\right) \in$ $\{0,1\}^{t-1}$ denote the path parsed so far in the tree by the two players. Then, the adversary picks $x_{t}=x_{b_{t}}$. After the learner reveals her choice $\widehat{y}_{t}$, the worst case adversary chooses as a response the branch of the Littlestone tree which does not correspond to the learner's choice (the adversary may even have two choices). By the definition of the tree, this chosen label is valid since there exists some $h \in \mathcal{H}$ that realizes the path $\left(x_{b_{1}}, \ldots, x_{b_{t-1}}, x_{t}\right)$. Moreover, this choice provokes a mistake to the learning player and this is true for any round. Hence, there is a strategy for the adversary that forces any learner to make a mistake in every round.
For the other direction, assume that the class $\mathcal{H}$ does not have an infinite multiclass Littlestone tree. Before we describe the winning strategy of the learner, we need to introduce the notion of ordinal multiclass Littlestone dimension. We will assign an ordinal to every finite multiclass Littlestone tree. For some preliminaries on ordinals and transfinite recursion, we refer to $\left[\mathrm{BHM}^{+} 21\right]$. The rank is defined by a partial order $\prec$. We set $t^{\prime} \prec t$ if $t^{\prime}$ is a multiclass Littlestone tree that extends $t$ by one level, i.e., $t$ is obtained from $t^{\prime}$ by removing its leaves. A multiclass Littlestone tree $t$ is minimal if it cannot be extended to a multiclass Littlestone tree of larger depth. For such a tree, we set $\operatorname{rank}(t)=0$. If the tree $t$ is non-minimal, then it can be extended and this is quantified using transfinite recursion by

$$
\operatorname{rank}(t)=\sup \left\{\operatorname{rank}\left(t^{\prime}\right)+1: t^{\prime} \prec t\right\}
$$

The rank is well-defined as long as $\mathcal{H}$ has no infinite multiclass Littlestone tree (since $\prec$ is wellfounded). In particular, we define

$$
\overline{\operatorname{Ldim}}_{k}(\mathcal{H})= \begin{cases}-1 & \text { if } \mathcal{H} \text { is empty } \\ \Omega & \text { if } \mathcal{H} \text { has an infinite multiclass Littlestone tree } \\ \operatorname{rank}(\emptyset) & \text { otherwise }\end{cases}
$$

The strategy is chosen so that $\overline{\operatorname{Ldim}}_{k}\left(\mathcal{H}_{x_{1}, y_{1}, \ldots, x_{t}, y_{t}}\right)$ decreases in every round and the learner that follows this strategy will win the game, since the ordinals do not admit an infinite decreasing chain. We note that this statement at first is purely existential via the theory of Gale-Stewart games. We next shortly provide a "constructive" way to compute the winning strategy of the learning player in the set of games we consider. Let us describe the winning strategy: The learner invokes the ordinal (multiclass) Standard Optimal Algorithm and chooses the label $y_{t}$ (given $x_{t}$ ) that maximizes the ordinal multiclass Littlestone dimension, i.e., $y_{t}=\operatorname{argmax}_{y \in[k]} \overline{\operatorname{Ldim}}_{k}\left(V_{t}^{y}\right)$, where $=V_{t}^{y}=\left\{h \in \mathcal{H}_{x_{1}, y_{1}, \ldots, x_{t-1}, y_{t-1}}: h\left(x_{t}\right)=y\right\}$. The ordinal SOA at round $t=1,2, \ldots$ with initial set $V_{0}=\mathcal{H}$ works as follows:

1. Receive $x_{t}$.
2. For any $y \in[k]$, let $V_{t}^{y}=\left\{h \in V_{t-1}: h\left(x_{t}\right)=y\right\}$.
3. Predict $\widehat{y}_{t} \in \operatorname{argmax}_{y \in[k]} \overline{\operatorname{Ldim}}_{k}\left(V_{t}^{y}\right)$, where $\overline{\operatorname{Ldim}}_{k}$ is the ordinal multiclass Littlestone dimension.
4. Receive true answer $y_{t}$ and set $V_{t}=V_{t}^{y_{t}}$.

This algorithm drives the game in a win-win phenomenon for the learner in every round: If the adversary forces the learner to a mistake, then she will "prune" the tree and set the learner closer
to winning the game. Otherwise, the learner will be correct and will not incur any loss. In order to show that the ordinal SOA makes a finite number of mistakes, we couple the online game with a Gale-Stewart game. The idea is that every time the learner makes a mistake in the online game on point $x_{t}$, we advance the Gale-Stewart game by one round where we pretend that $\xi_{\tau}=x_{t}, y_{\tau}^{(0)}=$ $\widehat{y}_{t}, y_{\tau}^{(1)}=y_{t}, \eta_{\tau}=y_{t}$. Notice that if the learner makes an infinite number of mistakes in the online game using the ordinal SOA, then the Gale-Stewart game can proceed infinitely. Hence, to conclude the proof, we need to show that in this coupled game, there is some finite point $\tau^{*}$ such that $\mathcal{H}_{\xi_{1}, y_{1}^{\left(\eta_{1}\right)}, \ldots, \xi_{\tau^{\star}}, y_{\tau^{\star}}^{\left(\eta_{\left.\tau^{\star}\right)}\right)}}=\emptyset$. The following result helps us establish that. In fact, the next lemma follows from $\left[\mathrm{BHM}^{+} 21\right]$ (Proposition B.8) by choosing the value of the game being the ordinal multiclass Littlestone dimension.

Lemma 3 (See Proposition B. 8 of $\left[\mathrm{BHM}^{+} 21\right]$ ). Assume that $\mathcal{H}$ does not contain an infinite multiclass Littlestone tree. Then, for any choices of the adversary $\kappa_{1}, \ldots, \kappa_{t-1}$ up to round $t$ and for any choice $\kappa_{t}=\left(\xi_{t}, y_{t}^{(0)}, y_{t}^{(1)}\right)$ in round there is a choice $\eta_{t}$ of the learner such that

$$
\overline{\operatorname{Ldim}}_{k}\left(\mathcal{H}_{\xi_{1}, y_{1}^{\left(\eta_{1}\right)}, \ldots, \xi_{t-1}, y_{t-1}^{\left(\eta_{t-1}\right)}, \xi_{t}, y_{t}^{\left(\eta_{t}\right)}}\right)<\overline{\operatorname{Ldim}}_{k}\left(\mathcal{H}_{\xi_{1}, y_{1}^{\left(\eta_{1}\right)}, \ldots, \xi_{t-1}, y_{t-1}^{\left(\eta_{t-1}\right)}}\right) .
$$

The previous result shows that for every $\xi_{t}$ there is at most one label $\ell_{t} \in[k]$ such that

$$
\overline{\operatorname{Ldim}}_{k}\left(\mathcal{H}_{\xi_{1}, y_{1}^{\left(\eta_{1}\right)}, \ldots, \xi_{t-1}, y_{t-1}^{\left(\eta_{t-1}\right)}, \xi_{t}, \ell_{t}}\right)=\overline{\operatorname{Ldim}}_{k}\left(\mathcal{H}_{\xi_{1}, y_{1}^{\left(\eta_{1}\right)}, \ldots, \xi_{t-1}, y_{t-1}^{\left(\eta_{t-1}\right)}}\right) .
$$

Indeed, assume that there are two such labels $\ell_{t}, \ell_{t}^{\prime}$ for some $\xi_{t}$. Then, if the adversary proposes the point $\left(\xi_{t}, \ell_{t}, \ell_{t}^{\prime}\right)$, there is no choice $\eta_{t}$ of the learner that decreases that ordinal Littlestone dimension in this round, which leads to a contradiction. Hence, the learner can pick any label as long as it is not the one that maximizes the ordinal Littlestone dimension. This is exactly how the coupled Gale-Stewart game proceeds, so we know that every time the learner makes a mistake in the online game the ordinal Littlestone dimension of the coupled game decreases. Since ordinals that are less than $\Omega$ do not admit infinitely decreasing chains, we get the desired result.

## C.1.2 Moving from the Adversarial Setting to the Probabilistic Setting

The measurability of the winning strategies and of the learning algorithm developped in the previous section constitutes an important detail, extensively discussed in $\left[\mathrm{BHM}^{+} 21\right]$, in order to move from the adversarial setting to the probabilistic one. We provide the next useful result.
Lemma 4. Let $\mathcal{X}$ be Polish, $k \in \mathbb{N}$ be a finite constant and $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ be measurable. Then, the Gale-Stewart game $\mathcal{G}$ of Figure 4 has a universally measurable winning strategy.

Proof. It suffices to prove that the set $\mathcal{W}$ of the winning strategies for the learning player in the Gale-Stewart game is coanalytic (see Lemma 2). Equivalently, we will prove that the set of winning strategies of the adversary $\mathcal{W}^{c}$ is analytic, where

$$
\mathcal{W}^{c}=\left\{(\boldsymbol{\kappa}, \boldsymbol{\eta}) \in(\mathcal{X} \times[k] \times[k] \times\{0,1\})^{\infty}: \mathcal{H}_{\xi_{1}, y_{1}^{\left(\eta_{1}\right)}, \ldots, \xi_{t}, y_{t}^{\left(\eta_{t}\right)}} \neq \emptyset \text { for all } t<\infty\right\}
$$

This set is equal to

$$
\mathcal{W}^{c}=\bigcap_{1 \leq \tau<\infty} \bigcup_{\theta \in \Theta} \bigcap_{1 \leq t \leq \tau}\left\{(\boldsymbol{\kappa}, \boldsymbol{\eta}) \in(\mathcal{X} \times[k] \times[k] \times\{0,1\})^{\infty}: h\left(\theta, \xi_{t}\right)=y_{t}^{\left(\eta_{t}\right)}\right\}
$$

The set $\left\{(\theta, \boldsymbol{\kappa}, \boldsymbol{\eta}): h\left(\theta, \xi_{i}\right)=y_{i}^{\left(\eta_{j}\right)}\right\}=\bigcup_{\left(y_{i}^{(0)}, y_{i}^{(1)}\right) \in[k] \times[k], y_{i}^{(0)} \neq y_{i}^{(1)}}\left\{(\theta, \boldsymbol{\xi}, \boldsymbol{\eta}): h\left(\theta, \xi_{i}\right)=y_{i}^{\left(\eta_{j}\right)}\right\}$ is a Borel set using the standard measurability assumption of Definition 9. The set $\mathcal{W}^{c}$ is analytic since the two intersections are over countable sets and the union is a projection of a Borel set.

Crucially the above result states that the winning strategy $\eta_{t}$ of the learning player is measurable. However, the previous proof made use of the ordinal multiclass SOA algorithm, whose measurability is not directly implied. To this end, we modify the adversarial algorithm to handle the measurability issue. The modification follows:

1. Initialize $\tau \leftarrow 1, G=\operatorname{Clique}(V=[k]), f(\cdot, \cdot, \cdot) \leftarrow \eta_{1}(\cdot, \cdot, \cdot) \quad \triangleright \tau$ is the mistake counter
2. For every round $t \geq 1$ :
(a) Observe $x_{t}$
(b) For any $y \neq y^{\prime}$ with $y, y^{\prime} \in[k]$, orient the edge $\left(y, y^{\prime}\right)$ of $G$ according to $f\left(x_{t}, y, y^{\prime}\right)$
(c) Let $G^{\prime}$ the directed clique
(d) Predict $\widehat{y}_{t} \leftarrow \operatorname{argmax}_{y \in[k]} \operatorname{outdeg}\left(y ; G^{\prime}\right)$
(e) If $\widehat{y}_{t} \neq y_{t}$, let $\xi_{\tau} \leftarrow x_{t}, f(\cdot, \cdot, \cdot) \leftarrow \eta_{\tau+1}\left(x_{1}, y_{1}, \ldots, x_{\tau}, y_{\tau}, \cdot, \cdot, \cdot\right), \tau \leftarrow \tau+1$

Figure 5: Measurable Modification of Online Learning Algorithm for Exponential Rates

The above algorithm makes use of a tournament procedure. The algorithm is a measurable function since (i) the winning strategy of the learner is measurable and (ii) the countable maximum of measurable functions is measurable. This algorithm can be used in order to show that if $\mathcal{H}$ does not have an infinite multiclass Littlestone tree, then the above algorithm makes only a finite number of mistakes against any adversary. Essentially, this is due to the fact that when the winning strategy has converged to a zero-mistake prediction rule (which occurs after a finite number of mistakes), the tournament procedure will always output the correct label for the observed example. Hence, the algorithm will eventually make a finite number of mistakes in the adversarial setting.
The algorithm of Figure 5 works in the adversarial setting. We first show that it also applies to the probabilistic setting (and this is why we require the above measurability discussion). The proof is quite similar to Lemma 4.3 of $\left[\mathrm{BHM}^{+} 21\right]$.
Lemma 5 (From Adversarial to Probabilistic). For any distribution $P$ over $\mathcal{X} \times[k]$ and for the learning algorithm $\widehat{y}_{t}: \mathcal{X} \rightarrow[0,1]$ of Theorem 9, we have

$$
\mathbf{P r}_{S_{t}}\left[\underset{(x, y) \sim P}{\mathbf{P r}}\left[\widehat{y}_{t}(x) \neq y\right]>0\right] \rightarrow 0 \text { as } t \rightarrow \infty
$$

where $S_{t}$ is the training set $\left(x_{1}, y_{1}, \ldots, x_{t-1}, y_{t-1}\right)$ of the algorithm.
Proof. Since the distribution $P$ is realizable, there exists a sequence of functions $h_{k} \in \mathcal{H}$ so that

$$
\underset{(x, y) \sim P}{\mathbf{P r}}\left[h_{k}(x) \neq y\right]<\frac{1}{2^{k}} .
$$

Let us fix $t \geq 1$. We have that

$$
\sum_{k=1}^{\infty} \operatorname{Pr}\left[\exists s \leq t: h_{k}\left(X_{s}\right) \neq Y_{s}\right] \leq t \sum_{k=1}^{\infty} \operatorname{Pr}_{(X, Y) \sim P}\left[h_{k}(X) \neq Y\right]<\infty
$$

where the first inequality is due to union bound. By Borel-Cantelli, with probability one, there exists for every $t \geq 1$ a hypothesis $h \in \mathcal{H}$ so that $h\left(X_{s}\right)=Y_{s}$ for all $s \leq t$. Hence, the sequence $X_{1}, Y_{1}, X_{2}, Y_{2}, \ldots$ is a valid input for the online learning game with probability one. In particular, we make use of the following statement: If $\mathcal{H}$ does not have an infinite multiclass Littlestone tree, then there is a strategy for the learner that makes only finitely many mistakes against any adversary. This is proved in Theorem 9. The existence of a winning strategy $\widehat{y}_{t}$ for the learning player implies that the time $T$ where the player makes a mistake is

$$
T=\sup \left\{s \in \mathbb{N}: \widehat{y}_{s-1}\left(X_{s}\right) \neq Y_{s}\right\}
$$

is a random variable that is finite with probability one. Moreover, the online learner is selected so that it is changed only when a loss is observed. This means that $\widehat{y}_{s}=\widehat{y}_{t}$ for all rounds $s \geq t \geq T$.

We now employ the law of large numbers in order to understand the asymptotic behavior of the online learner:

$$
\operatorname{Pr}\left[\underset{(x, y) \sim P}{\operatorname{Pr}}\left[\widehat{y}_{t}(x) \neq y\right]=0\right]=\operatorname{Pr}\left[\lim _{S \rightarrow \infty} \frac{1}{S} \sum_{s=t+1}^{t+S} 1\left\{\widehat{y}_{t}\left(X_{s}\right) \neq Y_{s}\right\}=0\right]
$$

and this probability is at least the probability of this event and of the event that $T \leq t$, i.e.,

$$
\operatorname{Pr}\left[\underset{(x, y) \sim P}{\operatorname{Pr}}\left[\widehat{y}_{t}(x) \neq y\right]=0\right] \geq \operatorname{Pr}\left[\lim _{S \rightarrow \infty} \frac{1}{S} \sum_{s=t+1}^{t+S} 1\left\{\widehat{y}_{t}\left(X_{s}\right) \neq Y_{s}\right\}=0, T \leq t\right]=\operatorname{Pr}[T \leq t]
$$

where the last inequality follows from the observation that since $s \geq t$ and $t$ is greater than the critical time $T$ then the first event occurs with probability one. This implies that

$$
\operatorname{Pr}\left[\underset{(x, y) \sim P}{\operatorname{Pr}}\left[\widehat{y}_{t}(x) \neq y\right]=0\right] \geq \operatorname{Pr}[T \leq t]
$$

and so

$$
\operatorname{Pr}\left[\underset{(x, y) \sim P}{\mathbf{P r}}\left[\widehat{y}_{t}(x) \neq y\right]>0\right] \leq \lim _{t \rightarrow \infty} 1-\operatorname{Pr}[T \leq t]=0 .
$$

The above result guarantees that the expected error of the learning algorithm tends to zero as $t$ goes to infinity, i.e., we have established that that $\mathbf{E}\left[\operatorname{er}\left(\widehat{y}_{t}\right)\right] \rightarrow 0$ as $t \rightarrow \infty$. This means that the ordinal SOA is a consistent algorithm in the statistical setting. However, this fact is not enough to establish the exponential convergence rate. We follow the approach of $\left[\mathrm{BHM}^{+} 21\right]$ to come up with an algorithm that achieves this guarantee. The first step is to observe that there is some distribution-dependent $t^{\star}$ such that $\operatorname{Pr}\left[\operatorname{er}\left(\widehat{y}_{t^{\star}}\right)>0\right]<1 / 4$. If we were to know this $t^{\star}$ one way to get the exponential rates is the following: We divide the training set into $\Theta\left(n / t^{\star}\right)$ batches and we get one classifier $\widehat{y}_{t^{\star}}^{i}$ for every batch. Afterwards, we output as classifier the multiclass majority vote among the classifiers.
Theorem 10 (Exponential Rates). Assume that class $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ does not have an infinite multiclass Littlestone tree. Then, $\mathcal{H}$ admits a learning algorithm that achieves an exponentially fast rate.

Proof. Consider the sequence $\widehat{t}_{n}$ which satisfies the properties of Lemma 6. Consider the collection of the learners $\mathbb{A}_{i, n}:=\widehat{y}_{\widehat{t}_{n}}^{i}: \mathcal{X} \rightarrow[k]$ for any $1 \leq i \leq\left\lfloor\frac{n}{2 \widehat{t_{n}}}\right\rfloor$ and $n \in \mathbb{N}$. Let us fix a time $t \in T^{\star}$, where $T^{\star}$ is the set of good estimates of the critical time $t^{\star}$ (see Lemma 6). We have that

$$
\operatorname{Pr}\left[\frac{1}{\left\lfloor\frac{n}{2 t}\right\rfloor} \sum_{i=1}^{\left\lfloor\frac{n}{2 t}\right\rfloor} \mathbf{1}\left\{\underset{(x, y) \sim P}{\mathbf{P r}}\left[\widehat{y}_{t}^{i}(x) \neq y\right]>0\right\}>\frac{7}{16}\right] \leq \exp \left(-\left\lfloor n /\left(2 t^{\star}\right)\right\rfloor / 128\right)
$$

using Hoeffding's inequality. The above probability is over the sequence of all the training sets and essentially states that the "bad" event that the misclassification error is non-zero holds for the majority of the trained algorithms $\widehat{y}_{t}^{i}$ with exponentially small probability. Conversely, except on an event of exponentially small probability, we have that $\operatorname{Pr}_{(x, y) \sim P}\left[\widehat{y}_{t}^{i}(x) \neq y\right]=0$ for the majority of $i$. Recall that the above discussion holds for a chosen good $t$. We have to understand how well our learners $\mathbb{A}_{i, n}$ perform. To this end, we have that

$$
\begin{aligned}
& \operatorname{Pr}\left[\underset{(x, y) \sim P}{\mathbf{P r}}\left[\mathbb{A}_{i, n}(x) \neq y\right]>0 \text { for the majority of } i \leq\left\lfloor n /\left(2 \widehat{t}_{n}\right)\right\rfloor\right] \\
& \leq \mathbf{P r}\left[\widehat{t}_{n} \notin T^{\star}\right]+\mathbf{P r}\left[\exists t \in T^{\star}: \underset{(x, y) \sim P}{\mathbf{P r}}\left[\widehat{y}_{t}^{i}(x) \neq y\right]>0 \text { for the majority of } i \leq\left\lfloor n /\left(2 \widehat{t}_{n}\right)\right\rfloor\right] \\
& \leq C(P) \exp (-c(P) \cdot n)+t^{\star} \cdot \exp \left(-\left\lfloor n /\left(2 t^{\star}\right)\right\rfloor / 128\right) .
\end{aligned}
$$

This implies that the majority of our learners $\mathbb{A}_{i, n}$ will not incur loss except on an event of exponentially small probability. As a result, the majority vote of these classifiers is almost surely correct on a random sample from the distribution $P$ over $\mathcal{X} \times[k]$. Hence, we have that

$$
\begin{aligned}
\mathbf{E}\left[\underset{(x, y) \sim P}{\mathbf{P r}}\left[\operatorname{Maj}\left(\left(\mathbb{A}_{i, n}(x)\right)_{i}\right) \neq y\right]\right] & \leq \operatorname{Pr}\left[\underset{(x, y) \sim P}{\operatorname{Pr}}\left[\operatorname{Maj}\left(\left(\mathbb{A}_{i, n}(x)\right)_{i}\right) \neq y\right]>0\right] \\
& \leq C(P) \exp (-c(P) \cdot n)+t^{\star} \cdot \exp \left(-\left\lfloor n /\left(2 t^{\star}\right)\right\rfloor / 128\right)
\end{aligned}
$$

This concludes the proof.

Lemma 6. For any $n \in \mathbb{N}$, there exists a universally measurable $\widehat{t}_{n}=\widehat{t}_{n}\left(X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}\right)$ whose definition does not depend on $P$ so that the following holds. Set the critical time $t^{\star} \in \mathbb{N}$ be such that

$$
\operatorname{Pr}\left[\underset{(x, y) \sim P}{\operatorname{Pr}}\left[\widehat{y}_{t^{\star}}(x) \neq y\right]>0\right] \leq 1 / 8
$$

where the probability is over the training set of the algorithm $\widehat{y}_{t}$. There exist $C, c>0$ that depend on $P, t^{\star}$ but not $n$ so that

$$
\operatorname{Pr}\left[\widehat{t}_{n} \in T^{\star}\right] \geq 1-C e^{-c n}
$$

where the probability is over the training of the estimator $\widehat{t}_{n}$ and $T^{\star}$ is the set

$$
T^{\star}=\left\{1 \leq t \leq t^{\star}: \operatorname{Pr}\left[\underset{(x, y) \sim P}{\mathbf{P r}}\left[\widehat{y}_{t}(x) \neq y\right]>0\right] \leq 3 / 8\right\}
$$

where the probability is over the training of $\widehat{y}_{t}$.
Proof. We split the training set into two sets. The idea is to use the first one to train the learning algorithm and the other set to estimate the generalization error. For each $1 \leq t \leq\left\lfloor\frac{n}{2}\right\rfloor$ and $1 \leq i \leq\left\lfloor\frac{n}{2 t}\right\rfloor$, we let

$$
\widehat{y}_{t}^{i}(x)=\widehat{Y}_{t+1}\left(X_{(i-1) t+1}, Y_{(i-1) t+1}, X_{i t}, Y_{i t}, x\right)
$$

be the output of the learning algorithm that is trained on batch $i$ of the data. For every fixed $t$, the data that the classifiers $\left\{\widehat{y}_{t}^{i}\right\}_{i \leq\lfloor n / 2 t\rfloor}$ are trained on are independent of each other and of the second half of the training set. This means that we can view every $\left\{\widehat{y}_{t}^{i}\right\}_{i \leq\lfloor n / 2 t\rfloor}$ as an independent draw of the distribution of $\widehat{y}_{t}$. To estimate the generalization error of the algorithm we use the second half of the training set. We let

$$
\hat{e}_{t}=\frac{1}{\lfloor n / 2 t\rfloor} \sum_{i=1}^{\lfloor n / 2 t\rfloor} \mathbf{1}\left\{\widehat{y}_{t}^{i}\left(X_{s}\right) \neq Y_{s} \text { for some } n / 2 \leq s \leq n\right\}
$$

Now observe that, with probability one,

$$
\hat{e}_{t} \leq e_{t}=\frac{1}{\lfloor n / 2 t\rfloor} \sum_{i=1}^{\lfloor n / 2 t\rfloor} \mathbf{1}\left\{\underset{(X, Y) \sim P}{\mathbf{P r}}\left[\widehat{y}_{t}^{i}(X) \neq Y\right]>0\right\}
$$

We define $\widehat{t}_{n}=\inf \left\{t \leq\lfloor n / 2\rfloor: \widehat{e}_{t}<1 / 4\right\}$, where we assume that $\inf \emptyset=\infty$.
We now want to bound the probability that $\widehat{t_{n}}>t^{\star}$. Using Hoeffding's inequality we get that

$$
\begin{aligned}
\operatorname{Pr}\left[\widehat{t}_{n}>t^{\star}\right] \leq \operatorname{Pr}\left[\widehat{e}_{t^{\star}} \geq \frac{1}{4}\right] \leq & \operatorname{Pr}\left[e_{t^{\star}} \geq \frac{1}{4}\right]= \\
& =\operatorname{Pr}\left[e_{t^{\star}}-\frac{1}{8} \geq \frac{1}{8}\right]=\operatorname{Pr}\left[e_{t^{\star}}-\mathbf{E}\left[e_{t^{\star}}\right] \geq \frac{1}{8}\right] \leq e^{-\left\lfloor n / 2 t^{\star}\right\rfloor / 32}
\end{aligned}
$$

This implies that $\widehat{t_{n}} \leq t^{\star}$ except for an event with exponentially small probability.
Moreover, for all $1 \leq t \leq t^{\star}$ that $\operatorname{Pr}\left[\operatorname{Pr}_{(x, y) \sim P}\left[\widehat{y}_{t}(x) \neq y\right]>0\right]>\frac{3}{8}$, there is some $\epsilon>0$ such that $\operatorname{Pr}\left[\operatorname{Pr}_{(x, y) \sim P}\left[\widehat{y}_{t}(x) \neq y\right]>\epsilon\right]>\frac{1}{4}+\frac{1}{16}$ (this holds by continuity). Now fix some $1 \leq t \leq t^{\star}$ such that $\operatorname{Pr}\left[\operatorname{Pr}_{(x, y) \sim P}\left[\widehat{y}_{t}(x) \neq y\right]>0\right]>\frac{3}{8}$ (if it exists). Then, using Hoeffding's inequality again we get that

$$
\operatorname{Pr}\left[\frac{1}{\lfloor n / 2 t\rfloor} \sum_{i=1}^{\lfloor n / 2 t\rfloor} \mathbf{1}\left\{\underset{(x, y) \sim P}{\operatorname{Pr}}\left[\widehat{y}_{t}^{i}(x) \neq y\right]>\epsilon\right\}<\frac{1}{4}\right] \leq e^{\left\lfloor n / 2 t^{\star}\right\rfloor / 128}
$$

Whenever $f$ is a function such that $\operatorname{Pr}_{(x, y) \sim P}[f(x) \neq y]>\epsilon$, then

$$
\operatorname{Pr}\left[f\left(X_{s}\right) \neq Y_{s} \text { for some } n / 2 \leq s \leq n\right] \geq 1-(1-\epsilon)^{n / 2}
$$

As we mentioned before, $\left\{\widehat{y}_{t}^{i}\right\}_{i \leq\lfloor n / 2 t\rfloor}$ are independent of $\left(X_{s}, Y_{s}\right)_{s>n / 2}$. Thus, applying a union bound we get that the probability that all $\widehat{y}_{t}^{i}$ that have $\operatorname{Pr}_{(x, y) \sim P}\left[\widehat{y}_{t}^{i}(x) \neq y\right]>\epsilon$ make at least one error on the second half of the training set is
$\operatorname{Pr}\left[\sum_{i=1}^{\lfloor n / 2 t\rfloor} \mathbf{1}\left\{\underset{(x, y) \sim P}{\operatorname{Pr}}\left[\widehat{y}_{t}^{i}(x) \neq y\right]>\epsilon\right\} \leq \sum_{i=1}^{\lfloor n / 2 t\rfloor} \mathbf{1}\left\{\widehat{y}_{t}^{i}\left(X_{s}\right) \neq Y_{s}\right.\right.$ for some $\left.\left.n / 2<s \leq n\right\}\right] \geq 1-\left\lfloor\frac{n}{2 t}\right\rfloor(1-\epsilon)^{n / 2}$.
Thus, we get that

$$
\operatorname{Pr}\left[\widehat{t}_{n}=t\right] \leq \operatorname{Pr}\left[\widehat{e}_{t}<\frac{1}{4}\right] \leq\left\lfloor\frac{n}{2}\right\rfloor(1-\epsilon)^{n / 2}+e^{-\left\lfloor\frac{n}{2 t^{\star}}\right\rfloor / 128}
$$

Using the previous estimates and applying a union bound, we get that

$$
\operatorname{Pr}\left[\widehat{t_{n}} \notin T^{\star}\right] \leq e^{-\left\lfloor n / 2 t^{\star}\right\rfloor / 32}+t^{\star}\left\lfloor\frac{n}{2}\right\rfloor(1-\epsilon)^{n / 2}+t^{\star} e^{-\left\lfloor n / 2 t^{\star}\right\rfloor / 128} \leq C e^{-c n}
$$

for some constants $C, c>0$. Note that $C=C\left(P, t^{\star}\right)$ and $c=c\left(P, t^{\star}\right)$.

## C. 2 Infinite Multiclass Littlestone Trees and Rates

We next show that if $\mathcal{H}$ has an infinite mutliclass Littlestone tree, then there exists a significant drop in the rate: any learning algorithm $\mathcal{H}$ cannot be faster than linear. Our proof follows the approach in $\left[\mathrm{BHM}^{+} 21\right]$.
Theorem 11. Assume that $\mathcal{H} \subseteq\{0,1, \ldots, k\}^{\mathcal{X}}$ has an infinite multiclass Littlestone tree. Then, for any learning algorithm $\widehat{h}_{n}$, there exists a realizable distribution $P$ over $\mathcal{X} \times[k]$ such that $\mathbf{E}\left[\operatorname{er}\left(\widehat{h}_{n}\right)\right] \geq \Omega(1 / n)$ for infinitely many $n$. This means that $\mathcal{H}$ is not learnable at rate faster than linear, i.e., $R(n) \geq 1 / n$.

Let us provide some intuition. We will make use of the probabilistic method. We are going to define a distribution over probability distributions so that, with positive probability over this choice of the random object, any learning algorithm will have an expected error of order $\Omega(1 / n)$. This positive probability implies that there exists such a distribution and hence the above result holds true. The key idea is that we are going to associate any random distribution $P_{\boldsymbol{y}}$ with a branch $\boldsymbol{y}$ of the infinite multiclass Littlestone tree. Given a finite number of samples from this distribution, only a finite part of this infinite path will be discovered; hence any algorithm after the "revealed path" must guess whether this random path goes left or right. This implies that the algorithm will err with probability $1 / 2$ when it observes a point that lies deeper in the branch than the training examples.

Proof. Fix any learner $\widehat{h}_{n}$ and an infinite multiclass Littlestone tree for $\mathcal{H}$. Fix also a random branch $\boldsymbol{y}=\left(y_{1}, y_{2}, \ldots\right)$ of this tree, where the sequence is an i.i.d. sequence of fair Bernoulli coins. We introduce the random distribution over $\mathcal{X} \times\{0,1, \ldots, k\}$ as

$$
P_{\boldsymbol{y}}\left(\left(x_{\boldsymbol{y}_{\leq \ell}}, z_{\ell+1}\right)\right)=\frac{1}{2^{\ell+1}}, \ell \geq 0
$$

where $z_{\ell+1} \in\{0,1, \ldots, k\}$ is the label of the edge connecting $x_{\boldsymbol{y}_{\leq \ell}}$ to its child according to the chosen path $\boldsymbol{y}$. For any $n<\infty$, there exists a hypothesis $h \in \mathcal{H}$ so that

$$
h\left(x_{\boldsymbol{y}_{\leq \ell}}\right)=z_{\ell+1}
$$

for $0 \leq \ell \leq n$. This is due to the construction of a multiclass Littlestone tree. We have that

$$
\operatorname{er}_{\boldsymbol{y}}(h)=\underset{(x, z) \sim P_{y}}{\mathbf{P r}_{y}}[h(x) \neq z] \leq \sum_{\ell>n} 2^{-\ell-1}
$$

which goes to 0 as $n \rightarrow \infty$. This implies that $P_{\boldsymbol{y}}$ is realizable for every infinite branch $\boldsymbol{y} \in\{0,1\}^{\infty}$. Let us draw $(X, Z),\left(X_{1}, Z_{1}\right),\left(X_{2}, Z_{2}\right), \ldots$ i.i.d. samples from $P_{\boldsymbol{y}}$. Moreover, the mapping $y \rightarrow P_{y}$ is measurable. The first sample corresponds to the test sample and the other samples deal with the training phase. Moreover, let $T, T_{1}, T_{2}, \ldots$ be i.i.d. Geometric random variables with success probability $1 / 2$ starting at 0 . We can set

1. $X=x_{\boldsymbol{y} \leq T}, Z=z_{T+1}$ and
2. $X_{i}=x_{\boldsymbol{y} \leq T_{i}}, Z_{i}=z_{T_{i}+1}$.

Crucially, on the event that $\left\{T=\ell, \max \left\{T_{1}, \ldots, T_{n}\right\}<\ell\right\}$, the value of $\widehat{h}_{n}(X)$ is conditionally independent of $z_{\ell+1}$ given $X,\left(X_{1}, Z_{1}\right), \ldots,\left(X_{n}, Z_{n}\right)$. We next have that
$\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z, T=\ell, \max \left\{T_{1}, \ldots, T_{n}\right\}<\ell\right]=\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z_{\ell+1}, T=\ell, \max \left\{T_{1}, \ldots, T_{n}\right\}<\ell\right]$.
This is equal to

$$
\mathbf{E}\left[\mathbf{P r}_{Z}\left[\widehat{h}_{n}(X) \neq Z \mid X,\left(X_{1}, Z_{1}\right), \ldots,\left(X_{n}, Z_{n}\right)\right] \mathbf{1}\left\{T=\ell, \max \left\{T_{1}, \ldots, T_{n}\right\}<\ell\right\}\right]
$$

Now conditional on this event, any algorithm will err with probability $1 / 2$ (since it will guess the value at random). Hence, this quantity is lower bounded by

$$
\frac{1}{2} \operatorname{Pr}\left[T=\ell, \max \left\{T_{1}, \ldots, T_{n}\right\}<\ell\right]=2^{-\ell-2}\left(1-2^{-\ell}\right)^{n}
$$

We are free now to pick $\ell$. Choosing $\ell=\ell_{n}:=\lceil 1+\log (n)\rceil$, we have that $1 / 2^{\ell}>1 /(4 n)$ and $\left(1-2^{-\ell}\right)^{n} \geq 1 / 2$. Our goal is to apply the reverse Fatou lemma. This can be done since almost surely, we have that

$$
n \mathbf{P r}\left[\widehat{h}_{n}(X) \neq Z, T=\ell_{n} \mid \boldsymbol{y}\right] \leq n \mathbf{P r}\left[T=\ell_{n}\right]=n 2^{-\ell_{n}-1} \leq 1 / 4
$$

Hence, we can apply the reverse Fatou lemma and get

$$
\mathbf{E}\left[\limsup _{n \rightarrow \infty} n \operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z, T=\ell_{n} \mid \boldsymbol{y}\right]\right] \geq \limsup _{n \rightarrow \infty} n \operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z, T=\ell_{n}\right]>1 / 32
$$

But, almost surely, it holds that

$$
\mathbf{E}\left[\mathrm{er}_{\boldsymbol{y}}\left(\widehat{h}_{n}\right) \mid \boldsymbol{y}\right]=\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z \mid \boldsymbol{y}\right] \geq \operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z, T=\ell_{n} \mid \boldsymbol{y}\right]
$$

So, combining the above inequalities

$$
\mathbf{E}\left[\limsup _{n \rightarrow \infty} n \mathbf{E}\left[\operatorname{er}_{\boldsymbol{y}}\left(\widehat{h}_{n}\right)\right]\right]>\Omega(1)
$$

Hence, there must exist a realization of $\boldsymbol{y}$ so that $\mathbf{E}\left[\operatorname{er}_{\boldsymbol{y}}\left(\widehat{h}_{n}\right)\right]=\Omega(1 / n)$ infinitely often. Choosing $P=P_{\boldsymbol{y}}$ completes the proof.

The above result states that a class which does not have exponential rates cannot be learned faster than linearly. However, it is not clear if even linear rates are achievable. The next section deals with this case.

## C. 3 Linear Learning Rates and Natarajan-Littlestone Trees

We introduce a novel combinatorial measure, the Natarajan-Littlestone (NL) tree, which essentially combines the structure of the Natarajan dimension and the Littlestone dimension.
Definition 12. A Natarajan-Littlestone (NL) tree for $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ of depth $d \leq \infty$ consists of a tree

$$
\bigcup_{0 \leq \ell<d}\left\{x_{u} \in \mathcal{X}^{\ell+1}, u \in\{0,1\} \times\{0,1\}^{2} \times \ldots \times\{0,1\}^{\ell}\right\}
$$

and two colorings $s^{(0)}, s^{(1)}$ mapping each position $u^{i} \in u$ for any node with pattern $u \in\{0,1\} \times$ $\ldots \times\{0,1\}^{\ell}$ for $i \in\{0,1, \ldots, \ell\}$ and $\ell \in\{0,1, \ldots, d-1\}$ of the tree to some color $\{0,1, \ldots, k\}$ such that for every finite level $n<d$, the subtree $T_{n}=\cup_{0 \leq \ell \leq n}\left\{x_{u}=\left(x_{u}^{0}, \ldots, x_{u}^{\ell}\right): u \in\{0,1\} \times\right.$ $\left.\{0,1\}^{2} \times \ldots \times\{0,1\}^{\ell}\right\}$ satisfies the following:

1. At any point $x_{u}^{i} \in x_{u} \in T_{n}$, it holds $s^{(0)}\left(x_{u}^{i}\right) \neq s^{(1)}\left(x_{u}^{i}\right)$ and
2. for any path $\boldsymbol{y} \in\{0,1\} \times \ldots \times\{0,1\}^{n+1}$, there exists a concept $h \in \mathcal{H}$ so that $h\left(x_{\boldsymbol{y}_{\leq \ell}}^{i}\right)=$ $s^{(0)}\left(x_{\boldsymbol{y}_{\leq \ell}}^{i}\right)$ if $y_{\ell+1}^{i}=1$ and $h\left(x_{\boldsymbol{y}_{\leq \ell}}^{i}\right)=s^{(1)}\left(x_{\boldsymbol{y}_{\leq \ell}}^{i}\right)$ otherwise, for all $0 \leq i \leq \ell$ and $0 \leq \ell \leq n$, where

$$
\boldsymbol{y}_{\leq \ell}=\left(y_{1}^{0},\left(y_{2}^{0}, y_{2}^{1}\right), \ldots,\left(y_{\ell}^{0}, \ldots, y_{\ell}^{\ell-1}\right)\right), x_{\boldsymbol{y}_{\leq \ell}}=\left(x_{\boldsymbol{y}_{\leq \ell}}^{0}, \ldots, x_{\boldsymbol{y}_{\leq \ell}^{\ell}}^{\ell}\right) .
$$

We say that $\mathcal{H}$ has an infinite $N L$ tree if it has a $N L$ tree of depth $d=\infty$.
We note that in the above definition we identify the color $s^{(0)}\left(x_{\boldsymbol{y} \leq \ell}^{i}\right)$ with the (unique) position of this point $x_{\boldsymbol{y} \leq \ell}^{i}$ (since typically the coloring is over positions). For short, we will call the colorings $s^{(0)}, s^{(1)}$ used in the above definition everywhere different and denote by $s^{(0)} \neq s^{(1)}$. As a sanity check, one can verify that if $\mathcal{H}$ has an infinite NL tree, then it has an infinite multiclass Littlestone tree. In fact, one can construct the latter tree by choosing only one point from any node of the infinite NL tree.

The main result of this section is the next theorem.
Theorem 12. Assume that $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ does not have an infinite Natarajan-Littlestone tree. Then, there exists an algorithm that learns $\mathcal{H}$ at a linear rate.

## C.3.1 The Natarajan-Littlestone Game

In a similar manner as in the exponential rates setting, we need to come up with a function that is correct after a finite number of steps $n^{*}$ and this implies that some appropriate data-dependent complexity measure of $\mathcal{H}$ is finite. The first step towards establishing the desired result is to introduce the following two-player game between an adversary and a learning player.

1. The adversary picks $t$ points and two everywhere different colorings of these $t$ points $\xi_{t}=\left(\xi_{t}^{(0)}, \ldots, \xi_{t}^{(t-1)}, s_{t}^{(0)}, s_{t}^{(1)}\right) \in \mathcal{X}^{t} \times[k]^{t} \times[k]^{t}$ and reveals them to the learner.
2. The learner chooses a pattern $\eta_{t}=\left(\eta_{t}^{(0)}, \ldots, \eta_{t}^{(t-1)}\right) \in\{0,1\}^{t}$.

Figure 6: Adversarial Setting - 2-Player NL Game

Let $\Xi_{t}=\left\{\xi_{1}^{(0)}, \xi_{2}^{(0)}, \xi_{2}^{(1)}, \ldots, \xi_{t}^{(0)}, \ldots, \xi_{t}^{(t-1)}, s_{1}^{(0)}, s_{1}^{(1)}, \ldots, s_{t}^{(0)}, s_{t}^{(1)}\right\}$ be the sets of all points and colorings chosen by the adversary after $t$ rounds, for some finite integer $t$. The learning player wins the game in round $t$ if the class $\mathcal{H}_{t}=\mathcal{H}_{\xi_{1}, \eta_{1}, \ldots, \xi_{t}, \eta_{t}}=\emptyset$, where $\mathcal{H}_{0}=\mathcal{H}$ and $\mathcal{H}_{t}=\mathcal{H}_{\xi_{1}, \eta_{1}, \ldots, \xi_{t}, \eta_{t}}$ is the set

$$
\left\{h \in \mathcal{H} \text { so that } \begin{array}{l}
h\left(\xi_{z}^{(i)}\right)=s_{z}^{(0)}\left(\xi_{z}^{(i)}\right) \text { if } \eta_{z}^{(i)}=0 \\
h\left(\xi_{z}^{(i)}\right)=s_{z}^{(1)}\left(\xi_{z}^{(i)}\right) \text { if } \eta_{z}^{(i)}=1
\end{array} \text { for } 0 \leq i<z, z \in[1 . . t]\right\} .
$$

Note that the collection of valid functions decreases as the game proceeds, i.e., $\mathcal{H}_{t} \subseteq \mathcal{H}_{t-1}$. The next lemma guarantees the existence of a winning strategy for the learner in the game.
Lemma 7. If $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ has no infinite $N L$ tree, then there is a universally measurable winning strategy for the learning player in the game $\mathcal{G}$.

Proof. The first point is that the 2-player game of Figure 6 is Gale-Stewart. This follows by the observation that the set of winning sequences of the learning player is finitely decidable, since the membership of $(\boldsymbol{\xi}, \boldsymbol{\eta})$ in the set is witnessed by a finite subsequence. Hence, exactly one of the two players has a winning strategy in the game (see Appendix B.6).
The second point is that the class $\mathcal{H}$ has an infinite NL tree if and only if the adversary player has a winning strategy in the game. Suppose that $\mathcal{H}$ has an infinite NL tree. The adversary can adopt the strategy iteratively by setting $\xi_{t}\left(\eta_{1}, \ldots, \eta_{t-1}\right)=\left(x_{\eta_{1}, \ldots, \eta_{t-1}}, s_{\eta_{1}, \ldots, \eta_{t-1}}^{(0)}, s_{\eta_{1}, \ldots, \eta_{t-1}}^{(1)}\right) \in$ $\mathcal{X}^{t} \times[k]^{t} \times[k]^{t}$. Hence, the adversary traverses the infinite tree and, by construction of the NL tree, the set $\mathcal{H}_{\xi_{1}, \eta_{1}, \ldots, \xi_{t}, \eta_{t}}$ never gets empty for any sequence of patterns $\boldsymbol{\eta}$ and $t<\infty$. Thus, this is a winning strategy for the adversary player. In the opposite direction, assume that the adversary has a winning
strategy. Then, she can construct an infinite NL tree by setting $\left(x_{\eta_{1}, \ldots, \eta_{t-1}}, s_{\eta_{1}, \ldots, \eta_{t-1}}^{(0)}, s_{\eta_{1}, \ldots, \eta_{t-1}}^{(1)}\right)=$ $\xi_{t}\left(\eta_{1}, \ldots, \eta_{t-1}\right)$ for any possible binary pattern. Note that since the strategy is winning, the class $\mathcal{H}_{\xi_{1}, \eta_{1}, \ldots, \xi_{t}, \eta_{t}} \neq \emptyset$ for any $t \in \mathbb{N}$ and this means that for any level of the tree, there exists two everywhere different colorings that witness the patterns.

Using the second point, we have that the learning player has a winning strategy in the game if the class $\mathcal{H}$ has no infinite NL tree. To show that this strategy is also universally measurable, it suffices to show that the set of winning sequences for the learning player is coanalytic. The set of winning strategies for the adversary is
$\mathcal{W}^{c}=\left\{\left(\left(\boldsymbol{\xi}, \boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}\right) \in \bigcup_{t=1}^{\infty}\left(\left(\mathcal{X}^{t} \times[k]^{t} \times[k]^{t}\right) \times\{0,1\}^{t}\right): \mathcal{H}_{\xi_{1}, \eta_{1}, \ldots, \xi_{t}, \eta_{t}} \neq \emptyset\right.$ for all $\left.t<\infty\right\}$.
The set $\mathcal{W}^{c}$ is equal to

$$
\bigcap_{1 \leq t<\infty} \bigcup_{\theta \in \Theta} \bigcap_{1 \leq \ell \leq t}\left\{\left(\left(\boldsymbol{\xi}, \boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}\right): \begin{array}{ll}
h\left(\theta, \xi_{z}^{(i)}\right)=s_{z}^{(0)}\left(\xi_{z}^{(i)}\right) & \text { if } \eta_{z}^{(i)}=0 \\
h\left(\theta, \xi_{z}^{(i)}\right)=s_{z}^{(1)}\left(\xi_{z}^{(i)}\right) & \text { if } \eta_{z}^{(i)}=1
\end{array} \text { for } 0 \leq i<z, z \in[\ell]\right\}
$$

In words, this set contains all the states of the game, for any timestep $t$, so that there exists a parameter in $\Theta$ and some $k$-colorings $s^{(0)}, s^{(1)}$ which are everywhere different, that are $N$-consistent with some hypothesis in the class. Also, note that the set

$$
\begin{aligned}
& \left\{\left(\left(\boldsymbol{\xi}, \boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}\right): \begin{array}{l}
\left.h\left(\theta, \xi_{z}^{(i)}\right)=s_{z}^{(0)}\left(\xi_{z}^{(i)}\right) \quad \begin{array}{l}
\text { if } \eta_{z}^{(i)}=0 \\
h\left(\theta, \xi_{z}^{(i)}\right)=s_{z}^{(1)}\left(\xi_{z}^{(i)}\right) \\
\text { if } \eta_{z}^{(i)}=1
\end{array} \text { for } 0 \leq i<z, z \in[\ell]\right\}= \\
\bigcap_{1 \leq z \leq \ell} \bigcap_{0 \leq i<z}\left\{\left(\left(\boldsymbol{\xi}, \boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}\right): h\left(\theta, \xi_{z}^{(i)}\right)=s_{z}^{(0)}\left(\xi_{z}^{(i)}\right) \text { if } \eta_{z}^{(i)}=0\right\} \cap \\
\left\{\left(\left(\boldsymbol{\xi}, \boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}\right): h\left(\theta, \xi_{z}^{(i)}\right)=s_{z}^{(1)}\left(\xi_{z}^{(i)}\right) \text { if } \eta_{z}^{(i)}=1\right\}= \\
\bigcap_{1 \leq z \leq \ell} \bigcap_{0 \leq i<z}\left\{\left(\left(\boldsymbol{\xi}, \boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}\right): h\left(\theta, \xi_{z}^{(i)}\right)=s_{z}^{\left(\eta_{z}^{(i)}\right)}\left(\xi_{z}^{(i)}\right)\right\} .
\end{array} .\right.
\end{aligned}
$$

Recall that a Borel set is any set in a topological space that can be formed from open sets through the operations of countable union, countable intersection, and relative complement. By the measurability assumption of Definition 9, the set $\left\{\left(\theta,\left(\boldsymbol{\xi}, \boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}\right): h\left(\theta, \xi_{z}^{(i)}\right)=s_{z}^{\left(\eta_{z}^{(i)}\right)}\left(\xi_{z}^{(i)}\right)\right\}$ which corresponds to the subspace of the product space of $\Theta$ and the space of all infinite sequences $\left(\left(\boldsymbol{\xi}, \boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}\right)$ whose $(z, i)$-th term is consistent with some hypothesis $h$ is Borel measurable (note that the sequence $\left(\boldsymbol{s}^{(0)}, \boldsymbol{s}^{(1)}\right), \boldsymbol{\eta}$ ) uniquely induces a sequence of labels over $[k]$ on which we apply the measurability definition). Moreover, the following hold: (i) the two intersections over $t$ and $\ell$ are countable, (ii) the union over $\theta$ is a projection of a Borel set and (iii) the union over the colorings is countable (it is even finite given $t$ ). This implies that $\mathcal{W}$ is coanalytic.

## C.3.2 Pattern Avoidance Algorithm

Our learning algorithm for the linear rates case will be built over the notion of data patterns and pattern avoidance functions. In the exponential rates setting, we focused on the case where the learner is successful if she makes a finite number of mistakes. In the case of linear rates, this is replaced with controlling the model complexity, i.e., we have to understand how expressive the class $\mathcal{H}$ is. We will say that a sequence $x_{1}, y_{1}, x_{2}, y_{2}, \ldots$ in $(\mathcal{X} \times[k])^{\infty}$ is consistent with $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ if, for every finite $t<\infty$, there exists $h \in \mathcal{H}$ such that $h\left(x_{i}\right)=y_{i}$ for all $i \leq t$.
We next introduce the crucial notion of pattern. Intuitively, the expressivity of a hypothesis class $\mathcal{H}$ is proportional to the length of the "patterns" realized by concepts in $\mathcal{H}$. Hence, by controlling the length of the realizable patterns, we can obtain learners for $\mathcal{H}$.
Definition 13. Given a sequence $S=\left(x_{1}, y_{2}, x_{2}, y_{2}, \ldots\right) \in(\mathcal{X} \times[k])^{\infty}$ that is consistent with $\mathcal{H}$, the binary string $b_{t} \in\{0,1\}^{t}$ of length $t$ is a Natarajan-Littlestone pattern (or simply a pattern) if there exists a subsequence of consecutive terms $S^{\prime}=\left(x_{z+i}, y_{z+i}\right)_{i \in[t]}$ of length $t$ for some $z \in \mathbb{N}$ so that there exist two $k$-colorings everywhere different $s^{(0)}, s^{(1)}$ of the elements $\left(x_{z+i}\right)_{i \in[t]}$ with $s^{(0)}\left(x_{z+i}\right)=y_{z+i}$ if $b_{t}=0$ and $s^{(1)}\left(x_{z+i}\right)=y_{z+i}$ if $b_{t}=1$.

We are going to use the winning strategy of the learning player of the game of Figure 6 in order to design an algorithm that avoids NL patterns in the data. This is exactly the intuition behind the design of the Gale-Stewart game of Figure 6. The learner predicts a binary string $\eta_{t} \in\{0,1\}^{t}$ with goal that this pattern is not realized by the class $\mathcal{H}$, i.e., it is not a NL pattern with respect to a sequence consistent with $\mathcal{H}$. Intuitively, the reason we are aiming for forbidden patterns is the following: before witnessing the realizable sequence we have no control over the complexity of the hypothesis class $\mathcal{H}$, since its Natarajan dimension can be infinite. However, leveraging the fact that there is no infinite NL-tree we show that for every realizable sequence $S$ there is some (sequence-dependent) $n^{*}(S)$ such that the class $\mathcal{H}_{x_{1}, y_{1}, \ldots, x_{n}, y_{n^{*}}}$ of the concepts that agree on the first $n^{*}$ terms has Natarajan dimension that is bounded by $O\left(n^{*}\right)$. This is because if we can produce a forbidden pattern for every $x$-tuple of length $n^{*}$ we know that the Natarajan dimension of this class is bounded by $n^{*}$.

The algorithm will get as input a data sequence consistent with $\mathcal{H}$ and will identify patterns in the data. In particular, the algorithm works as follows: For any finite pattern length $t$, the algorithm traverses the data sequence in consecutive blocks of length $t$ and in each such block it tries all the everywhere different colorings for the points and uses the universally measurable winning strategy of the learning player to obtain a guess for a forbidden NL pattern for each of these colorings. The algorithm then checks if at least one of the pairs of colorings and forbidden patterns are actually realized by the data. If this is the case (i.e., the guess was false), the algorithm adds a new point $\xi_{t}$ in its list of "bad points" and continues searching for patterns of larger length. If the guess was correct and the pattern is not realized in this block, it continues with the same pattern size in the next block.

The algorithm operates as follows. We remark that we can construct the strategy $\eta_{t}$ of the learner by invoking a notion of ordinal NL dimension in a similar manner as in the exponential rates case.

1. Initialize the pattern length $\ell_{0}=1$.
2. At every time step $t \geq 1$ :
(a) Set $L:=\ell_{t-1}$ and create a list $S_{L}$ of all the possible pairs of everywhere different colorings of $L$ elements.
(b) Set $j:=1$ be the index that traverses the list $S_{L}$.
(c) Set $s^{(0)}:=S_{L}[j](0), s^{(1)}:=S_{L}[j](1)$.
(d) Compute pattern $b_{L}=\eta_{L}\left(\xi_{1}, \ldots, \xi_{L-1}, x_{t-L+1}, \ldots, x_{t}, s^{(0)}, s^{(1)}\right) \in\{0,1\}^{L} \times$ $[k]^{L} \times[k]^{L}$.
(e) If it holds that, $s_{z}^{(0)}\left(x_{t-L+1+z}\right)=y_{t-L+1+z}$, if $b_{L}(z)=0$, and $s_{z}^{(1)}\left(x_{t-L+1+z}\right)=y_{t-L+1+z}$, otherwise, for $z \in\{0,1, \ldots, L-1\}$, then i. Set $\xi_{L}=\left(x_{t-L+1}, \ldots, x_{t}, s^{(0)}, s^{(1)}\right) / /$ The pattern is realized (not forbidden). ii. Update $\ell_{t}=L+1 / /$ Look for larger patterns.
(f) Else:
i. If $j<L$ then update $j:=j+1$ and go to step (c).
ii. Else shift the block, and move to the next timestep, i.e. $\ell_{t}=L, t:=t+1$.

Figure 7: Pattern Avoidance Algorithm

Hence, the above procedure defines a pattern avoidance function

$$
\begin{equation*}
\widehat{y}_{t-1}\left(x_{1}, \ldots, x_{\ell_{t-1}}, s^{(0)}, s^{(1)}\right):=\eta_{\ell_{t-1}}\left(\xi_{1}, \ldots, \xi_{\ell_{t-1}-1}, x_{1}, \ldots, x_{\ell_{t-1}}, s^{(0)}, s^{(1)}\right) \in\{0,1\}^{\ell_{t-1}} \times[k]^{\ell_{t-1}} \times[k]^{\ell_{t-1}} . \tag{1}
\end{equation*}
$$

The intuition is that as long as the class $\mathcal{H}$ does not have an infinite NL tree and a consistent sequence is provided to the algorithm, then whenever the "if" statement gets True, the game of Figure 6 progresses and the learner gets closer to winning the game. The finiteness of the NL tree gives the next key lemma.

Lemma 8. For any sequence $x_{1}, y_{1}, x_{2}, y_{2}, \ldots$ that is consistent with $\mathcal{H}$, the pattern avoidance algorithm of Figure 7, in a finite number of steps, rules out NL patterns in the data sequence, in the sense that the "if" statement in Line (c) is false and $\ell_{t}=\ell_{t-1}<\infty, \widehat{y}_{t}=\widehat{y}_{t-1}$ for all sufficiently large $t$. Moreover, the mappings $\ell_{t}$ and $\widehat{y}_{t}$ are universally measurable.

Proof. Consider an infinite sequence of times $t=t_{1}, t_{2}, \ldots$ and assume that the "if" condition is true for any time-step in this sequence. Since $\eta_{t}$ is a winning strategy for the learning player in the Gale-Stewart game $\mathcal{G}$, we have that there exists some finite time $t^{\star}$ so that $\mathcal{H}_{\xi_{1}, \eta_{1}, \ldots, \xi_{t^{\star}}, \eta_{t^{\star}}}=\emptyset$. By the structure of the "if" condition, we have that $\xi_{i}=\left(x_{t_{i}-\ell_{t_{i}}-1}+1, \ldots, x_{t_{i}}\right)$ and $\eta_{i}=\left(y_{t_{i}-\ell_{t_{i}}-1}+\right.$ $1, \ldots, y_{t_{i}}$ ). But since the sequence is consistent with $\mathcal{H}$, the class $\mathcal{H}_{\xi_{1}, \eta_{1}, \ldots, \xi_{t^{\star}}, \eta_{t^{\star}}}$ must contain the hypothesis that makes it consistent which is a contradiction since the class is empty. For the measurability of the strategies, we refer to Remark 5.4 of $\left[\mathrm{BHM}^{+} 21\right]$.

## C.3.3 Asymptotic Behavior of the Algorithm

As in the case of exponential rates, our first goal towards the design of our algorithms is an asymptotic result. We first have to move from the adversarial to a probabilistic setting. Fix a realizable distribution $P$ over $\mathcal{X} \times[k]$ and let $\left(X_{1}, Y_{1}\right),\left(X_{2}, Y_{2}\right), \ldots$ be i.i.d. samples from $P$. As it is clear in this section for the upper bounds, we assume that $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ does not have an infinite NL tree.

For any integer $\ell \in \mathbb{N}$ and any universally measurable pattern avoidance function $g: \mathcal{X}^{\ell} \times[k]^{\ell} \times$ $[k]^{\ell} \rightarrow\{0,1\}^{\ell}$, we consider the error
$\operatorname{per}(g):=\operatorname{per}_{\ell}(g)=\underset{\left(x_{1}, y_{1}, \ldots, x_{\ell}, y_{\ell}\right)}{\mathbf{P r}}\left[g\right.$ fails to avoid some NL pattern realized by the data $\left.x_{1}, y_{1}, \ldots, x_{\ell}, y_{\ell}\right]$,
i.e., there exist some colorings $s^{(0)}, s^{(1)}$ everywhere different that witness the pattern $g\left(x_{1}, \ldots, x_{\ell}, s^{(0)}, s^{(1)}\right)$ with colors $\left(y_{1}, \ldots, y_{\ell}\right)$. The following lemma establishes this result. It follows the approach in $\left[\mathrm{BHM}^{+} 21\right]$.
Lemma 9. For the algorithm of Equation (1), it holds that

$$
\operatorname{Pr}\left[\operatorname{per}\left(\widehat{y}_{t}\right)>0\right] \rightarrow 0 \text { as } t \rightarrow \infty,
$$

where the probability is over the random data that are used to train the algorithm.
Proof. As in the proof of the asymptotic result in the exponential rates, since the distribution $P$ is realizable, we can get that the data sequence $X_{1}, Y_{1}, X_{2}, Y_{2}, \ldots$ is consistent with $\mathcal{H}$ almost surely We can now employ Lemma 8 and get that the random variable

$$
T=\sup \{s \geq 1: \text { the Pattern Avoidance Algorithm of Figure } 7 \text { gets in the "if" statement }\}
$$

is finite almost surely and $\widehat{y}_{s}=\widehat{y}_{t}$ and $\ell_{s}=\ell_{t}$ for all $s \geq t \geq T$. We can apply the law of large nymber for $m$-dependent sequences and get that the quantity $\operatorname{Pr}\left[\operatorname{per}_{\ell_{t}}\left(\widehat{y}_{t}\right)=0\right]$ is equal to
$\operatorname{Pr}\left[\lim _{S \rightarrow \infty} \frac{1}{S} \sum_{s=t+1}^{t+S} 1\left\{\widehat{y}_{t}\right.\right.$ fails to avoid some NL pattern realized by $\left.\left.\left(X_{s}, Y_{s}, \ldots, X_{s+\ell_{t}-1}, Y_{s+\ell_{t}-1}\right)\right\}=0\right]$.
We can decrease the probability of the right-hand side by taking the intersection of the right-hand side event with the event $\{T \leq t\}$. Hence, we have that

$$
\operatorname{Pr}\left[\operatorname{per}_{\ell_{t}}\left(\widehat{y}_{t}\right)=0\right] \geq \mathbf{P r}[T \leq t] \rightarrow 1, \text { as } t \rightarrow \infty,
$$

since $T$ is finite with probability one.

## C.3.4 Linear Learning Rates

Given some pattern avoidance function that is correct on any tuple of size $t$, we can essentially use the 1 -inclusion graph in a similar manner as if the Natarajan dimension of $\mathcal{H}$ was bounded by $t$. This is established in the following lemma.
Lemma 10 (Learning NL pattern classes). Fix $t \geq 1$. Let $g: \mathcal{X}^{t} \times[k]^{t} \times[k]^{t} \rightarrow\{0,1\}^{t}$ be the universally measurable NL pattern avoidance function of Figure 7. For any $n \geq 1$, there exists a universally measurable classifier $\widehat{Y}_{n}^{g}:(\mathcal{X} \times[k])^{n-1} \times \mathcal{X} \rightarrow[k]$ such that for every training set $\left(x_{1}, y_{1}, \ldots, x_{n}, y_{n}\right) \in(\mathcal{X} \times[k])^{n}$ where $g\left(x_{i_{1}}, \ldots, x_{i_{t}}, s^{(0)}, s^{(1)}\right)$ is an avoiding NL pattern for all pairwise distinct indices $1 \leq i_{1}, \ldots, i_{t} \leq n$ and everywhere different colorings $s^{(0)}, s^{(1)}$, the classifier achieves a linear permutation bound, i.e.,

$$
\operatorname{Pr}_{\sigma \sim \mathcal{U}\left(\mathbb{S}_{n}\right)}\left[\widehat{Y}_{n}^{g}\left(x_{\sigma(1)}, y_{\sigma(1)}, \ldots, x_{\sigma(n-1)}, y_{\sigma(n-1)}, x_{\sigma(n)}\right) \neq y_{\sigma(n)}\right] \leq \frac{t \log (k)}{n}
$$

Note that the above algorithm is transductive. This hints the use of the one-inclusion learning algortihm.

Proof. Fix $n \geq 1$ and set $X=[n]$. Also let $F$ be a set of colorings $f: X \rightarrow[k]$. We can apply the one-inclusion hypergraph algorithm (see Lemma 1) to get that there exists an algorithm $\mathbb{A}$ such that

$$
\operatorname{Pr}_{\sigma \sim \mathcal{U}\left(\mathbb{S}_{n}\right)}[\mathbb{A}(F, \sigma(1), f(\sigma(1)), \ldots, \sigma(n-1), f(\sigma(n-1)), \sigma(n)) \neq f(\sigma(n))] \leq \frac{\operatorname{Ndim}(F) \log (k)}{n}
$$

for any coloring $f \in F$ and $F \in 2^{[k]^{X}}$. By construction of the algorithm $\mathbb{A}$, the output of the mapping is preserved under relabeling of $X$, i.e.,

$$
\mathbb{A}\left(F, \sigma(1), y_{1}, \ldots, \sigma(n-1), y_{n-1}, \sigma(n)\right)=\mathbb{A}\left(F \circ \sigma, 1, y_{1}, \ldots, n-1, y_{n-1}, n\right)
$$

where $F \circ \sigma=\{f \circ \sigma: f \in F\}$. Moreover, the mapping $\mathbb{A}$ is measurable since its domain is finite (since the number of colors $k$ is a constant).
We will make use of the next result.
Claim 2. Consider the pattern avoidance mapping $g: \mathcal{X}^{t} \times[k]^{t} \times[k]^{t} \rightarrow\{0,1\}^{t}$ of Figure 7. Consider the class $F$ of $k$-colorings of $[n]$ and fix a sequence $\boldsymbol{x}=\left(x_{1}, \ldots, x_{n}\right) \in \mathcal{X}^{n}$. Fix $0<t \leq n$. Define the subset $F_{\boldsymbol{x}}$ that contains the colorings $f:[n] \rightarrow[k]$ that satisfy the following property: For all subsets $\left(i_{1}, \ldots, i_{t}\right)$ of $[n]$ of size $t$ and for all colorings $s^{(0)}, s^{(1)}:[n] \rightarrow[k]$ which are everywhere different, at least one of the next $t$ conditions is violated:

$$
\begin{aligned}
& f\left(i_{j}\right)=s^{(0)}\left(x_{i_{j}}\right) \text { if } g\left(x_{i_{1}}, \ldots, x_{i_{t}}, s^{(0)}, s^{(1)}\right)\left[x_{i_{j}}\right]=0 \\
& f\left(i_{j}\right)=s^{(1)}\left(x_{i_{j}}\right) \text { if } g\left(x_{i_{1}}, \ldots, x_{i_{t}}, s^{(0)}, s^{(1)}\right)\left[x_{i_{j}}\right]=1 .
\end{aligned}
$$

Then, we have that $\operatorname{Ndim}\left(F_{\boldsymbol{x}}\right)<t$.
Proof. It must be the case that the Natarajan dimension cannot be more than $t-1$ since we cannot $N$-shatter any $t$-subset of $[n]$.

Given any input sequence $\left(x_{1}, y_{1}, \ldots, x_{n}, y_{n}\right)$, we introduce the concept class $F_{x}$ as in Claim 2. Moreover, consider the mapping $g: \mathcal{X}^{t} \times[k]^{t} \times[k]^{t} \rightarrow\{0,1\}^{t}$ defined to be the pattern avoidance function generated by Figure 7. We can introduce a data-dependent classifier

$$
\widehat{Y}_{n}\left(g ; x_{1}, y_{1}, \ldots, x_{n-1}, y_{n-1}, x_{n}\right)=\mathbb{A}\left(F_{x}, 1, y_{1}, \ldots, n-1, y_{n-1}, n\right)
$$

Note that the above mapping is universally measurable due to Lemma 8. By relabeling we have that
$\widehat{Y}_{n}\left(g ; x_{\sigma(1)}, y_{\sigma(1)}, \ldots, x_{\sigma(n-1)}, y_{\sigma(n-1)}, x_{\sigma(n)}\right)=\mathbb{A}\left(F_{x}, \sigma(1), y_{\sigma(1)}, \ldots, \sigma(n-1), y_{\sigma(n-1)}, \sigma(n)\right)$.
By the assumption of the lemma about the input $\left(x_{1}, y_{1}, \ldots, x_{n}, y_{n}\right)$, we have that the coloring $y(i)=y_{i} \in[k]$ satisfies $y \in F_{x}$ (due to the construction of Claim 2). Hence, for such a sequence, we get that

$$
\operatorname{Pr}_{\sigma \sim \mathcal{U}\left(\mathbb{S}_{n}\right)}\left[\widehat{Y}_{n}^{g}\left(x_{\sigma(1)}, y_{\sigma(1)}, \ldots, x_{\sigma(n-1)}, y_{\sigma(n-1)}, x_{\sigma(n)}\right) \neq y_{\sigma(n)}\right] \leq \frac{\operatorname{Ndim}\left(F_{x}\right) \log (k)}{n}
$$

However, due to the conclusion of Claim 2, we have that the above rate is of order $\log (k) \cdot t / n$, since by construction $\operatorname{Ndim}\left(F_{x}\right)<t$.

To make use of the results we just stated, we need to come up with a pattern avoidance function that is correct on any tuple of size $t$. However, we can only establish that our pattern avoidance function is eventually correct, so there is no bound on $t$ that is given to the learner. To deal with this problem, we keep track of some $\widehat{t}_{n}$ such that, with high probability, the error of the pattern avoidance function is small. Then, we divide our sample into $n / \widehat{t}_{n}$ batches of size $\widehat{t}_{n}$. For each batch $i$ we generate some pattern avoidance function $g^{i}$, whose error probability is small. To achieve the linear rates we run the algorithm from Lemma 10 for every different batch $i$ and function $g^{i}$ to get a classifier $\widehat{h}^{i}$. We output the majority vote of the classifiers $\widehat{h}^{i}$.

Lemma 11. For any $n \in \mathbb{N}$, Consider a training set $\left\{\left(X_{i}, Y_{i}\right)\right\}$ consisting of $n$ points i.i.d. drawn from $P$. Then there exists a universally measurable $\widehat{t}_{n}=\widehat{t}_{n}\left(X_{1}, Y_{1}, \ldots, X_{\lfloor n / 2\rfloor}, Y_{\lfloor n / 2\rfloor}\right)$ whose definition does not depend on $P$ so that the following holds. Set the critical time $t^{\star} \in \mathbb{N}$ be such that

$$
\operatorname{Pr}\left[\operatorname{per}\left(\widehat{y}_{t^{\star}}\right)>0\right] \leq 1 / 8,
$$

where the probability is over the training set of the algorithm $\widehat{y}_{t}$. Then, there exist $C, c>0$ that depend on $P, t^{\star}$ but not $n$ so that

$$
\operatorname{Pr}\left[\widehat{t}_{n} \in T^{\star}\right] \geq 1-C e^{-c n}
$$

where the probability is over the training of the estimator $\widehat{t}_{n}$ and $T^{\star}$ is the set

$$
T^{\star}=\left\{1 \leq t \leq t^{\star}: \operatorname{Pr}\left[\operatorname{per}\left(\widehat{y}_{t^{\star}}\right)>0\right] \leq 3 / 8\right\}
$$

where the probability is over the training of $\widehat{y}_{t}$.
Proof. We consider the strategies

$$
\ell_{t}=L_{t}\left(X_{1}, Y_{1}, \ldots, X_{t}, Y_{t}\right)
$$

and

$$
\widehat{y}_{t}\left(z_{1}, \ldots, z_{\ell_{t}}, s^{(0)}, s^{(1)}\right)=\widehat{Y}_{t}\left(X_{1}, Y_{1}, \ldots, X_{t}, Y_{t}, z_{1}, \ldots, z_{\ell_{t}}, s^{(0)}, s^{(1)}\right),{ }^{6}
$$

which can be obtained by the pattern avoidance strategies. We remark that the algorithms behind $L_{t}$ and $\widehat{Y}_{t}$ try out all the possible colorings as described in the pattern avoidance algorithm. As a first step we decompose the training set into four parts. We will use the first quarter of the training examples as follows: For any $1 \leq t \leq\lfloor n / 4\rfloor$ and $1 \leq i \leq\lfloor n /(4 t)\rfloor=: \widehat{N}$, we invoke the above two strategies to get

$$
\ell_{t}^{i}=L_{t}\left(X_{(i-1) t+1}, Y_{(i-1) t+1}, \ldots, X_{i t}, Y_{i t}\right)
$$

and

$$
\widehat{y}_{t}^{i}\left(z_{1}, \ldots, z_{\ell_{t}}, s^{(0)}, s^{(1)}\right)=\widehat{Y}_{t}\left(X_{(i-1) t+1}, Y_{(i-1) t+1}, \ldots, X_{i t}, Y_{i t}, z_{1}, \ldots, z_{\ell_{t}}, s^{(0)}, s^{(1)}\right)
$$

For each $t$ in the decomposition, we estimate the value $\operatorname{Pr}\left[\operatorname{per}\left(\widetilde{y}_{t}\right)>0\right]$ using our estimates as
$\widehat{e}_{t}=\frac{1}{\widehat{N}} \sum_{i \in[\widehat{N}]} 1\left\{\widehat{y}_{t}^{i}\right.$ fails to avoid some NL pattern realized by $\left(X_{s+1}, Y_{s+1}, \ldots, X_{s+\ell_{t}^{i}}, Y_{s+\ell_{t}^{i}}\right)$ for some $\left.n / 4 \leq s \leq n / 2-\ell_{t}^{i}\right\}$
We note that almost surely $\widehat{e}_{t} \leq e_{t}=\sum_{i \in[\widehat{N}]} \mathbf{1}\left\{\operatorname{per}\left(\widehat{y}_{t}^{i}\right)>0\right\} / \widehat{N}$. Moreover, we set

$$
\widehat{t}_{n}=\inf \left\{t \leq\lfloor n / 4\rfloor: \widehat{e}_{t}<1 / 4\right\}
$$

and we set $\inf \emptyset=\infty$. Set the critical time $t^{\star} \in \mathbb{N}$ be such that

$$
\operatorname{Pr}\left[\operatorname{per}\left(\widehat{y}_{t^{\star}}\right)>0\right] \leq 1 / 8
$$

where the probability is over the training set of the algorithm $\widehat{y}_{t}$. It holds that

$$
\operatorname{Pr}\left[\widehat{t}_{n}>t^{\star}\right] \leq \operatorname{Pr}\left[\widehat{e}_{t^{\star}} \geq 1 / 4\right] \leq \operatorname{Pr}\left[e_{t^{\star}}-\mathbf{E}\left[e_{t^{\star}}\right] \geq 1 / 8\right] \leq \exp \left(-N^{\star} / 32\right)
$$

where $N^{\star}=\left\lfloor n /\left(4 t^{\star}\right)\right\rfloor$. By continuity, there exists $\epsilon>0$ such that for all $1 \leq t \leq t^{\star}$ such that $\operatorname{Pr}\left[\operatorname{per}\left(\widehat{y}_{t}\right)\right]>3 / 8$, we have that $\operatorname{Pr}\left[\operatorname{per}\left(\widehat{y}_{t}\right)>\epsilon\right]>1 / 4+1 / 16$. Fix $1 \leq t \leq t^{\star}$ with $\operatorname{Pr}\left[\operatorname{per}\left(\widehat{y}_{t}\right)>0\right]>3 / 8$ (if such an index exists). Standard concentration inequalities yield

$$
\operatorname{Pr}\left[\sum_{i \in[\widehat{N}]} 1\left\{\operatorname{per}\left(\widehat{y}_{t}^{i}\right)>\epsilon\right\}<\frac{\widehat{N}}{4}\right] \leq \exp \left(-N^{\star} / 128\right)
$$

We remark that any NL pattern avoidance function $g$ which satisfies $\operatorname{per}(g)>\epsilon$, we have that
$\operatorname{Pr}\left[g\right.$ fails to avoid some NL pattern realized by $\left(X_{s+1}, Y_{s+1}, \ldots, X_{s+\ell}, Y_{s+\ell}\right)$ for some $\left.n / 4 \leq s \leq n / 2-\ell\right] \geq 1-p$,

[^0]where $p=p(\epsilon, n, \ell)=(1-\epsilon)^{\lfloor(n-4) /(4 \ell)\rfloor}$ since there are $\lfloor(n-4) /(4 \ell)\rfloor$ disjoint intervals of length $\ell$ in $[n / 4+1, n / 2] \cap \mathbb{N}$.
We now pass to the second quarter of the dataset to test our guesses using the above results. First, the estimates $\left(\ell_{t}^{i}, \widehat{y}_{t}^{i}\right)$ for $1 \leq i \leq \widehat{N}$ are independent of $\left(X_{s}, Y_{s}\right)$ for $s>n / 4$. Second, using a union bound conditionally on the first quarter of the dataset gives that the probability that every guess function $\widehat{y}_{t}^{i}$ with $\operatorname{per}^{\ell_{t}^{i}}\left(\widehat{y}_{t}^{i}\right)>\epsilon$ (we let per ${ }^{\ell}$ denote that the probability is over a sample from the distribution $P^{\otimes \ell}$ ) makes an error on the second quarter of the dataset is
$$
\left.\operatorname{Pr}\left[(\forall i) \mathbf{1}\left\{\operatorname{per}^{\ell_{t}^{i}}\left(\widehat{y}_{t}^{i}\right)>\epsilon\right\} \leq \mathbf{1}\left\{E_{i, t}\right\}\right\}\right] \geq 1-\widehat{N}(1-\epsilon)^{\left\lfloor(n-4) /\left(4 t^{\star}\right)\right\rfloor}
$$
since $\ell_{t}^{i} \leq t^{\star}$, where $E_{i, t}$ is the event that the $(i, t)$-pattern avoidance estimate $\widehat{y}_{t}^{i}$ fails to avoid some NL pattern realized by the data sequence $\left(X_{s+1}, Y_{s+1}, \ldots, X_{s+\ell_{t}^{i}}, Y_{s+\ell_{t}^{i}}\right)$ for some $n / 4 \leq s \leq$ $n / 2-\ell_{t}^{i}$. This yields that
$$
\operatorname{Pr}\left[\widehat{t}_{n}=t\right] \leq \operatorname{Pr}\left[\widehat{e}_{t}<1 / 4\right] \leq\lfloor n / 4\rfloor(1-\epsilon)^{\left\lfloor(n-4) /\left(4 t^{\star}\right)\right\rfloor}+\exp \left(-N^{\star} / 32\right)
$$

Taking a union bound over the elements of $T^{\star}$, where

$$
T^{\star}=\left\{1 \leq t \leq t^{\star}: \operatorname{Pr}\left[\operatorname{per}\left(\widehat{y}_{t^{\star}}\right)>0\right] \leq 3 / 8\right\}
$$

we obtain that

$$
\operatorname{Pr}\left[\widehat{t_{n}} \notin T^{\star}\right] \leq \exp \left(-N^{\star} / 32\right)+t^{\star} \cdot\left(\lfloor n / 4\rfloor(1-\epsilon)^{\left\lfloor(n-4) /\left(4 t^{\star}\right)\right\rfloor}+\exp \left(-N^{\star} / 32\right)\right)
$$

This concludes the proof since there exist $C, c>0$ so that

$$
\operatorname{Pr}\left[\widehat{t}_{n} \in T^{\star}\right] \geq 1-C \exp (-c n)
$$

Theorem 13 (Linear Rates). Assume that class $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ does not have an infinite NatarajanLittlestone tree. Then, $\mathcal{H}$ admits a learning algorithm that achieves an optimal linear rate.

Proof. The learning algorithm works as follows: It first computes the estimate $\widehat{t}_{n}$ introduced in Lemma 11. Then it splits the data into two halves: the first half is used to compute the pattern avoidance functions $g^{i}:=\widehat{y}_{\widehat{t}_{n}}^{i}$ for $1 \leq i \leq\left\lfloor n /\left(4 \widehat{t}_{n}\right)\right\rfloor$ and the second half is used in order to apply Lemma 10 and get classifiers $\widehat{y}^{i}$ with $\widehat{y}^{i}(x):=\widehat{Y}_{\lfloor n / 2\rfloor+2}^{g^{i}}\left(X_{\lceil n / 2\rceil}, Y_{\lceil n / 2\rceil}, \ldots, X_{n}, Y_{n}, x\right)$. Finally, the algorithm outputs

$$
\widehat{h}_{n}=\operatorname{Maj}\left(\widehat{y}^{1}, \widehat{y}^{2}, \ldots, \widehat{y}^{\left\lfloor n /\left(4 \widehat{t}_{n}\right)\right\rfloor}\right)
$$

and aims to get $\mathbf{E}\left[\operatorname{Pr}_{(X, Y)}\left[\widehat{h}_{n}(X) \neq Y\right]\right] \leq C / n$ for some constant $C$, where the expectation is over the training set used for the predescribed steps. Set $\widehat{N}=\left\lfloor n /\left(4 \widehat{t}_{n}\right)\right\rfloor$ and $N^{\star}=\left\lfloor n /\left(4 t^{\star}\right)\right\rfloor$. With the notation of Lemma 11, we get

$$
\operatorname{Pr}\left[\frac{1}{\widehat{N}} \sum_{i \in[\widehat{N}]} \mathbf{1}\left\{\operatorname{per}\left(\widehat{y}_{\widehat{t}_{n}}^{i}\right)>0\right\}>\frac{1}{100 k}, \widehat{t}_{n} \in T^{\star}\right] \leq t^{\star} \exp \left(-C_{1} \cdot N^{\star}\right)
$$

i.e., the strict majority of the pattern avoidance functions have zero error with high probability where $C_{1}$ is a fixed constant depending on the (uniformly bounded) number of labels $k \in \mathbb{N}$. Using Lemma 11, we get

$$
\begin{aligned}
\mathbf{E}\left[\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Y\right]\right] & \leq \mathbf{P r}\left[\operatorname{Maj}\left(\widehat{y}^{1}(X), \widehat{y}^{2}(X), \ldots, \widehat{y}^{\left\lfloor n /\left(4 \widehat{t}_{n}\right)\right\rfloor}(X)\right) \neq Y\right] \\
& \leq C \exp (-c n)+t^{\star} \exp \left(-C_{1} \cdot n^{\star}\right)+p,
\end{aligned}
$$

where
$p=\operatorname{Pr}\left[\widehat{t}_{n} \in T^{\star}, \sum_{i \in[\widehat{N}]} 1\left\{\operatorname{per}\left(\widehat{y}_{\widehat{t}_{n}}^{i}\right)=0\right\} \geq(100 k-1) \widehat{N} /(100 k), \operatorname{Maj}\left(\widehat{y}^{1}(X), \widehat{y}^{2}(X), \ldots, \widehat{y}^{\left\lfloor n /\left(4 \widehat{t}_{n}\right)\right\rfloor}(X)\right) \neq Y\right]$.

We have that

$$
\operatorname{Pr}\left[\operatorname{Maj}\left(\widehat{y}^{1}(X), \widehat{y}^{2}(X), \ldots, \widehat{y}^{\left\lfloor n /\left(4 \hat{t}_{n}\right)\right\rfloor}(X)\right) \neq Y\right] \leq \operatorname{Pr}\left[\sum_{i \in[\widehat{N}]} \mathbf{1}\left\{\widehat{y}^{i}(X) \neq Y\right\} \geq \widehat{N} / k\right],
$$

since whenever the majority makes a mistake, the right-hand side event occurs. Since $k$ is a fixed constant, any two sets containing at least $1 / k$ and $1-1 /(100 k)$ fractions of $\left\{1,2, \ldots,\left\lceil n / \widehat{t}_{n}\right\rceil\right\}$, must have at least a $\Theta(1)$ fraction in their intersection. Hence, we get that

$$
p \leq \operatorname{Pr}\left[\widehat{t}_{n} \in T^{\star}, \frac{1}{\widehat{N}} \sum_{i \in[\widehat{N}]} \mathbf{1}\left\{\widehat{y}^{i}(X) \neq Y\right\} \cdot \mathbf{1}\left\{\operatorname{per}\left(\widehat{y}_{\widehat{t}_{n}}^{i}\right)=0\right\} \geq \Theta(1)\right]
$$

and using Markov's inequality

$$
p \leq \Theta(1) \cdot \mathbf{E}\left[\frac{1}{\widehat{N}} \sum_{i \in[\widehat{N}]} \mathbf{1}\left\{\widehat{t}_{n} \in T^{\star}\right\} \cdot \mathbf{1}\left\{\widehat{y}^{i}(X) \neq Y\right\} \cdot \mathbf{1}\left\{\operatorname{per}\left(\widehat{y}_{\widehat{t}_{n}}^{i}\right)=0\right\}\right]
$$

We can now apply Lemma 10 conditionally on the first half of the data and get

$$
\mathbf{E}\left[\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Y\right]\right] \leq C \exp (-c n)+t^{\star} \exp \left(-n^{\star} / 128\right)+\Theta(1) \cdot \log (k) \cdot \mathbf{E}\left[\mathbf{1}\left\{\widehat{t}_{n} \in T^{\star}\right\} \cdot \frac{1}{\widehat{N}} \sum_{i \in[\widehat{N}]} \frac{\ell_{\widehat{t}_{n}}^{i}}{\lfloor n / 2\rfloor+2}\right]
$$

Now we have that $\Theta(1) \log (k)=\Theta(1)$ and $\ell_{\widehat{t}_{n}}^{i} \leq \widehat{t}_{n}+1 \leq t^{\star}+1$, since $\widehat{t}_{n} \in T^{\star}$. This gives

$$
\mathbf{E}\left[\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Y\right]\right] \leq C / n
$$

## C. 4 Arbitrarily Slow Rates

The last step that is needed to establish the characterization of learnability in the multiclass setting is to show that if $\mathcal{H}$ has an infinite Natarajan Littlestone tree it is learnable at an arbitrarily slow rate. The following theorem establishes that. The idea of the proof is similar as in the setting with the infinite multiclass Littlestone tree. Intuitively, the reason that in this regime we get arbitrarily slow rates whereas in the other one we get linear rates is that the branching factor now is exponential in the depth of the tree, whereas before it remained constant.
Theorem 14. Assume that $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ has an infinite Natarajan-Littlestone tree. Then, any algorithm that learns $\mathcal{H}$ requires arbitrarily slow rates.

Proof. For the proof, fix a vanishing rate $R(t) \rightarrow 0$, fix any learning algorithm $\widehat{h}_{n}$ for $\mathcal{H}$ and let $\left\{x_{u}\right\}$ be an infinite NL tree for $\mathcal{H}$. Consider a random branch of the tree $y=\left(y_{1}, y_{2}, \ldots\right)$ where the pattern $y_{\ell}=\left(y_{\ell}^{0}, \ldots, y_{\ell}^{\ell-1}\right) \in\{0,1\}^{\ell}$ is chosen uniformly at random from the $\ell$-dimensional Boolean hypercube for any $\ell \in \mathbb{N}$. Fix a finite level $n \in \mathbb{N}$. By the structure of the NL tree, we know that there exist two different colorings $s_{1}, s_{2}$ and a hypothesis $h \in \mathcal{H}$ so that $h$ agrees either with $s_{1}$ if the pattern bit says 1 or with $s_{2}$ otherwise. Our first goal in this universal lower bound is to construct a realizable distribution. We define the random distribution that assigns non-zero mass to the points of the branch $y$ with labels consistent with $h$, i.e.,

$$
P_{y}\left(\left(x_{y \leq \ell-1}^{i}, h\left(x_{y \leq \ell-1}^{i}\right)\right)\right)=\frac{p_{\ell}}{\ell} \text { for } 0 \leq i<\ell, \ell \in \mathbb{N}
$$

where $p_{\ell}$ is a sequence of probabilities so that $\sum_{\ell \in \mathbb{N}} p_{\ell}=1$ that we will select later in the proof. Intuitively, the distribution $P_{y}$ chooses the node of the infinite branch at level $\ell$ with probability $p_{\ell}$; this node contains $\ell$ points of $\mathcal{X}$ and one of them is chosen uniformly at random.
By the structure of the infinite NL tree, we get that such a labeling $h \in \mathcal{H}$ exists for any level $n \in \mathbb{N}$ and this labeling is consistent with all the previous levels $1 \leq n^{\prime}<n$. Thus, we get that

$$
\underset{(x, z) \sim P_{y}}{\mathbf{P r}}[h(x) \neq z] \leq \sum_{\ell>n} p_{\ell}
$$

Hence, as $n \rightarrow \infty$, we get that $P_{y}$ is realizable for any realization of the random branch $y$. Moreover, the mapping $y \rightarrow P_{y}$ is measurable.
We now have to lower bound the loss of the potential learner $\widehat{h}_{n}$ using the probability measure $P_{y}$. Let $(X, Z),\left(X_{1}, Z_{1}\right),\left(X_{2}, Z_{2}\right), \ldots \in(\mathcal{X} \times[k])^{\infty}$ be a collection of i.i.d. samples from $P_{y}$. Equivalently, we can write

1. $X=x_{y \leq T-1}^{I}$ and $Z=z_{T}^{I}$ for two random variables $(T, I)$ with joint distribution $\operatorname{Pr}[T=$ $\ell, I=i]=\frac{p_{\ell}}{\ell}$ for $0 \leq i<\ell, \ell \in \mathbb{N}$.
2. For $j \in \mathbb{N}$, set $X_{j}=x_{y \leq T_{j}-1}^{I_{j}}$ and $Z_{j}=z_{T_{j}}^{I_{j}}$ for two random variables $\left(T_{j}, I_{j}\right)$ with joint distribution $\operatorname{Pr}\left[T_{j}=\ell, \bar{I}_{j}=i\right]=\frac{p_{\ell}}{\ell}$ for $0 \leq i<\ell, \ell \in \mathbb{N}$.

We underline that in the above the random variables $(T, I),\left(T_{1}, I_{1}\right),\left(T_{2}, I_{2}\right)$ are i.i.d. and independent of the random branch $y$. Our goal is to lower bound the error of any learning algorithm: For all $n$ and $\ell$,
$\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z, T=\ell\right] \geq \sum_{i=0}^{\ell-1} \operatorname{Pr}\left[\widehat{h}_{n}\left(x_{y \leq \ell-1}^{i}\right) \neq z_{\ell}^{i}, T=\ell, I=i, T_{1}, \ldots, T_{n} \leq \ell,\left(T_{1}, I_{1}\right), \ldots,\left(T_{n}, I_{n}\right) \neq(\ell, i)\right]$,
where in the right hand side the probability is decreased by additionally requiring that the whole training set is concentrated before the level $\ell+1$ and the testing example is not contained in the training set. Consider this event $E_{n, \ell, i}=\left\{T=\ell, I=i, T_{1}, \ldots, T_{n} \leq \ell,\left(T_{1}, I_{1}\right), \ldots,\left(T_{n}, I_{n}\right) \leq(\ell, i)\right\}$. If we condition on $E_{n, \ell, i}$, we have that the prediction of $\widehat{h}_{n}(X)=\widehat{h}_{n}\left(x_{y \leq \ell-1}^{i}\right)$ is independent of the label $z_{\ell}^{i}$. Hence, we have that

$$
\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z, T=\ell\right] \geq \sum_{i=0}^{\ell-1} \operatorname{Pr}\left[\widehat{h}_{n}\left(x_{y \leq \ell-1}^{i}\right) \neq z_{\ell}^{i} \mid E_{n, \ell, i}\right] \operatorname{Pr}\left[E_{n, \ell, i}\right] \geq \frac{1}{2} \sum_{i=0}^{\ell-1} \operatorname{Pr}\left[E_{n, \ell, i}\right]
$$

By the choice of the randomness over $T, I, T_{1}, I_{1}, \ldots$, we get that

$$
\operatorname{Pr}\left[\widehat{h}_{n}(X) \neq Z, T=\ell\right] \geq \frac{p_{\ell}}{2}\left(1-\sum_{m>\ell} p_{m}-\frac{p_{\ell}}{\ell}\right)^{n}
$$

To conclude the proof we have to choose the sequence of probabilities $\left(p_{\ell}\right)$ and relate it to the vanishing rate $R$. By combining Lemma 5.12 of $\left[\mathrm{BHM}^{+} 21\right]$ and by applying the reverse Fatou's lemma (see e.g., the end of the proof of Theorem 5.11 of $\left[\mathrm{BHM}^{+} 21\right]$ ), the proof is concluded.

## C. 5 A Sufficient Condition for Linear Rates for Multiclass Learning

Another approach to come up with an algorithm that works in the multiclass setting is to use the algorithm that was developed in $\left[\mathrm{BHM}^{+} 21\right]$ for the binary setting. Towards this end, we define a slightly different combinatorial measure for a class $\mathcal{H}$.
Definition 14. A Graph-Littlestone (GL) tree for $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ of depth $d \leq \infty$ consists of a tree

$$
\bigcup_{0 \leq \ell<d}\left\{x_{u} \in \mathcal{X}^{\ell+1}, u \in\{0,1\} \times\{0,1\}^{2} \times \ldots \times\{0,1\}^{\ell}\right\}
$$

and a coloring s mapping each position $u^{i} \in u$ for any node with pattern $u \in\{0,1\} \times \ldots \times\{0,1\}^{\ell}$ for $i \in\{0,1, \ldots, \ell\}$ and $\ell \in\{0,1, \ldots, d-1\}$ of the tree to some color $\{0,1, \ldots, k\}$ such that for every finite level $n<d$, the subtree $T_{n}=\cup_{0 \leq \ell \leq n}\left\{x_{u}: u \in\{0,1\} \times\{0,1\}^{2} \times \ldots \times\{0,1\}^{\ell}\right\}$ satisfies the following:

1. For any path $\boldsymbol{y} \in\{0,1\} \times \ldots \times\{0,1\}^{n+1}$, there exists a concept $h \in \mathcal{H}$ so that $h\left(x_{\boldsymbol{y}_{\leq \ell}}^{i}\right)=$ $s\left(x_{\boldsymbol{y}_{\leq \ell}}^{i}\right)$ if $y_{\ell+1}^{i}=1$ and $h\left(x_{\boldsymbol{y}_{\leq \ell}}^{i}\right) \neq s\left(x_{\boldsymbol{y}_{\leq \ell}}^{i}\right)$ otherwise, for all $0 \leq i \leq \ell$ and $0 \leq \ell \leq n$, where

$$
\boldsymbol{y}_{\leq \ell}=\left(y_{1}^{0},\left(y_{2}^{0}, y_{2}^{1}\right), \ldots,\left(y_{\ell}^{0}, \ldots, y_{\ell}^{\ell-1}\right)\right), x_{\boldsymbol{y}_{\leq \ell}}=\left(x_{\boldsymbol{y}_{\leq \ell}}^{0}, \ldots, x_{\boldsymbol{y}_{\leq \ell}}^{\ell}\right) .
$$

We say that $\mathcal{H}$ has an infinite GL tree if it has a GL tree of depth $d=\infty$.
We note that in the above definition we identify the color $s\left(x_{\boldsymbol{y} \leq \ell}^{i}\right)$ with the (unique) position of this point $x_{\boldsymbol{y} \leq \ell}^{i}$ (since typically the coloring is over positions). One can again verify that if $\mathcal{H}$ has an infinite GL tree, then it has an infinite multiclass Littlestone tree. We show that if a class $\mathcal{H}$ does not have an infinite GL tree, then it is possible to learn $\mathcal{H}$ at a linear rate. The proof of this statement is via a reduction to the binary setting. In fact, we invoke the well-studied "one versus all" approach, where we are trying to learn $k$ different classes $\mathcal{H}_{i}$ that distinguish between points that belong to the $i$-th class and points that belong to some class $j \neq i$.
Theorem 15. Assume that $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ does not have an infinite Graph-Littlestone tree. Then, there exists an algorithm that learns $\mathcal{H}$ at a linear rate.

Proof. Consider $k$ binary classes $\mathcal{H}_{i}$ induced by the hypothesis class $\mathcal{H}$ where $\mathcal{H}_{i}=\{x \mapsto 1\{h(x)=$ $i\}: h \in \mathcal{H}\}$, i.e., for $h \in \mathcal{H}, \widetilde{h} \in \mathcal{H}_{i}$ we have that $\widetilde{h}(x)=1$ if and only if $h(x)=i$.
Claim 3. Assume that $\mathcal{H}$ does not have an infinite GL tree. Then for any $i \in[k]$ the class $\mathcal{H}_{i}$ is learnable at a linear rate.

Proof. Fix $i \in[k]$. It suffices to show that the class $\mathcal{H}_{i}$ does not have an infinite VCL tree (VCL trees are introduced in $\left[\mathrm{BHM}^{+} 21\right]$ and, intuitively, a VCL tree is a GL tree with $k=1$ and $s=1$ everywhere). Towards contradiction, assume that $\mathcal{H}_{i}$ admits an infinite VCL tree $\mathcal{T}=\left\{x_{u}\right\}$. We can construct the following Graph-Littlestone tree: we use the same nodes as in $\mathcal{T}$ and for any node in the tree we use the coloring that colors each point with the color $i$. Fix an arbitrary level $n$ and a path $y$. We know that there exists a binary hypothesis $\widetilde{h} \in \mathcal{H}_{i}$ that realizes this path in the VCL tree. Due to the construction of the class $\mathcal{H}_{i}$, we have that there exists a mapping $h \in \mathcal{H}($ where $\widetilde{h}=1\{h(\cdot)=i\})$ that realizes the path in the constructed tree. This property holds for any path and any level. Hence, we have constructed an infinite GL tree for $\mathcal{H}$ which yields a contradiction.

Assume that we get $k$ binary classifiers $\widehat{h}_{n}^{(i)}$, one for each class $\mathcal{H}_{i}$. Finally, we set $\widehat{h}_{n}(x)=$ $\operatorname{argmax}_{i} \widehat{h}_{n}^{i}(x)$ for any $x \in \mathcal{X}$. The above claim concludes the proof since

$$
\mathbf{E} \underset{(x, y) \sim P}{\mathbf{P r}}\left[\widehat{h}_{n}(x) \neq y\right]=\mathbf{E} \underset{(x, y) \sim P}{\mathbf{P r}}\left[\exists i \in[k]: \widehat{h}_{n}^{(i)}(x) \neq \mathbf{1}\{y=i\}\right] \leq \sum_{i \in[k]} \mathbf{E}\left[\operatorname{err}\left(\widehat{h}_{n}^{(i)}\right)\right] \leq 1 / n
$$

since $k$ is a fixed constant.

We close this section with an open question. Is the GL tree roughly speaking equivalent to the NL tree? This should remind the reader the connection between the Graph and the Natarajan dimension in the uniform multiclass PAC learning. In fact, it holds that $\operatorname{Ndim}(\mathcal{H}) \leq \operatorname{Gdim}(\mathcal{H}) \leq$ $\operatorname{Ndim}(\mathcal{H}) \cdot O(\log (k))$.
Open Question 1. Let $\mathcal{H} \subseteq[k]^{\mathcal{X}}$ for some fixed constant $k \in \mathbb{N}$.

1. Is it true that $\mathcal{H}$ has an infinite NL tree if and only if it has an infinite GL tree?
2. Is it possible to obtain an analogue of the inequality between the Graph and the Natarajan dimensions for the ordinal GL and NL dimensions?

One approach to tackle this problem is to show that if $\mathcal{H}$ has an infinite GL tree, then it is learnable at an arbitrarily slow rate. It is not clear to us that the current proof can be modified to work in this setting. If we try to follow the same steps we cannot guarantee the realizability of the sequence. If we pick an infinite random path and let $p_{i}$ be the path up to depth $i$, we know that for each such depth there exists some $h_{i}$ that realizes this path, i.e., if $b_{i}=1 \Longrightarrow h\left(x_{i}\right)=y^{h_{i}}=s\left(x_{i}\right)$ and if $b_{i}=0 \Longrightarrow h\left(x_{i}\right)=y^{h_{i}} \neq s\left(x_{i}\right)$. Now if consider some depth $j>i$ then the hypothesis $h_{j}$ will agree with $h_{i}$ on every $x_{i}$ with $b_{i}=1$. However, we cannot guarantee that $h_{j}\left(x_{i}\right)=h_{i}\left(x_{i}\right)$ if $b_{i}=0$. Hence, the straightforward way to modify the proof to account for that would be to consider a family of different distributions. Note that this is not allowed in the universal learning setting, since we fix the distribution.

## D Multiclass Learning for Partial Concepts

In this section, we provide our results on multiclass learnability in the context of partial concept classes.

## D. 1 PAC Multiclass Learnability for Partial Concepts: The Proof of Theorem 6

We show the following theorem, which implies Theorem 3. For reader's convenience, we restate Theorem 6.
Theorem. For any partial concept class $\mathcal{H} \subseteq\{0,1, \ldots, k, \star\}^{\mathcal{X}}$ with $\operatorname{Ndim}(\mathcal{H}) \leq \infty$, the sample complexity of PAC learning the class $\mathcal{H}$ satisfies

$$
\mathcal{M}(\epsilon, \delta)=O\left(\frac{\mathrm{Ndim}(\mathcal{H}) \log (k)}{\epsilon} \log (1 / \delta)\right) \text { and } \mathcal{M}(\epsilon, \delta)=\Omega\left(\frac{\mathrm{Ndim}(\mathcal{H})+\log (1 / \delta)}{\epsilon}\right)
$$

In particular, if $\operatorname{Ndim}(\mathcal{H})=\infty$, then $\mathcal{H}$ is not PAC learnable.
Proof. Our algorithm will make use of the one-inclusion hypergraph algorithm whose utility is provided by Lemma 1 for total concepts. We first show the next lemma for the one-inclusion hypergraph predictor for partial concepts.
Lemma 12. Fix a positive constant $k$. For any partial concept class $\mathcal{H} \subseteq\{0,1, \ldots, k, \star\}^{\mathcal{X}}$ with $\operatorname{Ndim}(\mathcal{H})<\infty$, there exists an algorithm $\mathbb{A}:(\mathcal{X} \times[k])^{*} \times \mathcal{X} \rightarrow[k]$ such that, for any $n \in \mathbb{N}$ and any sequence $\left\{\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)\right\} \in(\mathcal{X} \times[k])^{n}$ that is realizable with respect to $\mathcal{H}$,

$$
\operatorname{Pr}_{\sigma \sim \mathcal{U}\left(\mathbb{S}_{n}\right)}\left[\mathbb{A}\left(x_{\sigma}(1), y_{\sigma}(1), \ldots, x_{\sigma}(n-1), y_{\sigma}(n-1), x_{\sigma}(n)\right) \neq y_{\sigma}(n)\right] \leq \frac{\operatorname{Ndim}(\mathcal{H}) \log (k)}{n}
$$

Proof. Fix $n \in \mathbb{N}$. Consider a set of points $S=\left\{x_{1}, \ldots, x_{n}\right\}$ and let $S_{d}$ be the set of distinct elements of the sequence $S$. Define the hypothesis class $\mathcal{H}_{S_{d}}$ that contains all the total functions $h: S_{d} \rightarrow[k]$ such that the sequence $\left\{(x, h(x)): x \in S_{d}\right\}$ is realizable with respect to $\mathcal{H}$.
CASE A: Assume that $\mathcal{H}_{S_{d}} \neq \emptyset$. This is a total concept class and so let $\mathbb{A}_{S_{d}}$ be the algorithm guaranteed to exist by Lemma 1 with $X=S_{d}$ and $\mathcal{H}=\mathcal{H}_{S_{d}}$. For any $y_{1}, \ldots, y_{n} \in[k]$ so that the training sequence $\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)$ is realizable with respect to $\mathcal{H}$ (and so realizable with respect to $\mathcal{H}_{S_{d}}$ ), define

$$
\mathbb{A}\left(x_{1}, y_{1}, \ldots, x_{n-1}, y_{n-1}, x_{n}\right)=\mathbb{A}_{S_{d}}\left(\mathcal{H}_{S_{d}}, x_{1}, y_{1}, \ldots, x_{n-1}, y_{n-1}, x_{n}\right)
$$

Moreover, we can consider any permutation of the sequence $x_{1}, \ldots, x_{n}$ and let the feature space $S_{d}$ and the hypothesis class $\mathcal{H}_{S_{d}}$ the same. Finally, we have that $\operatorname{Ndim}\left(\mathcal{H}_{S_{d}}\right) \leq \operatorname{Ndim}(\mathcal{H})$. This gives the desired bound.
Case B: Assume that $\mathcal{H}_{S_{d}}$ is empty. In this case, set $\mathcal{A}\left(x_{1}, y_{1}, \ldots, x_{n-1}, y_{n-1}, x_{n}\right)=0$ for all sequences $\left(x_{1}, \ldots, x_{n}\right) \in \mathcal{X}^{n}$ and $\left(y_{1}, \ldots, y_{n-1}\right) \in[k]^{n-1}$ so that $\left\{h \in \mathcal{H}: h\left(x_{i}\right)=y_{i}\right.$ with $i<$ $n$ and $\left.h\left(x_{n}\right) \in[k]\right\}=\emptyset$.

Let us now focus on the upper bound given that $\operatorname{Ndim}(\mathcal{H})<\infty$. For any distribution $P$ realizable with respect to $\mathcal{H}$ and for a sequence of $n$ labeled i.i.d. examples from $P$, we define the strategy $\widehat{h}_{n}(\cdot)=\mathbb{A}\left(X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}, \cdot\right)$ and so
$\mathbf{E}\left[\operatorname{er}_{P}\left(\widehat{h}_{n}\right)\right]=\underset{\left(X_{i}, Y_{i}\right)_{i \leq n}}{\mathbf{E}}\left[X_{\left(X_{n+1}, Y_{n+1}\right)}^{\mathbf{P r}}\left[\mathbb{A}\left(X_{1}, Y_{1}, \ldots, X_{n}, Y_{n}, X_{n+1}\right) \neq Y_{n+1}\right]\right] \leq \frac{\mathrm{Ndim}(\mathcal{H}) \Theta(\log (k))}{n+1}$.
We next have to convert this algorithm which guarantees an expected error bounded by $N \operatorname{dim}(\mathcal{H}) \Theta(\log (k)) /(n+1)$ into an algorithm that guarantees a bound on the error with probability at least $1-\delta$. In order to boost the algorithm, we use a standard boosting algorithm by decomposing the dataset into $\log (1 / \delta)$ parts and using Chernoff bounds. For the details we refer to the boosting trick of [HLW94] and the proof of Theorem 34.(i) of [AHHM22].
Let $\operatorname{Ndim}(\mathcal{H})=\infty$. We will show that $\mathcal{H}$ is not PAC learnable. For any $\ell \leq \operatorname{Ndim}(\mathcal{H})$, let $\mathcal{X}_{\ell}=\left\{x_{1}, \ldots, x_{\ell}\right\}$ be a set $N$-shattered by $\mathcal{H}$ using the function $f$. Let $\mathcal{H}_{\ell}$ be the class of all total functions $\mathcal{X}_{\ell} \rightarrow\{0,1, \ldots, \ell\}$, any distribution $P$ on $\mathcal{X}_{\ell} \times\{0,1, \ldots, \ell\}$ realizable with respect
to $\mathcal{H}_{\ell}$ can be extended to a distribution on $\mathcal{X} \times\{0,1, \ldots, \ell\}$ realizable with respect to $\mathcal{H}$ with $P\left(\left(\mathcal{X} \backslash \mathcal{X}_{k}\right) \times\{0,1, \ldots, \ell\}\right)=0$. Thus, any lower bound on the sample complexity of PAC learning the total concept class $\mathcal{H}_{\ell}$ is also a lower bound on the sample complexity of learning the partial class $\mathcal{H}$. This gives the desired lower bound. Also, the partial concept classes with infinite Natarajan dimension are not PAC learnable.

## D. 2 The Proof of Proposition 2

Our goal is to understand when the ERM principle succeeds in the partial setting. To address this natural and important task, we pose the following question: What is the difference between learning partial concepts with output in $\{0,1, \ldots, k, \star\}$ and learning total concepts with labels in $\{0,1, \ldots, k, k+1\}$, where $k$ is a positive integer? Conceptually, the key difference between partial concepts with $k+1$ labels and $(k+2)$-label multiclass classification has to do with the support of the distribution: In the latter, the learning problem is a distribution over $\mathcal{X} \times\{0,1, \ldots, k, k+1\}$, while in the former we have a distribution only over $\mathcal{X} \times\{0,1, \ldots, k\}$ (recall Definition 4 ). We address this question in Proposition 2. We show that any partial concept class $\mathcal{H}$ is learnable in the $(k+2)$-label setting if and only if $\mathcal{H}$ is learnable in the partial setting and the VC dimension of the set family $\{\operatorname{supp}(h): h \in \mathcal{H}\}$ is finite. This result is helpful since it guarantees that when the VC dimension of the above family is bounded, the ERM principle provably holds and can be applied in the partial setting.
The next result gives a formal connection between the two settings.
Proposition 2. Any partial class $\mathcal{H} \subseteq\{0,1, \ldots, k, \star\}^{\mathcal{X}}$ is PAC learnable in the $(k+2)$ label multiclass setting if and only if $\overline{\mathcal{H}}$ is PAC learnable in the partial concepts setting and $\operatorname{VCdim}(\{\operatorname{supp}(h): h \in \mathcal{H}\})<\infty$.

The following proof is an adaptation of Proposition 23 of [AHHM22] to the multiclass setting.
Proof. Let us first assume that $\mathcal{H}$ is PAC learnable in the $(k+2)$-label multiclass setting. This implies that $\operatorname{Ndim}_{k+2}(\mathcal{H}) \leq \operatorname{Gdim}_{k+2}(\mathcal{H}) \leq O(\log (k+2)) \cdot \operatorname{Nim}_{k+2}(\mathcal{H})<\infty$. This implies that $\mathcal{H}$ is also learnable in the partial concepts setting by Theorem 3 (by the definitions of the extended Graph and Natarajan dimensions; intuitively in the partial setting, we can use one less color). Now consider the collection of sets $\mathbb{S}=\{\operatorname{supp}(h): h \in \mathcal{H}\}=\{\{x \in \mathcal{X}: h(x) \neq \star\}: h \in \mathcal{H}\}$. For any sequence of $d$ points $x_{1}, \ldots, x_{d}$ shattered by $\mathbb{S}$, let us take the hypothesis $f: \mathcal{X} \rightarrow\{0,1, \ldots, k, \star\}$ so that $f\left(x_{i}\right)=\star$ for any $i \in[d]$. As an implication, this sequence $x_{1}, \ldots, x_{d}$ is also shattered by $\{x \mapsto \mathbf{1}\{h(x)=f(x)\}: h \in \mathcal{H}\}$. By the definition of the graph dimension, we get that $\operatorname{Gdim}_{k+2}(\mathcal{H}):=\sup _{f: \mathcal{X} \rightarrow\{0,1, \ldots, k, \star\}} \operatorname{VCdim}(x \mapsto 1\{h(x)=f(x)\}: h \in \mathcal{H}) \geq \operatorname{VCdim}(\mathbb{S})$. Since $\mathcal{H}$ is PAC learnable in the $(k+2)$-label multiclass setting, we get that $\operatorname{VCdim}(\{\operatorname{supp}(h): h \in$ $\mathcal{H}\})<\infty$.
Let us now assume that $\mathcal{H}$ is PAC learnable in the partial concepts setting and $\operatorname{VCdim}(\{\operatorname{supp}(h)$ : $h \in \mathcal{H}\})<\infty$. We are going to show that the Natarajan dimension of $\mathcal{H}$ is not infinite in the $(k+2)$ label multiclass setting. Fix a sequence $\left(x_{i}, y_{i}^{(0)}, y_{i}^{(1)}\right)_{i \in[d]} \in(\mathcal{X} \times\{0,1, \ldots, k, \star\} \times\{0,1, \ldots, k, \star\})^{d}$ for some $d \in \mathbb{N}$, as in the definition of the $(k+2)$-label Natarajan dimension. For this sequence, consider the set $\left\{x_{i}: \star \notin\left\{y_{i}^{(0)}, y_{i}^{(1)}\right\}\right\}$ with $y_{i}^{(0)} \neq y_{i}^{(1)}$. This set is shattered by the partial concept class $\mathcal{H}$ by the extension of the Natarajan dimension to the partial concepts setting. Moreover, the set $\left\{x_{i}: \star \in\left\{y_{i}^{(0)}, y_{i}^{(1)}\right\}\right\}$ is shattered by the set $\mathbb{S}=\{\operatorname{supp}(h): h \in \mathcal{H}\}$. This implies that $\operatorname{Ndim}_{k+2}(\mathcal{H}) \leq \operatorname{Ndim}(\mathcal{H})+\operatorname{VCdim}(\mathbb{S})<\infty$.

## D. 3 Multiclass Disambiguations

We extend the definition of [AHHM22] to multiclass partial concepts classes.
Definition 15 (e.g., [AHHM22]). A total concept class $\overline{\mathcal{H}} \subseteq[k]^{\mathcal{X}}$ is a special type of partial concept class such that every $h \in \overline{\mathcal{H}}$ has range $\{0,1, \ldots, k\}$, i.e., is a total concept. A total concept class $\overline{\mathcal{H}}$ is said to disambiguate a partial concept class $\mathcal{H} \subseteq\{0,1, \ldots, k, \star\}^{\mathcal{X}}$ if every finite data sequence $S \in(\mathcal{X} \times[k])^{*}$ realizable with respect to $\mathcal{H}$ is also realizable with respect to $\overline{\mathcal{H}}$. In this case, $\overline{\mathcal{H}}$ is called a disambiguation of $\mathcal{H}$.

We will make use of the following key result in graph theory and communication complexity. We let $\chi(G)$ be the chromatic number of the simple graph $G$. We also set $\operatorname{bp}(G)$ be the biclique partition number of $G$, i.e., the minimum number of complete bipartite graphs needed to partition the edge set of $G$. The following result lies in the intersection of complexity theory and graph theory and is a result of numerous works [HS12, Ama14, AHHM22, SA15, Göö15, GPW18, BDHT17, BBDG ${ }^{+}$22]. Motivation was the Alon-Saks-Seymour problem in graph theory, which asks: How large a gap can there be between the chromatic number of a graph and its biclique partition number?
Proposition 3 (Biclique Partition and Chromatic Number $\left[\mathrm{BBDG}^{+} 22\right]$ ). For any $n \in \mathbb{N}$, there exists a simple graph $G$ with $\operatorname{bp}(G)=n$ such that

$$
\chi(G) \geq n^{(\log (n))^{1-\epsilon(n)}}
$$

where $\epsilon(n)$ is a sequence that tends to 0 as $n \rightarrow \infty$.
Let us consider the binary classification setting. In the work of [AHHM22], it was shown that the (combinatorial version) of the Sauer-Shelah-Perles (SSP) lemma fails in the partial concepts setting. In the total concepts setting, this lemma controls the size of a concept class $\mathcal{H}$ in terms of its VC dimension. Another variant of this lemma controls the growth function of the class. Given a set $C=\left\{x_{1}, \ldots, x_{m}\right\} \subseteq \mathcal{X}$, the growth function of $\mathcal{H} \subseteq\{0,1\}^{\mathcal{X}}$ with respect to $C$ is the cardinality of the set $\Pi_{\mathcal{H}}(C)$ of binary patterns realized by hypotheses in $\mathcal{H}$ when projected to $C$, i.e.,

$$
\Pi_{\mathcal{H}}(C)=\left\{\left(h\left(x_{1}\right), \ldots, h\left(x_{m}\right)\right): h \in \mathcal{H}\right\} \subseteq\{0,1\}^{m} .
$$

We define the growth function of $\mathcal{H}$ at $m \in \mathbb{N}$ as

$$
\Pi_{\mathcal{H}}(m)=\sup _{C \subseteq \mathcal{X}:|C|=m}\left|\Pi_{\mathcal{H}}(C)\right|
$$

This definition naturally extends to partial concepts where we still look only for binary patterns. Interestingly, while the combinatorial [Sau72] and the growth function [SSBD14] versions of the SSP lemma are both true in the total concepts setting, this is not true in the partial case. We show that the growth function variant still holds, while the combinatorial one fails [AHHM22].
Lemma 13 (Growth Function - SSP Lemma for Partial Concepts). Let $\mathcal{H} \subseteq\{0,1, \star\}^{\mathcal{X}}$ be a partial concept class with finite VC dimension. For any $m \in \mathbb{N}$, it holds that

$$
\Pi_{\mathcal{H}}(m) \leq \sum_{i=0}^{\text {VCdim }(\mathcal{H})}\binom{m}{i} .
$$

Proof. Set $d=\operatorname{VCdim}(\mathcal{H})$. Consider a set $C=\left\{x_{1}, \ldots, x_{m}\right\} \subseteq \mathcal{X}$ of $m$ points and define $\Pi_{\mathcal{H}}(C)=\left\{\left(h\left(x_{1}\right), \ldots, h\left(x_{m}\right)\right): h \in \mathcal{H}\right\} \subseteq\{0,1\}^{m}$ (where we ignore any vector that contains the $\star$ symbol). Note that if $m \leq d$, then it holds $\left|\Pi_{\mathcal{H}}(C)\right|=2^{m}$, by the definition of the VC dimension in the partial setting. We are going to prove the next claim.
Claim 4. For any $C=\left\{x_{1}, \ldots, x_{m}\right\}$ and any binary partial concept class $\mathcal{H}$, we have that

$$
\left|\Pi_{\mathcal{H}}(C)\right| \leq \mid\{B \subseteq C: \mathcal{H} \text { shatters } B\} \mid
$$

The above claim suffices since the RHS is at most $\sum_{i=0}^{V C d i m}(\mathcal{H})\binom{m}{i}$. Now we prove the above claim. For $m=1$, the result holds. Assume that the claim is true for sets of size $\ell<m$ and let us prove it for sets of size $m$. Fix some partial binary concept class $\mathcal{H} \subseteq\{0,1, \star\}^{\mathcal{X}}$ and set $C=\left\{x_{1}, \ldots, x_{m}\right\}$. Let $C^{\prime}=C \backslash\left\{x_{1}\right\}$ and define

$$
Y_{0}=\left\{\left(y_{2}, \ldots, y_{m}\right) \in\{0,1\}^{m-1}:\left(0, y_{2}, \ldots, y_{m}\right) \in \Pi_{\mathcal{H}}(C) \vee\left(1, y_{2}, \ldots, y_{m}\right) \in \Pi_{\mathcal{H}}(C)\right\}
$$

and

$$
Y_{1}=\left\{\left(y_{2}, \ldots, y_{m}\right) \in\{0,1\}^{m-1}:\left(0, y_{2}, \ldots, y_{m}\right) \in \Pi_{\mathcal{H}}(C) \wedge\left(1, y_{2}, \ldots, y_{m}\right) \in \Pi_{\mathcal{H}}(C)\right\}
$$

Note that $\left|\Pi_{\mathcal{H}}(C)\right|=\left|Y_{0}\right|+\left|Y_{1}\right|$ (due to double counting). Our first observation is that $\left|Y_{0}\right| \leq$ $\left|\Pi_{\mathcal{H}}\left(C^{\prime}\right)\right|$, since there may exist some $h \in \mathcal{H}$ which is undefined at $x_{1}$ and that generates a pattern that is not contained in $Y_{0}$. Using the inductive hypothesis on $\mathcal{H}$ and $C^{\prime}$ for the second inequality and the definition of $C^{\prime}$ for the third equality, we get

$$
\left|Y_{0}\right| \leq\left|\Pi_{\mathcal{H}}\left(C^{\prime}\right)\right| \leq \mid\left\{B \subseteq C^{\prime}: \mathcal{H} \text { shatters } B\right\}|=|\left\{B \subseteq C: x_{1} \notin B \wedge \mathcal{H} \text { shatters } B\right\} \mid
$$

Let us now set

$$
\mathcal{H}^{\prime}=\left\{h \in \mathcal{H}: \exists h^{\prime} \in \mathcal{H} \text { such that }\left(h\left(x_{1}\right), h\left(C^{\prime}\right)\right)=\left(1-h^{\prime}\left(x_{1}\right), h^{\prime}\left(C^{\prime}\right)\right) \wedge \star \notin h(C)\right\}
$$

i.e., $\mathcal{H}^{\prime}$ contains all the pairs of hypotheses that (i) are well defined over $C$, (ii) agree on $C^{\prime}$ and (iii) differ on $x_{1}$. Note that if $\mathcal{H}^{\prime}$ shatters a set $B \subseteq C^{\prime}$ then it also shatters $B \cup\left\{x_{1}\right\}$ and vice versa. Moreover we have that $Y_{1}=\Pi_{\mathcal{H}^{\prime}}\left(C^{\prime}\right)$. By the inductive hypothesis on $\mathcal{H}^{\prime}$ and $C^{\prime}$, we have that

$$
\left|Y_{1}\right|=\left|\Pi_{\mathcal{H}^{\prime}}\left(C^{\prime}\right)\right| \leq \mid\left\{B \subseteq C^{\prime}: \mathcal{H}^{\prime} \text { shatters } B\right\} \mid .
$$

This gives that

$$
\left|Y_{1}\right| \leq \mid\left\{B \subseteq C^{\prime}: \mathcal{H}^{\prime} \text { shatters } B \cup\left\{x_{1}\right\}\right\}|=|\left\{B \subseteq C: x_{1} \in B \wedge \mathcal{H}^{\prime} \text { shatters } B\right\} \mid
$$

Finally since $\mathcal{H}^{\prime}$ lies inside $\mathcal{H}$, we get

$$
\left|Y_{1}\right| \leq \mid\left\{B \subseteq C: x_{1} \in B \wedge \mathcal{H} \text { shatters } B\right\} \mid
$$

Combining our observations for $Y_{0}$ and $Y_{1}$, we get that

$$
\left|\Pi_{\mathcal{H}}(C)\right| \leq \mid\{B \subseteq C: \mathcal{H} \text { shatters } B\} \mid
$$

We invoke the above SSP variant to prove a bound for the Natarajan dimension of a partial concept class with multiple labels.
Lemma 14. Let $\mathcal{H}_{\text {bin }} \subseteq\{0,1, \star\}^{\mathcal{X}}$ be a partial concept class with finite VC dimension and $\mathcal{H}=$ $\left\{h=r\left(h_{1}, \ldots, h_{\ell}\right): h_{i} \in \mathcal{H}_{\text {bin }}\right\}$ for some $r:\{0,1\}^{\ell} \rightarrow[k]$. It holds that

$$
\operatorname{Ndim}(\mathcal{H}) \leq \widetilde{O}\left(k \cdot \operatorname{VCdim}\left(\mathcal{H}_{\text {bin }}\right)\right)
$$

Proof. Let the VC dimension of the binary concept class be $d$. Let $S \subseteq \mathcal{X}$ be a shattered set by the partial concept class $\mathcal{H}$. Hence we have that

$$
\left|\Pi_{\mathcal{H}}(S)\right| \geq 2^{|S|}
$$

Now any hypothesis $h \in \mathcal{H}$ is identified by $k$ binary partial concepts from $\mathcal{H}_{\text {bin }}$. We have that

$$
\left|\Pi_{\mathcal{H}}(S)\right| \leq\left|\Pi_{\mathcal{H}_{\mathrm{bin}}}(S)\right|^{k}
$$

by the structure of $\mathcal{H}$. This implies that

$$
\left|\Pi_{\mathcal{H}}(S)\right| \leq O\left(|S|^{d}\right),
$$

by the properties of the growth function. This gives that $|S| \leq \widetilde{O}(d k)$ and concludes the proof.
The next theorem is one of the main results of this section; it essentially states that there exists some simple (with small Natarajan dimension) partial concept class $\mathcal{H}^{\star}$ in the multiclass classification setting which cannot be disambiguated in the sense that any extension of $\mathcal{H}^{\star}$ to a total concept class has unbounded Natarajan dimension. Hence, for this class, there is no way to assign labels to the undefined points and preserve the expressivity of the induced collection of total classifiers.
Theorem 16 (Disambiguation). Fix $k \in \mathbb{N}$. For any $n \in \mathbb{N}$, there exists a partial concept class $\mathcal{H}_{n} \subseteq\{0,1, \ldots, k, \star\}^{[n]}$ with $\operatorname{Ndim}\left(\mathcal{H}_{n}\right)=O_{k}(1)$ such that any disambiguation $\overline{\mathcal{H}}$ of $\mathcal{H}_{n}$ has size at least $n^{\log (n)^{1-o(1)}}$, where the $o(1)$ term tends to 0 as $n \rightarrow \infty$. This implies that there exists $\mathcal{H}_{\infty} \subseteq\{0,1, \ldots, k, \star\}^{\mathbb{N}}$ with $\operatorname{Ndim}\left(\mathcal{H}_{\infty}\right)=O_{k}(1)$ but $\operatorname{Nim}(\overline{\mathcal{H}})=\infty$ for any disambiguation $\overline{\mathcal{H}}$ of $\mathcal{H}_{\infty}$.

Proof. The proof is essentially a tensorization of the construction of [AHHM22]. Fix $k, n \in \mathbb{N}$. Set $L=\log _{2}(k+1)$ and assume that $L \in \mathbb{N}$ without loss of generality. Let us consider the partial concept class

$$
\mathcal{H}_{n}=\mathcal{H}_{n}^{(1)} \times \ldots \times \mathcal{H}_{n}^{(L)}
$$

where $\mathcal{H}_{n}^{(i)} \subseteq\{0,1, \star\}^{[n]}$. Let us now explain the partial concepts that lie in $\mathcal{H}_{n}^{(i)}$ and subsequently in $\mathcal{H}_{n}$. To this end, we invoke Proposition 3 which lies in the intersection of combinatorics and
complexity theory. Consider $L$ independent and disjoint copies of the graph promised by Proposition 3 and let $G$ be the union of these $L$ graphs. Fix $i \in[L]$. Define the partial class $\mathcal{H}_{n}^{(i)} \subseteq\{0,1, \star\}^{[n]}$ using the graph $G^{(i)}$ and its $n$ bipartite complete graphs $B_{j}^{(i)}=\left(L_{j}^{(i)}, R_{j}^{(i)}, E_{j}^{(i)}\right)$ with $j \in[n]$. The class contains $\left|V\left(G^{(i)}\right)\right|$ concepts, each one identified by a vertex $v \in V\left(G^{(i)}\right)$ with

$$
c_{v}^{(i)}(j)=\left\{\begin{array}{cc}
0 & \text { if } v \in L_{j}^{(i)} \\
1 & \text { if } v \in R_{j}^{(i)} \\
\star & \text { otherwise. }
\end{array}\right\}
$$

Using Lemma 31 from [AHHM22], we have that $\operatorname{VCdim}\left(\mathcal{H}_{n}^{(i)}\right)=1$.
We overload the " + " notation by setting $g+\star=\star$ for any $g \in \mathbb{N}$. The partial concept class $\mathcal{H}_{n}$ contains all the partial concepts $h_{v_{1}, \ldots, v_{L}}(i)=\sum_{j \in[L]} 2^{L} \cdot c_{v_{j}}^{(j)}(i) \in\{0,1, \ldots, k, \star\}$. Hence we have that $\mathcal{H}_{n} \subseteq\{0,1, \ldots, k, \star\}^{[n]}$. We can use the growth function of the partial concepts setting and get that $\operatorname{Ndim}\left(\mathcal{H}_{n}\right)=O\left(\log (k+1) \cdot \operatorname{VCdim}\left(\mathcal{H}_{n}^{(1)}\right)\right)=O_{k}(1)$.
Consider some disambiguation $\overline{\mathcal{H}} \subseteq[k]^{[n]}$ of $\mathcal{H}_{n}$. The class $\overline{\mathcal{H}}$ induces $L$ disambiguations $\overline{\mathcal{H}^{(i)}}$ for the binary partial classes $\mathcal{H}_{n}^{(i)}$. Then $\overline{\mathcal{H}}$ defines a coloring of $G$ using $\min _{i}\left|\overline{\mathcal{H}^{(i)}}\right|$ colors. Proposition 3 implies that

$$
\min _{i}\left|\overline{\mathcal{H}^{(i)}}\right| \geq n^{\log (n)^{1-o(1)}}
$$

Finally, we can consider the class $\mathcal{H}_{\infty}$ as the disjoint union of $\mathcal{H}_{n}$. Each $\mathcal{H}_{n}$ has domain $\mathcal{X}_{n}$, where the domains $\mathcal{X}_{n}$ are mutually disjoint and $\mathcal{H}_{\infty}$ is the union $\bigcup_{n} \widetilde{\mathcal{H}}_{n}$, where $\widetilde{\mathcal{H}}_{n}$ is obtained from $\mathcal{H}_{n}$ by adding $\star$ outside of its domain. Then $\operatorname{Ndim}\left(\mathcal{H}_{\infty}\right)=O_{k}(1)$. Since the size of its disambiguations is unbounded, then the multiclass Sauer-Shelah-Perles Lemma [BCHL95] implies that the Natarajan dimension of any disambiguation of $\mathcal{H}_{\infty}$ is infinite.


[^0]:    ${ }^{6}$ Notice that $L_{t}, \widehat{Y}_{t}$ depend also on the colorings that were used throughout the execution of the algorithm. We try the colorings in some fixed order and we slightly abuse the notation to drop that dependence.

