

Online Learning and Solving Infinite Games with an ERM Oracle

Angelos Assos^{*} Idan Attias[†] Yuval Dagan[‡] Constantinos Daskalakis[§]
Maxwell Fishelson[¶]

July 11, 2023

Abstract

While ERM suffices to attain near-optimal generalization error in the stochastic learning setting, this is not known to be the case in the online learning setting, where algorithms for general concept classes rely on computationally inefficient oracles such as the Standard Optimal Algorithm (SOA). In this work, we propose an algorithm for online binary classification setting that relies solely on ERM oracle calls, and show that it has finite regret in the realizable setting and sublinearly growing regret in the agnostic setting. We bound the regret in terms of the Littlestone and threshold dimensions of the underlying concept class.

We obtain similar results for nonparametric games, where the ERM oracle can be interpreted as a best response oracle, finding the best response of a player to a given history of play of the other players. In this setting, we provide learning algorithms that only rely on best response oracles and converge to approximate-minimax equilibria in two-player zero-sum games and approximate coarse correlated equilibria in multi-player general-sum games, as long as the game has a bounded fat-threshold dimension. Our algorithms apply to both binary-valued and real-valued games and can be viewed as providing justification for the wide use of double oracle and multiple oracle algorithms in the practice of solving large games.

1 Introduction

The advent of Deep Learning has exacerbated the importance of learning models which involve a large number of parameters or are non-parametric. Non-parametric learning is learning at its fullest generality. We make no assumption about the structure of our decision space, working with potentially infinite and non-continuous hypothesis classes. From a theoretical standpoint, most study of non-parametric learning has focused on the stochastic setting, where one learns a model given independent observations from some distribution. This study has led to important developments in—both frequentist and Bayesian—non-parametric Statistics, including the discovery of notions of complexity of hypotheses classes, such as the celebrated VC and fat-shattering dimensions, which tightly capture the number of observations from a distribution that is necessary to select a hypothesis whose prediction error under the distribution is approximately optimal. In fact, this is achieved via Empirical Risk Minimization (ERM), the simple method, which given a hypothesis class $\mathcal{H}$, of functions mapping a feature set $\mathcal{X}$ to a label set $\mathcal{Y} \subseteq \mathbb{R}$, and a set of observations $\{(x_i, y_i)\}_{i=1}^N$, outputs:¹

$$(\text{ERM}): \quad \arg \min_{h \in \mathcal{H}} \sum_{i=1}^N |h(x_i) - y_i|. \quad (1)$$

The main goal of this paper is to advance our understanding of non-parametric learning in the more general setting of *online learning*, which is general enough to capture a variety of other learning settings as special cases and has found applications in a diversity of fields, including optimization and *game theory*, which is also in the focus of this paper. We will consider a fairly general online learning setting wherein a learner interacts with an adversary over a number of rounds. In each round $t = 1, 2, \dots, T$, the learner picks a distribution $\hat{\mathcal{D}}_t$ over functions $h : \mathcal{X} \rightarrow \mathcal{Y}$, where $\mathcal{X}$ is a feature set and $\mathcal{Y} \subseteq \mathbb{R}$ is a label set, then the adversary picks a feature-label pair $(x_t, y_t) \in \mathcal{X} \times \mathcal{Y}$, and then the learner draws a sample $\hat{h}_t \sim \hat{\mathcal{D}}_t$ and suffers loss $|\hat{h}_t(x_t) - y_t|$. The learner's

^{*}MIT CSAIL; assos@mit.edu.

[†]Department of Computer Science, Ben-Gurion University; idanatti@post.bgu.ac.il.

[‡]MIT CSAIL; yuvaldag@gmail.com.

[§]MIT CSAIL; costis@csail.mit.edu.

[¶]MIT CSAIL; maxfish@mit.edu.

¹To be formal, if the $\inf_{h \in \mathcal{H}} \sum_{i=1}^N |h(x_i) - y_i|$ is not attained, we allow some small optimization error $\epsilon > 0$ from an ERM oracle, namely allowing it to output any hypothesis whose loss is within ϵ from the infimum.

losses add up over rounds and the learner’s goal is to make her total loss over several rounds as small as possible compared to some benchmark loss computed with hindsight information.

There are many variations to the theme depending on what distribution $\hat{\mathcal{D}}_t$ the learner is allowed to use, what data (x_t, y_t) the adversary is allowed to supply, what benchmark loss the learner competes against, etc. In this work, we consider the common setting where there is some class $\mathcal{H}$ of hypotheses from $\mathcal{X}$ to $\mathcal{Y}$, and the performance of the learner is measured against the optimal hindsight error made by functions in this class: $\inf_{h \in \mathcal{H}} \sum_{t=1}^T |h(x_t) - y_t|$, i.e. the goal of the learner is to minimize the following quantity, called *regret*, in expectation or with good probability:

$$\sum_{t=1}^T |\hat{h}_t(x_t) - y_t| - \inf_{h \in \mathcal{H}} \sum_{t=1}^T |h(x_t) - y_t|. \quad (2)$$

When the distributions $\hat{\mathcal{D}}_t$ chosen by the learner are supported on hypotheses from $\mathcal{H}$ the learner is called (*randomized*) *proper*, otherwise the learner is called *improper*. When the pairs (x_t, y_t) chosen by the adversary satisfy that $y_t = h(x_t)$ for some $h \in \mathcal{H}$ the setting is called *realizable* otherwise the setting is called *agnostic*. Finally, when $\mathcal{Y} = \{0, 1\}$, we call $\mathcal{H}$ a *0-1-valued or binary-valued concept class* and the learning task *online classification*. Otherwise, when $\mathcal{Y} = \mathbb{R}$ or some subset such as $[0, 1]$, $\mathcal{H}$ is called *real-valued* and the learning task is called *online regression*.

Similar to the stochastic setting, there has been extensive work in the online setting on developing complexity measures of concept classes, which suffice to characterize or bound the optimal regret (2) that is attainable by a proper or improper learner, in the realizable or agnostic setting, and for binary-valued or real-valued functions. For example, the celebrated work of Littlestone [1988] characterizes the optimal regret bound attainable by an improper learner, in the realizable online classification setting, in terms of a complexity measure of the concept class $\mathcal{H}$ that is now known as *Littlestone dimension* of $\mathcal{H}$; see Definition A.3. More recent work by Daskalakis and Golowich [2022] provides near-matching bounds for the optimal regret of proper learners in the same setting, and generalizes this result to proper learners in the realizable online regression setting, providing regret bounds in terms of the natural generalization to this setting of Littlestone dimension, called *sequential fat-shattering dimension*; see Definition A.9. In the agnostic setting, Ben-David et al. [2009] and Alon et al. [2021] characterize the optimal regret of improper learners in the online classification setting in terms of Littlestone dimension, while in the online regression setting Rakhlin et al. [2010] characterize the optimal regret of proper learners in terms of the sequential Rademacher complexity as well as in terms of the sequential fat-shattering dimension, the latter characterization being recently tightened by Block et al. [2021]. Finally, Hanneke et al. [2021] obtain near-optimal proper learners for online classification in the agnostic setting via more constructive arguments compared to Rakhlin et al. [2010] and Block et al. [2021].

In contrast to the stochastic setting, however, our understanding of attainable regret in the online learning setting is quite more limited, in the sense that the afore-described works, which bound or characterize the optimal regret, are either non-constructive or make oracle queries to the Standard Optimal Algorithm (SOA) proposed by Littlestone [1988] or generalizations thereof [Daskalakis and Golowich, 2022]. Indeed, learning algorithms for non-parametric hypothesis classes must have access to some oracle in order to interface with a potentially infinite menu of hypotheses. We must keep in mind, though, that our goal in studying learning in this general setting is to say something meaningful about specific learning tasks. Thus, our selection of oracle model should be informed by the tasks to which we hope to apply our non-parametric algorithm. SOA and its generalizations involve computing the Littlestone dimension or the sequential fat-shattering dimension of concept classes defined by the online learner in the course of its interaction with the adversary, which are challenging computations, even when the concept class and the set of features are finite [Manurangsi and Rubinfeld, 2017, Manurangsi, 2022]. Thus, the non-parametric learning algorithms coming from the SOA oracle model are utterly useless for any practical applications.

On the other hand, the stochastic learning setting is studied under the more standard ERM oracle model (1). The learning algorithms here enjoy both success in their guaranteed performance and practical feasibility in their more realistic oracle assumption. Ideally, we would like to construct online learning algorithms using this standard oracle too. One of the main questions we ask is thus the following:

Goal 1.1 *In the non-parametric online learning setting, do there exist learning algorithms whose steps run in finite time given access to an ERM oracle as well as standard arithmetic operations, and whose regret is finite or sublinearly growing with T ?*

One of our main contributions is to provide positive answers to this question for the online classification setting, as summarized in Table 1. In the realizable setting, we provide an improper learner whose regret is finite and a proper learner whose regret grows sublinearly in the number of rounds; see Theorem 3.3. In the agnostic setting,

Setting	Time per iter.	ERM calls/iter.	Regret
Realizable, improper	$\min(t, 4^m, C^d)$	1	$\min(4^m, C^d)$
Realizable, proper	t	1	$T^{\frac{2m+2}{2m+3}}$
Agnostic, improper	$t^{\min(4^m, C^d)}$	$t^{\min(4^m, C^d)}$	$\sqrt{T \min(4^m, C^d)}$
Agnostic, proper	$e^{o(T)}$	T^d	$T^{\frac{2m+3}{2m+4}}$

Table 1: The table describes the complexities of our online learner, Algorithm 1 (and its extension to the agnostic setting), in various settings, up to polylogarithmic factors. We use the following parameters: $m = \text{tr}(\mathcal{H})$, the threshold dimension of the concept class; $d = \text{Lit}(\mathcal{H})$, the Littlestone dimension of $\mathcal{H}$; T , the total number of iterations; and t , the current iteration count (in case the complexity is different for different t). The complexities are up to polylogarithmic factors and C is an absolute constant.

we provide an improper learner with sublinear regret as well as a proper learner whose regret is also sublinear but grows faster than that of the proper learner; see Theorem 3.5. The regret, time per iteration, and ERM oracle calls per iteration of our algorithms are bounded in terms of the Littlestone dimension of the concept class and/or the *threshold dimension* of the concept class, formally defined in Definition A.4 and related to the Littlestone dimension as per Lemma A.6. We note that our algorithms use a weaker oracle than the ERM oracle, called *consistent oracle*, formally given in Definition 3.1. This oracle takes as input a set of examples $(x_1, y_1), \dots, (x_n, y_n)$ and outputs some $h \in \mathcal{H}$ such that $y_i = h(x_i)$ for all i , if such h exists.

We note that, in the worst case, our algorithms require a number of iterations that is exponential in the Littlestone dimension. This is expected due to the lower bounds of Hazan and Koren [2016]. They show that there exist finite concept classes $\mathcal{H}$ such that $\tilde{\Omega}(\sqrt{|\mathcal{H}|})$ -many ERM and function evaluation queries are necessary to obtain sublinear regret in the agnostic proper setting. When $\mathcal{H}$ is finite, its Littlestone dimension is bounded by $\log|\mathcal{H}|$. Thus the total time in the last line of Table 1 should have exponential dependence on the Littlestone dimension.

One of the main applications of online learning is for the purpose of equilibrium learning in games. Indeed, the existence of agnostic, proper learners whose regret grows sublinearly in the number of rounds in the online regression setting with finite concept classes can be used to establish the existence as well as the distributed learnability of minimax equilibria in two-player zero-sum games with a finite number of actions per player and coarse correlated equilibria in multi-player general-sum games with a finite number of actions per player; see e.g. Cesa-Bianchi and Lugosi [2006]. This result has been recently generalized to non-parametric games, i.e. games wherein players have an infinite set of actions, under the condition that a collection of concept classes (one class per player) defined in terms of the game’s payoff matrix have finite Littlestone or sequential fat-shattering dimensions Hanneke et al. [2021], Daskalakis and Golowich [2022], Rakhlin et al. [2010]. However, the resulting algorithms for equilibrium learning in non-parametric games also involve SOA oracles. Our second goal in this paper is the following:

Goal 1.2 *Consider a family of non-parametric two-player zero-sum (respectively multi-player general-sum) games for which minimax (respectively coarse correlated) equilibria exist. For such family of games, are there algorithms for computing an ϵ -approximate minimax (respectively coarse correlated) equilibrium, which run in finite time given access to best-response oracles for each player (a.k.a. ERM calls) as well as standard arithmetic operations?*

Our other main contribution is to provide positive answers to this question, as summarized in Table 2. In particular, Theorem 4.2 provides an algorithm computing an ϵ -approximate minimax equilibrium of a two-player zero-sum game, whose number of iterations, time per iteration and number of best-response (a.k.a. ERM) calls are bounded in terms of the *fat-threshold dimension of the game*, as per Definition A.22. Theorem 4.3 provides similar results for approximate coarse correlated equilibrium computation. Our results apply to both binary-valued and real-valued games. It is important to note that our algorithm for solving zero-sum games is a variant of the double oracle algorithm, proposed by McMahan et al. [2003]. In our variant, the players grow the action sets they consider in alternating rounds of the algorithm as opposed to simultaneously in every round. So our Theorem 4.2 provides conditions under which our variant of the double oracle algorithm converges in games with infinite action spaces, answering a question raised by Gemp et al. [2022] and their references. In a similar vein, our algorithm for solving multi-player games provides a multi-player variant of the double oracle algorithm and conditions under which it converges. Such multi-oracle algorithms are used extensively in practice for equilibrium computation in large games such as multi-agent reinforcement learning; see e.g. the discussion on the policy-space response oracles (PSRO) algorithm in [Gemp et al., 2022] and its references. Again our work provides convergence guarantees when the action sets are infinite.

Setting	Time per iter.	BR calls/it.	#iterations
Minmax, 0-1 valued	t/ϵ^4	$\log t/\epsilon^2$	$C^{\text{Lit}(G)/\epsilon^2} \wedge \epsilon^{-C \text{VC}(G)^2 \text{tr}(G) \log \text{VC}(G)/\epsilon^4}$
Minmax, real valued	t/ϵ^4	$\log t/\epsilon^2$	$C^{\text{sfat}(G,\epsilon)/\epsilon^2} \wedge \epsilon^{-C I(G)^2 \text{fattr}(G,\epsilon)/\epsilon^5}$
CCE, 0-1 valued	kt/ϵ^2	$k \log t/\epsilon^2$	$C^{(k/\epsilon^3) \text{Lit}(G)}$
CCE, real valued	kt/ϵ^2	$k \log t/\epsilon^2$	$C^{(k/\epsilon^3) \text{sfat}(G,\epsilon)}$

Table 2: The table describes the time per iteration, the number of best-response calls per iteration and the number of iterations of our algorithms, up to polylogarithmic factors for finding an $O(\epsilon)$ -approximate Nash in a zero-sum two player game (minmax equilibrium) and Coarse Correlated Equilibrium (CCE) in general games G . Here, $C > 0$ is a universal constant, and Lit, VC, tr, sfat, fat, fattr denote Littlestone, VC, threshold, sequential fat, fat and fat-threshold dimensions of G , $I(G) = \int_0^1 \left(\sqrt{\text{fat}(G, \delta)} d\delta \right)^2$ and $\wedge$ denotes a minimum of two terms.

We want to highlight that we give the first algorithm for general concept classes and games *that can be implemented* with access to an ERM oracle. Though SOA-oracle algorithms only require polynomially-many iterations in the Littlestone dimension, the execution of a single iteration for even simple tasks can take exponentially long. Our ERM-oracle algorithms enjoy fast per-iteration time complexity, and often in terminate in far fewer than the worst-case exponentially-many iterations. It is no surprise that the algorithms that arise from the ERM-oracle model parallel the algorithms actually used in the practice of solving large games (known as double-oracle algorithms; discussed under Goal 2). In contrast, the algorithms that arise from the SOA-oracle model are utterly dissimilar from any practical algorithms. The thesis of this work is that, rather than making oracle assumptions based on what will guarantee a polynomial regret bound, we should instead select an oracle based on what is practically feasible, and then from there, see what regret guarantees are possible.

2 Preliminaries

We include below a shortened version of the preliminaries. See Section A for a full version.

Games A k -player game is a pair $(\mathcal{A}, u)$, where $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_k$ and $u = (u_1, \dots, u_k)$, where each $u_p: \mathcal{A} \rightarrow \mathbb{R}$. Each $\mathcal{A}_p$ is the set of *actions* (a.k.a. *strategies*) available to player p and each u_p is the *utility*, or *payoff*, *function* of player p , which maps the set of *action profiles* $\mathcal{A}$ to the reals. Each player's goal is to maximize their own utility. We denote by $\mathcal{A}_{-p}$ the Cartesian product of $\{\mathcal{A}_j\}_{j \neq p}$. Similarly, for any action $a = (a_1, \dots, a_k) \in \mathcal{A}$, denote by a_{-p} the Cartesian product of $\{a_j\}_{j \neq p}$. A *mixed strategy* for player p is a distribution over $\mathcal{A}_p$. A *zero-sum* game is a two-player game such that $u_1(a, b) = -u_2(a, b)$ for all $a \in \mathcal{A}$ and $b \in \mathcal{B}$. We sometimes compress our notation and represent a zero-sum game as $(\mathcal{A}, \mathcal{B}, u)$ where $u: \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ is a single function representing the utility function of player 2. Player 2 aims to maximize this utility while player 1 aims to minimize this utility.

Definition 2.1 (ϵ -Nash equilibrium and ϵ -CCE) Let $(\mathcal{A}, u)$ denote a game. An ϵ -approximate Nash equilibrium is a collection of probability measures, $\mu_1, \dots, \mu_k$, over $\mathcal{A}_1, \dots, \mathcal{A}_k$, respectively, such that for any player $p \in [k]$ and any $d_p \in \mathcal{A}_p$,

$$\mathbb{E}_{a \sim \mu_1 \times \dots \times \mu_k} [u_p(d_p, a_{-p})] \leq \mathbb{E}_{a \sim \mu_1 \times \dots \times \mu_k} [u_p(a)] - \epsilon.$$

A *Coarse Correlated Equilibrium* (CCE) is a joint measure μ over $\mathcal{A}$ such that for any player p and any $d_p \in \mathcal{A}_p$,

$$\mathbb{E}_{a \sim \mu} [u_p(d_p, a_{-p})] \leq \mathbb{E}_{a \sim \mu} [u_p(a)] - \epsilon.$$

Definition 2.2 (Minimax equilibrium) Given a zero-sum game $G = (\mathcal{A}, \mathcal{B}, u)^2$, we say that G has a *minimax equilibrium* if

$$\inf_{\mu_1 \in \Delta(\mathcal{A})} \sup_{\mu_2 \in \Delta(\mathcal{B})} \mathbb{E}_{a \sim \mu_1, b \sim \mu_2} [u(a, b)] = \sup_{\mu_2 \in \Delta(\mathcal{B})} \inf_{\mu_1 \in \Delta(\mathcal{A})} \mathbb{E}_{a \sim \mu_1, b \sim \mu_2} [u(a, b)] \quad (3)$$

where $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$ denote the set of all probability measures over $\mathcal{A}$ and $\mathcal{B}$ respectively. If Eq. (3) is satisfied, we say the probability measures μ_1, μ_2 optimizing Eq. (3) are a *minimax equilibrium* of G . Denote by $\text{Val}(G)$ *value* of the game, which is the value of both sides of Eq. (3).

²We slightly abuse notation and throughout denote a zero-sum game $(\mathcal{A} \times \mathcal{B}, u)$ as $(\mathcal{A}, \mathcal{B}, u)$

Dimensions Given a real-valued function-class $\mathcal{F}$ over a domain X and $\epsilon > 0$, define the ϵ -fat threshold dimension of $\mathcal{F}$, $\text{fattr}(\mathcal{F}, \epsilon)$, to be the largest $m \geq 0$ such that there exist $f_1, \dots, f_m \in \mathcal{F}$ and $x_1, \dots, x_m \in X$ and a threshold $\theta \in \mathbb{R}$ such that

$$\begin{aligned} f_i(x_j) &\geq \theta + \epsilon & \text{for all } i \leq j \in [d] \\ f_i(x_j) &\leq \theta & \text{for all } i > j \in [d] \end{aligned} \quad (4)$$

For 0-1 classes define the *threshold dimension* of $\mathcal{F}$ as $\text{tr}(\mathcal{F}) = \text{fattr}(\mathcal{F}, 1/2)$. We will also use the notion of VC and Littlestone dimension, $\text{VC}(\mathcal{F})$ and $\text{Lit}(\mathcal{F})$ for real 0-1 valued classes and their real-valued analogues, the fat and sequential fat-shattering dimensions, $\text{fat}(\mathcal{F}, \epsilon)$ and $\text{sfat}(\mathcal{F}, \epsilon)$ (all defined in Section A).

These dimensions can be extended to games: For each player p , her utility $u_p : \mathcal{A}_p \times \mathcal{A}_{-p} \rightarrow \mathbb{R}$ can be thought of as a concept class $\mathcal{F}_p$ over the domain set $X_p = \mathcal{A}_p$, whose concepts $f_{a_{-p}}$ are parametrized by elements $a_{-p} \in \mathcal{A}_{-p}$ and are defined by $f_{a_{-p}}(a_p) := u_p(a_p, a_{-p})$, for each $a_p \in \mathcal{A}_p$. We define the dimension of a game to be the maximal dimension over these utility function classes, where p ranges across all players.

Convex hulls and their dimensions. Denote by $\text{conv}(\mathcal{F})$ the *convex hull* of $\mathcal{F}$, namely the class of all convex combinations of elements from $\mathcal{F}$, by $\Delta(\mathcal{X})$ the set of all probability distributions over $\mathcal{X}$, by $\text{dconv}(\mathcal{F})$ the *dual convex hull* of $\mathcal{F}$ the set of all elements from $\mathcal{F}$ extended to the domain $\Delta(\mathcal{X})$ by taking an expectation, namely, $f(\mu) = \mathbb{E}_{x \sim \mu}[f(x)]$ for all $f \in \mathcal{F}$ and $\mu \in \Delta(X)$, and by $\text{conv} 2(\mathcal{F}) = \text{dconv}(\text{conv}(\mathcal{F}))$. Similarly, for a game G we denote by $\text{conv}(G)$ the game obtained from G where the action sets $\mathcal{A}_p$ are replaced by $\Delta(\mathcal{A}_p)$. We prove the following theorem, which bounds the threshold dimension of $\text{conv} 2(\mathcal{F})$ (see Appendix D for the proof):

Theorem 2.3 For $[0, 1]$ -valued concept classes $\mathcal{F}$, $\text{fattr}(\text{conv}(\mathcal{F}), \epsilon) \leq e^{C \text{sfat}(\mathcal{F}, \epsilon/C)/\epsilon^2}$ and same result holds when conv is replaced by dconv and $\text{conv} 2$ and when $\mathcal{F}$ is replaced by a zero-sum two-player game G . Moreover, for a zero-sum two-player game G , $\text{fattr}(\text{conv}(G), \epsilon) \leq e^{-CI(G)^2 \text{fattr}(G, \epsilon/C)/\epsilon^5}$, where $I(G) := \left(\int_0^1 \sqrt{\text{fat}(G, \epsilon)} d\epsilon \right)^2$ and $C > 0$ is a universal constant. For $\{0, 1\}$ -valued games, we have $\text{tr}(\text{conv}(G), \epsilon) \leq O\left((1/\epsilon)^{C \text{VC}(G)^2 \text{tr}(G) \log(\text{VC}(G))}\right)$

Throughout, we use the notation $\tilde{O}(\cdot)$ for omitting poly-logarithmic factors.

3 Online learning

We describe below our algorithm and results in the online learning setting. First, we define the oracles that the algorithm is allowed to use:

Definition 3.1 (Consistent oracle) For any set of pairs $\{(x_1, y_1), \dots, (x_t, y_t)\}$, the oracle outputs $h \in \mathcal{H}$ such that $h(x_i) = y_i$, for all $i \in [t]$, if exists (otherwise it is undefined).

Definition 3.2 (Value oracle) For any $h \in \mathcal{H}$ and $x \in X$, return $h(x)$.

We notice that the only access that the algorithm has to the function-class $\mathcal{H}$ is via the consistent oracle. Further, the only access to the set X is via the examples generated by the adversary. We present the following theorem on Algorithm 1 which only has access to these two oracles:

Theorem 3.3 (Realizable) Let $\mathcal{H}$ be a 0-1 valued concept class and assume the stream of examples is realizable by some $h^* \in \mathcal{H}$. Then, Algorithm 1, instantiated as an improper learner, has the following bound on its regret and on its number of calls to the consistent oracle: $\min(\mathcal{O}(4^{\text{tr}(\mathcal{H})}), e^{\mathcal{O}(\text{Lit}(\mathcal{H}))})$; if Algorithm 1 is instantiated as a randomized proper learner, the bound changes to $\tilde{O}\left(T^{\frac{2 \text{tr}(\mathcal{H}) + 2}{2 \text{tr}(\mathcal{H}) + 3}}\right)$. Here, $\text{tr}(\mathcal{H})$ and $\text{Lit}(\mathcal{H})$ are the threshold and Littlestone dimensions of $\mathcal{H}$, respectively.

We notice that Theorem 3.3 provides a new proof that $\text{Lit}(\mathcal{H}) \leq \mathcal{O}(4^{\text{tr}(\mathcal{H})})$: indeed, Littlestone dimension equals the smallest regret possible for any improper learner in the realizable setting.³ Notice that the best known bound is $\text{Lit}(\mathcal{H}) \leq 2^{\text{tr}(\mathcal{H})}$ [Alon et al., 2019, Theorem 3] [Hodges et al., 1997, Shelah, 1990]. Additionally, as a direct corollary, we obtain the first polynomial-time algorithm that, given a full description of the class, implements a no-regret learner whose mistake bound depends only on Littlestone's dimension.

Corollary 3.4 There is an online learner who has access to a table of size $n = |\mathcal{H}||\mathcal{X}|$ that describes a binary-valued concept-class $\mathcal{H}$ over $\mathcal{X}$, that has a mistake bound of $e^{\mathcal{O}(\text{Lit}(\mathcal{H}))}$ in the realizable setting and runs in time $O(n)$ per iteration.

³Equivalently, a constant bound on the regret is called *mistake bound*.

For comparison, SOA achieves a mistake bound of $\text{Lit}(\mathcal{H})$, however, its runtime is not polynomial in n : while SOA requires computation of Littlestone’s dimension, under hardness assumptions, it is impossible to even approximate Littlestone’s dimension in time polynomial in n as long as it is $\omega(1)$ [Frances and Litman, 1998, Manurangsi and Rubinstein, 2017, Manurangsi, 2022]. On the other hand, the *halving* algorithm [Shalev-Shwartz and Ben-David, 2014] takes time $O(|\mathcal{H}|)$ per iteration, however, its mistake bound is $O(\log |\mathcal{H}|)$ – this does not depend only on $\text{Lit}(\mathcal{H})$.

Next, we describe the result for the agnostic setting, which is obtained by applying the reduction of Ben-David et al. [2009] from the agnostic to the realizable setting, while using Algorithm 1 as the realizable learner in this reduction (see Section B for the proof):

Theorem 3.5 (Agnostic) *Let $\mathcal{H}$ be a 0-1 valued class. Then, there exists an improper learner which accesses $\mathcal{H}$ only via the consistent and value oracles, that achieves in the improper setting a regret of $\sqrt{T \min\left(\tilde{O}\left(4^{\text{tr}(\mathcal{H})}\right), e^{\mathcal{O}(\text{Lit}(\mathcal{H}))}\right)}$ and $\tilde{O}\left(T^{\frac{2 \text{tr}(\mathcal{H})+3}{2 \text{tr}(\mathcal{H})+4}}\right)$ in the proper setting.*

3.1 Algorithm

Below we describe the algorithm for the realizable setting. It has two variants, proper and improper. For convenience of notation, in the proper setting, we say that the algorithm plays a distribution μ^t over hypotheses $h \in \mathcal{H}$ in each iteration t and suffers loss $\text{Loss}(\mu^t, (x_t, y_t)) := \Pr_{h \sim \mu^t}[h(x_t) \neq y_t]$.⁴ In the improper setting, the algorithm is allowed to take a weighted majority vote over hypotheses and it will select the label that is predicted with the largest probability according to μ . Formally, the learner predicts $\hat{y}_t := \text{Maj}(\mu^t, x_t) = \begin{cases} 1, & \Pr_{h \sim \mu^t}[h(x_t) = 1] \geq 1/2 \\ 0, & \text{otherwise} \end{cases}$.

The algorithm proceeds in phases, where each phase consists of multiple rounds of prediction. In each phase, j , the algorithm plays a distribution over some pool of actions $\{h_1, \dots, h_j\}$. By the end of each phase, the algorithm adds a new action to the pool, h_{j+1} , which is taken as a hypothesis that is consistent with the whole history of elements (x, y) observed by the algorithm throughout all phases.

Next, we describe the phases. Fix a phase j , denote the distribution played by the algorithm at any round t of this phase by μ^t and we describe how to determine μ^t : first, μ^1 is the uniform distribution over all the hypotheses available in this phase: $\{h_1, \dots, h_j\}$. In rounds t when $\text{Loss}(\mu^t, (x_t, y_t)) \geq \epsilon$, the algorithm updates μ^t via a multiplicative-weight update. In the remaining rounds, no update is made, and $\mu^{t+1} \leftarrow \mu^t$. We call this type of update *Lazy multiplicative weights*. The phase j ends once T_j updates have been made. We rely on an auxiliary parameter α to determine the value of T_j . See Algorithm 1 for the main pseudocode and Algorithm 2 for the implementation of each phase.

Algorithm 1 Online Action Insertion

Input: A function class $\mathcal{H}$.

Parameters: $\epsilon, \alpha > 0$.

Subroutines: Consistent oracle (Definition 3.1), Lazy Multiplicative Weights (Algorithm 2).

1. Initialize active action set $\mathcal{H}_1 \leftarrow \{h_1\}$ for arbitrary $h_1 \in \mathcal{H}$.
 2. **For** phase $j = 1, 2, \dots, J$:
 - (a) Instantiate the algorithm `Lazy Multiplicative Weights` with the function-class $\mathcal{H}_j$ and parameters $T_j = \left\lceil \frac{C \log |\mathcal{H}_j|}{\alpha^2} \right\rceil$ and ϵ ($C > 0$ is a constant), in order to predict in the next classification rounds, until the execution of `Lazy Multiplicative Weights` terminates.
 - (b) Call the consistent oracle to obtain $h_{j+1} \in \mathcal{H}$ that is consistent with all pairs (x, y) observed in all previous rounds throughout all phases.
 - (c) Update the active action set $\mathcal{H}_{j+1} \leftarrow \mathcal{H}_j \cup \{h_{j+1}\}$.
-

⁴The alternative would be to sample one h from $\mathcal{H}$ - these two notions are equivalent if one is interested in an expected loss.

Algorithm 2 Lazy Multiplicative Weights

Input: A finite set of functions $\mathcal{G} = \{g_1, \dots, g_n\} \subseteq \{0, 1\}^{\mathcal{X}}$.

Parameters: K rounds with an update, accuracy parameter $\epsilon > 0$, learning rate $\eta = \sqrt{\frac{\log|\mathcal{G}|}{2K}}$.

1. Initialize a uniform probability distribution $\mu^1 = (\mu_1^1, \dots, \mu_n^1) = (\mu^1(g_1), \dots, \mu^1(g_n))$ over $\mathcal{G}$, where $\mu_i^1 \leftarrow \frac{1}{n}$ for all $i = 1, \dots, n$.
 2. $k \leftarrow 0$.
 3. **For** $t = 1, 2, \dots$:
 - (a) **If proper:** predict μ^t , observe (x_t, y_t) and suffer $\text{Loss}(\mu^t, (x_t, y_t))$.
Else (if improper): Observe x_t , predict $\hat{y}_t := \text{Maj}(\mu^t, x_t)$, observe y_t and suffer a loss $\mathbb{I}(\hat{y}_t \neq y_t)$.
 - (b) **If** $\text{Loss}(\mu^t, (x_t, y_t)) \geq \epsilon$:
 - i. $\forall i, \mu_i^{t+1} \leftarrow \mu_i^t \exp(-\eta|g_i(x_t) - y_t|) / \sum_{j=1}^n \mu_j^t \exp(-\eta|g_j(x_t) - y_t|)$.
 - ii. $k \leftarrow k + 1$.
 - iii. **If** $k = K$ then **Return**.
Else: make no update: $\mu^{t+1} \leftarrow \mu^t$.
-

3.2 Proof

Proper learner, bound in terms of $\text{tr}(\mathcal{H})$ We say that a distribution μ over $\mathcal{H}$ makes an ϵ -mistake on $(x, y) \in X \times \{0, 1\}$ if $\text{Loss}(\mu, (x, y)) \geq \epsilon$. We say that the algorithm makes an ϵ -mistake on iteration t if its prediction μ^t makes an ϵ -mistake on (x_t, y_t) . Our first goal would be to bound the number of ϵ mistakes made by the algorithm. This will be done by bounding the number of phases observed by the algorithm. Assume that the algorithm has observed more than J phases, and for every $j \in [J]$, denote by Z_j the set of elements (x, y) observed by the algorithm on phase j , on which the algorithm made an ϵ -mistake. For any $h \in \mathcal{H}$, denote $\text{Loss}(h, Z_j) = \frac{1}{|Z_j|} \sum_{(x, y) \in Z_j} \mathbb{I}(h(x) \neq y)$ as the fraction of mistakes of h on the elements of Z_j . Recall the pool of functions $h_1, \dots, h_J$ maintained by the algorithm in phase J . In the following lemma, we argue that the loss applied on the functions $h_1, \dots, h_J$ and the sets $Z_1, \dots, Z_J$, have the following threshold behavior:

Lemma 3.6 Let $h_1, \dots, h_J$ and $Z_1, \dots, Z_J$ be defined as in the preceding paragraph. Then, for all $i, j \in [J]$

$$\begin{aligned} \text{Loss}(h_i, Z_j) &\geq \epsilon - \alpha && \text{if } i \leq j \\ \text{Loss}(h_i, Z_j) &= 0 && \text{if } i > j \end{aligned} \tag{5}$$

Proof For $i > j$, $\text{Loss}(h_i, Z_j) = 0$ since h_i is the output of a consistent oracle that observed all the examples from previous rounds, including those from Z_j . Next, we proceed with $i \leq j$. To prove for $i \leq j$, fix round j , and notice that the Lazy multiplicative weights algorithm, restricted to rounds where the adversary played actions in Z_j , behaves exactly as the original multiplicative weights algorithm. There are T_j such rounds, and let $\mu^{t_1}, \dots, \mu^{t_{T_j}}$ be the learner's distributions over $\mathcal{H}_j$ for these rounds. From the Multiplicative Weights guarantee (see Lemma A.16) and by definition of T_j , the regret of the learner is upper bounded by $\mathcal{O}(\sqrt{T_j \log J}) \leq T_j \alpha$:

$$\begin{aligned} \epsilon T_j &\leq \sum_{k=1}^{T_j} \text{Loss}(\mu^{t_k}, (x_{t_k}, y_{t_k})) \leq \min_{h \in \{h_1, \dots, h_j\}} \sum_{k=1}^{T_j} |h(x_{t_k}) - y_{t_k}| + T_j \alpha \\ &= T_j \min_{h \in \{h_1, \dots, h_j\}} \text{Loss}(h, Z_j) + T_j \alpha, \end{aligned}$$

where the first inequality holds since the algorithm makes an ϵ -mistake in each of these rounds, the second due to the regret guarantee of multiplicative weights and the third since $Z_j = \{(x_{t_k}, y_{t_k})\}_{k \in [T_j]}$. By rearranging terms, this concludes that for all $i \leq j$, $\text{Loss}(h_i, Z_j) \geq \epsilon - \alpha$. $\blacksquare$

Let h^* be the function that is realizable with all the examples provided by the adversary throughout all phases, namely, $h^*(x_t) = y_t$ for observed (x_t, y_t) . Such h^* exists due to the realizability assumption. For $i \in [J]$, define the function $f_i(x) = (h_i \oplus h^*)(x) := \mathbb{I}(h_i(x) \neq h^*(x))$, and notice that Eq. (5) implies the following, where

$\text{unif}(Z_j)$ denotes a uniform distribution over Z_j and only x is sampled (rather than the pair (x, y)):

$$\begin{aligned}\mathbb{E}_{x \sim \text{unif}(Z_j)}[f_i(x)] &\geq \epsilon - \alpha & \text{if } i \leq j \\ \mathbb{E}_{x \sim \text{unif}(Z_j)}[f_i(x)] &= 0 & \text{if } i > j\end{aligned}\tag{6}$$

By a simple inductive argument, Eq. (6) is sufficient to imply a lower bound on the threshold dimension of $\{f_1, \dots, f_J\}$ in terms of J , which equivalently bounds the number of phases J in terms of $\text{tr}(\{f_1, \dots, f_J\})$:

Lemma 3.7 *Let $f_1, \dots, f_J: X \rightarrow \{0, 1\}$ and $Z_1, \dots, Z_J \subset X$ such that Eq. (6) holds, let $m \in \mathbb{N}$ and assume that $J \geq \sum_{k=1}^m \left(\frac{1}{\epsilon - \alpha}\right)^k$. Then, $\text{tr}(\{f_1, \dots, f_J\}) \geq m$.*

Proof We prove the claim by induction on m . For $m = 0$, trivially there exists a threshold game of size 0. For the induction step, assume that the claim holds for $m - 1$. To prove for m , we need to show that there exist functions $g_1, \dots, g_m \in \{f_1, \dots, f_J\}$ and $x_1, \dots, x_m \in X$ such that $g_i(x_j) = \mathbb{I}(i \leq j)$. We will start by determining x_m and g_m and the remaining elements f_i and g_i will be taken from the induction hypothesis. To determine x_m , we notice that from double summation and from the condition of this lemma,

$$\mathbb{E}_{x \sim \text{unif}(Z_J)}|\{i \in [J]: f_i(x) = 1\}| = \sum_{i=1}^J \mathbb{E}_{x \sim \text{unif}(Z_J)}[f_i(x)] \geq J(\epsilon - \alpha).$$

This implies that there exists $x' \in Z_J$ such that the set $I_{x'} := \{i \in [J]: f_i(x') = 1\}$ is of cardinality at least $J(\epsilon - \alpha)$. We take $x_m = x'$. The functions $g_1, \dots, g_m$ will be taken from $\{g_i\}_{i \in I_{x'}}$ and this will guarantee that $g_i(x_m) = 1 = \mathbb{I}(i \leq m)$. We denote $i' = \max I_{x'}$ and set $g_m = f_{i'}$. The remaining elements $x_1, \dots, x_{m-1}$ will be taken from the sets $\{Z_j\}_{j \in I'}$ where $I' = I_{x'} \setminus \{i'\}$. This will guarantee that for all $j < m$, $g_m(x_j) = 0 = \mathbb{I}(m \leq j)$, using Eq. (6). We would like to apply the induction hypothesis on the functions $\{f_i\}_{i \in I'}$ and the sets $\{Z_j\}_{j \in I'}$. Indeed, the induction hypothesis can be applied since $|I'| = |I_{x'}| - 1 \geq (\epsilon - \alpha)J - 1 \geq \sum_{k=1}^{m-1} \left(\frac{1}{\epsilon - \alpha}\right)^k$,

which follows by the computed bound on $|I_{x'}|$ and by the assumption $J \geq \sum_{k=1}^m \left(\frac{1}{\epsilon - \alpha}\right)^k$. Hence, the induction hypothesis yields that $\text{tr}(\{f_i\}_{i \in I'}) \geq m - 1$, which imply the existence of functions $g_1, \dots, g_{m-1} \in \{f_i\}_{i \in I'}$ and elements $x_1, \dots, x_{m-1} \in X$ such that $g_i(x_j) = \mathbb{I}(i \leq j)$ for all $i, j \in [m - 1]$. Together with g_m and x_m , the arguments above imply that $g_i(x_j) = \mathbb{I}(i \leq j)$ for all $i, j \in [m]$, which concludes the induction step. The proof follows. $\blacksquare$

We are now ready to bound the number of ϵ -mistakes of the algorithm. Combining Lemma 3.6, Eq. 6 and Lemma 3.7, we obtain that $J < \sum_{k=1}^{\text{tr}(\{f_1, \dots, f_J\})+1} (\epsilon - \alpha)^{-k} \leq 2(\epsilon - \alpha)^{-\text{tr}(\{f_1, \dots, f_J\})-1}$, assuming that $\epsilon \leq 1/2$. Recall that $f_i = h_i \oplus h^*$ and define $\mathcal{H} \oplus h^* := \{h \oplus h^*: h \in \mathcal{H}\}$. Lemma F.9 argues that $\text{tr}(\mathcal{H} \oplus h^*) \leq 2 \text{tr}(\mathcal{H}) + 1$, consequently, $J \leq 2(\epsilon - \alpha)^{-2 \text{tr}(\mathcal{H}) - 2}$. Recall that the number of ϵ -mistakes in each phase j is at most $\mathcal{O}(\log(j)/\alpha^2)$. Assuming that there are at least J phases, the total number of mistakes is bounded by $\mathcal{O}(J \log J / \alpha^2)$. Substituting the bound on J and optimizing over α , one obtains that the number of ϵ -mistakes is at most $\mathcal{O}(\text{tr}(\mathcal{H})^3 \epsilon^{-2 \text{tr}(\mathcal{H}) - 2})$ (see Lemma B.1). The total regret of the algorithm is bounded by the number of ϵ -mistakes, plus ϵT , in order to account for the loss for less-than- ϵ -mistakes. Setting $\epsilon = T^{-1/(2 \text{tr}(\mathcal{H}) + 3)}$ yields the bound of $\tilde{\mathcal{O}}\left(T^{\frac{2 \text{tr}(\mathcal{H}) + 2}{2 \text{tr}(\mathcal{H}) + 3}}\right)$.

Bounds for the improper learner. Recall that the improper learner takes a majority vote over μ^t , therefore, it makes a mistake only if μ^t makes a $1/2$ -mistake. Substituting $\epsilon = 1/2$ in the bound on the number of ϵ -mistakes of the proper learner, one obtains the desired bound of $\tilde{\mathcal{O}}(4^{\text{tr}(\mathcal{H})})$. For the bound in terms of Littlestone's dimension, we use the fact that Eq. (6) implies that $\text{fat}_{\text{tr}}(\text{dconv}(\mathcal{H} \oplus h^*), \epsilon - \alpha) \geq J$ (for the dual convex $\text{dconv}(\cdot)$ see Section 2). Indeed, the functions $f_1, \dots, f_J \in \mathcal{H} \oplus h^*$ and the elements $\left\{\sum_{x \in Z_j} \frac{x}{|Z_j|}\right\}_{j \in [J]} \in \text{conv}(X)$, satisfy Eq. (4) in the definition of the ϵ -fat threshold dimension. The bound in Theorem 2.3 on the threshold dimension of the dual convex concludes the proof. See Lemma B.2 for the proof of this argument.

4 Computing equilibria in games

In this section, we will provide an algorithm for computing approximate Nash equilibria in zero-sum two-player games and coarse correlated equilibria in general multi-player games. The algorithm is allowed to use the following oracles:

Definition 4.1 (Best-response and value oracles) An ϵ -best response oracle in a game $(\mathcal{A}_1, \dots, \mathcal{A}_k, u)$ receives a player p , a finitely-supported distribution μ_{-p} over the Cartesian product $\mathcal{A}_{-p} = \prod_{q \neq p} \mathcal{A}_q$ and outputs an action $\hat{a}_p \in \mathcal{A}_p$ that ϵ -maximizes the player's utility against a random sample from the distribution:

$$\mathbb{E}_{a_{-p} \sim \mu_{-p}}[u_p(\hat{a}_p, a_{-p})] \geq \sup_{a_p \in \mathcal{A}_p} \mathbb{E}_{a_{-p} \sim \mu_{-p}}[u_p(a_p, a_{-p})] - \epsilon.$$

A *value oracle* receives a player p and actions $(a_1, \dots, a_k)$ and outputs $u_p(a_1, \dots, a_k)$.

We notice that a *best-response oracle* can be viewed as an ERM oracle, that maximizes reward instead of minimizing loss. We note that the only access that the algorithms have to the game is via these oracles and they are not allowed to access the action sets apart via these oracles. This aims to capture the scenario that their action sets are large or perhaps even infinite. Algorithms are given under the assumption that the sequential fat-shattering dimension $\text{sfat}(G, \epsilon)$ of the game is finite. It has been shown that if the dimension is infinite, an equilibrium might not exist [Hanneke et al., 2021, Daskalakis and Golowich, 2022].⁵ In the following two sections, we present our two results: in Section 4.1 we present the result for computing a Nash equilibrium in two-player zero-sum games. In Section 4.2 we study general sum games and show how to compute a coarse correlated equilibrium, whereas computing Nash is a significantly harder problem: for finite games, it is PPAD-hard [Daskalakis et al., 2009, Chen et al., 2009] whereas CCE is poly-time computable [Hart and Mas-Colell, 2000, Hart and Schmeidler, 1989]

4.1 Approximating Nash equilibrium in zero-sum two-player games

Theorem 4.2 Let $G = (\mathcal{A}, \mathcal{B}, u)$ be a zero-sum two-player game, where $u: \mathcal{A} \times \mathcal{B} \rightarrow [0, 1]$ and let $\epsilon > 0$. There is an algorithm to find an $O(\epsilon)$ -Nash for this game using the following number of ϵ -best response oracle calls (assuming this number is finite):

$$\mathcal{O}\left(\min\left(e^{C \text{sfat}(G, \epsilon/C)/\epsilon^2}, (1/\epsilon)^{CI(G)^2 \text{fattr}(G, \epsilon/C)/\epsilon^5}\right); I(G) = \left(\int_0^1 \sqrt{\text{fat}(G, \epsilon)}\right)^2\right).$$

As a first step, we argue that it is possible to find an approximate Nash for a game $(\mathcal{A}, \mathcal{B}, u)$ where one of the players has a finite action set and random access to that set and the second player has an infinite action-set and an ϵ -best response oracle. We use the reduction from online learning to equilibrium computation which states that if two players play an algorithm whose regret behaves as $o(T)$ in a zero-sum game, then the pair of uniform distributions over their actions converges to a Nash equilibrium as $T \rightarrow \infty$. We use the common technique where the player with the finite number of actions can play a no-regret algorithm (such as multiplicative weight update) and the second player plays best-response, which is a no-regret algorithm as well. We notice that a modification of this technique does not work when both players play best response since best-response is no-regret only if it is played second, after observing the opponent's action.

Next, we will provide the algorithm to compute an ϵ -Nash in a zero-sum game $(\mathcal{A}, \mathcal{B}, u)$ where both $\mathcal{A}$ and $\mathcal{B}$ are possibly infinite or very large. The algorithm was inspired by the proof of existence of minimax equilibria by Hanneke et al. [2021]. It gradually accumulates actions for each of the two players until reaching a sufficiently large finite subgame whose ϵ -Nash equilibrium approximates the Nash of the complete game. In particular, at each iteration t , Player 1 will hold a finite subset of actions, $A_t \subseteq \mathcal{A}$, that grows as t increases, namely, $A_0 \subseteq A_1 \subseteq A_2 \subseteq \dots$. Similarly, Player 2 will hold finite sets $B_0 \subseteq B_1 \subseteq \dots$. The initial sets A_0 and B_0 are of cardinality 1 and contain a single arbitrary element from $\mathcal{A}$ and $\mathcal{B}$, respectively. Then, the players take turns adding actions to change the value of the game in their favor. In particular each iteration t begins where Players 1 and 2 hold the sets of actions A_{t-1} and B_{t-1} , respectively. Then, Player 1, whose aim is to minimize the utility u and the value Val , finds a set of actions $A_t \supseteq A_{t-1}$, such that $\text{Val}(A_t, B_{t-1}) \leq \text{Val}(A_{t-1}, B_{t-1}) - \Omega(\epsilon)$. This is done by computing an approximate Nash equilibrium for the game $(\mathcal{A}, B_{t-1}, u)$ where Player 1 is unrestricted and player 2 is restricted to her finite set B_t , and adding the support of this approximate Nash to A_{t-1} , thus creating A_t . Then, similarly, Player 2 responds by finding a set of actions, $B_t \supseteq B_{t-1}$ that increase the value of the game, namely, $\text{Val}(A_t, B_t) \geq \text{Val}(A_t, B_{t-1}) + \Omega(\epsilon)$ thus changing the value in her favor. The algorithm stops when no player can improve the value by more than ϵ . See Algorithm 3 for the pseudocode.

⁵There are some delicacies in the statement of when there exists, or there does not exist, an equilibrium and the exact conditions are not known. See Section A.4 and [Hanneke et al., 2021, Daskalakis and Golowich, 2022] for more discussion.

Algorithm 3 ϵ -approximate Nash Equilibrium for a zero-sum game

Input: A zero-sum game $(\mathcal{A}, \mathcal{B}, u)$, a parameter $\epsilon > 0$.

Subroutines:

- **Nash:** Receives sets of actions A and B for both players, and an $\epsilon > 0$, where either A is finite or B is. It outputs an ϵ -Nash for the subgame (A, B, u) (Algorithm 5)
- **Val:** Receives finite sets of actions of the players. Returns the value of this finite subgame.

1. $A_0 \leftarrow \{a\}, B_0 \leftarrow \{b\}$, where $a \in \mathcal{A}$ and $b \in \mathcal{B}$ are arbitrary actions.
 2. **For** $t = 1, 2, \dots$
 - (a) $(\mu^{t,1}, \mu^{t,2}) \leftarrow \text{Nash}(\mathcal{A}, B_{t-1}, \epsilon)$.
 - (b) $A_t \leftarrow A_{t-1} \cup \text{Support}(\mu^{t,1})$.
 - (c) $(\xi^{t,1}, \xi^{t,2}) \leftarrow \text{Nash}(A_t, \mathcal{B}, \epsilon)$.
 - (d) $B_t \leftarrow B_{t-1} \cup \text{Support}(\xi^{t,2})$.
 - (e) **If** $\text{Val}(A_t, B_{t-1}) \geq \text{Val}(A_{t-1}, B_{t-1}) - \epsilon$ **or** $\text{Val}(A_t, B_t) \leq \text{Val}(A_t, B_{t-1}) + \epsilon$:
 - i. **Return** $(\xi^{t,1}, \mu^{t,2})$.
-

The first statement that is proven is that the output of the algorithm is an approximate Nash equilibrium. Intuitively, this follows from the fact that when the game ends, no player can add action to drastically change their value. The second statement is that the algorithm eventually ends. This is proven using the fact that $\text{fattn}(\text{conv}(G), \epsilon)$ is finite, which follows from Theorem 2.3. Intuitively, the finiteness of $\text{fattn}(\text{conv}(G), \epsilon)$ implies that there exists no sequence of mixed strategies (i.e. distributions over actions) $\mu_1, \mu_2, \dots$ for player 1 and $\xi_1, \xi_2, \dots$ for player 2 and a threshold θ such that $u(\mu_i, \xi_j) \geq \theta + \epsilon$ if $i \leq j$ and $u(\mu_i, \xi_j) \leq \theta$ if $i > j$ – this can be shown to imply that the players cannot keep adding actions to increase the value by $\Omega(\epsilon)$ to their favor indefinitely.

Lastly, we explain how the two bounds on $\text{fattn}(\text{conv}(G), \epsilon)$ are derived in Theorem 2.3: the bound in terms of the threshold dimension of G is an adaptation of the beautiful technique of Hanneke et al. [2021] using Ramsey numbers, that was used to prove the existence of minimax in a wide class of zero-sum games. The bound in terms of the sequential fat-shattering dimension uses a standard technique of comparison to the sequential Rademacher complexity of Rakhlin et al. [2010].

4.2 CCE in multi-player general-sum games

We prove the following theorem for finding a CCE in a general game:

Theorem 4.3 *Let $G = (\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_k, u = (u_1, \dots, u_k))$ be a multi-player game. Assume that utilities are bounded $u_p : \mathcal{A} \rightarrow [0, 1]$, and let $\epsilon > 0$. Then, Algorithm 8 executed with parameters G, ϵ will compute an $O(\epsilon) - \text{CCE}$ for the game using the following number of ϵ -best response oracle calls: $O\left(e^{C(k/\epsilon^3) \text{sfat}(G, \epsilon/C)}\right)$.*

In order to extend our algorithm for finding a Nash in a two-player zero-sum for our setting, we use the brilliant reduction of Papadimitriou and Roughgarden [2008] from multiplayer games to a two-player game which we term *the CCE game*. Here, player 1 selects an entire strategy profile a and Player 2 selects a player p and an alternative action d_p for player p . The utility for player 2 in the CCE game, corresponds to the gain in utility (in the original game) made by player p deviating to action d_p when everyone is playing according to strategy profile a . This utility can be defined formally as a matrix whose entries are indexed by a and (p, d_p) , as follows:

Definition 4.4 (The CCE matrix of a game) For a game $G = (\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_k, u = (u_1, \dots, u_k))$, the CCE matrix $\mathbf{M}_{\text{CCE}}^G : \mathcal{A} \times \left(\bigcup_p \mathcal{A}_p\right)$ is defined, for $a \in \mathcal{A}$ and $(p, d_p) \in \bigcup_p \mathcal{A}_p$ as

$$\mathbf{M}_{\text{CCE}}^G[a, (p, d_p)] = u_p(d_p, a_{-p}) - u_p(a_p, a_{-p})$$

The goal of Player 1 is to minimize $\mathbf{M}_{\text{CCE}}^G[a, (p, d_p)]$. An existence of CCE in the original game implies that there exists a distribution μ^* over strategy profiles in that game, such that no deviation is profitable for any of the players. In particular, if Player 1 plays μ^* , this guarantees that the utility against any action of Player 2 is at most 0, which implies that the value of the CCE game, if exists, is at most 0. Similarly, the mixed-strategy μ

played by Player 1 in any ϵ -approximate Nash equilibrium for the **CCE** game, constitutes an ϵ -approximate CCE in the original game. Therefore, we would like to apply Algorithm 3 to find a Nash equilibrium of the **CCE** game. Yet, this requires two things: (1) Bounding the various dimensions of the **CCE** matrix in terms of those of the original game, which is proved using closure properties of these dimensions that appear in Appendix F; and (2) Implementing best-response oracles for the players of the **CCE** game, based on best-response and value oracles for the original game. We notice that since the utility of Player 2 corresponds to the deviations of players from a strategy-profile given by Player 1, the best deviation can be simulated using a best response oracle for the original game. For Player 1, we will not simulate a best-response oracle. Rather, recall that such an oracle is used only for computing an approximate Nash equilibrium for the half-infinite game, where Player 2 is restricted to play from finitely many actions and Player 1 is unrestricted. Hence, it is sufficient to compute such a Nash equilibrium, a task that can be reduced to a computation of a CCE in a finite subgame of the original game.⁶ In that finite game, the set of actions available to the players correspond to the actions of Player 1 in the **CCE** game.

Acknowledgements

Most of this work was completed during an internship at Archimedes AI Research Center. The authors would like to thank them for their kind support. Angelos Assos, Yuval Dagan, Costis Daskalakis, and Maxwell Fishelson are supported by NSF Awards CCF-1901292, DMS-2022448 and DMS2134108, a Simons Investigator Award, the Simons Collaboration on the Theory of Algorithmic Fairness and a DSTA grant. Idan Attias is supported by the Vatav Scholarship from the Israeli Council for Higher Education and by the Kreitman School of Advanced Graduate Studies. Yuval Dagan gratefully acknowledges the NSF's support of FODSI through grant DMS-2023505.

References

- Noga Alon, Roi Livni, Maryanthe Malliaris, and Shay Moran. Private pac learning implies finite littlestone dimension. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 852–860, 2019. 5, 13
- Noga Alon, Omri Ben-Eliezer, Yuval Dagan, Shay Moran, Moni Naor, and Eylon Yogev. Adversarial laws of large numbers and optimal regret in online classification. In *Proceedings of the 53rd annual ACM SIGACT symposium on theory of computing*, pages 447–455, 2021. 2
- Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of computing*, 8(1):121–164, 2012. 16
- Idan Attias, Aryeh Kontorovich, and Yishay Mansour. Improved generalization bounds for adversarially robust learning. *The Journal of Machine Learning Research*, 23(1):7897–7927, 2022. 20
- Vishal Balaji, Powers Lamb, Andrew Lott, Dhruv Patel, Alex Rice, Sakshi Singh, and Rose Ward. The pigeonhole principle and multicolor ramsey numbers. *arXiv preprint arXiv:2108.08410*, 2021. 24
- Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. Agnostic online learning. In *COLT*, volume 3, page 1, 2009. 2, 6, 19
- Adam Block, Yuval Dagan, and Alexander Rakhlin. Majorizing measures, sequential complexities, and online learning. In *Conference on Learning Theory*, pages 587–590. PMLR, 2021. 2, 27
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, Cambridge; New York, 2006. 3
- Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009. 9
- Constantinos Daskalakis and Noah Golowich. Fast rates for nonparametric online learning: from realizability to learning in games. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 846–859, 2022. 2, 3, 9, 14, 38
- Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The Complexity of Computing a Nash Equilibrium. *SIAM Journal on Computing*, 39(1):195, 2009. 9

⁶We note that computing a **CCE** does not quite yield a Nash equilibrium in the **CCE**-subgame, yet, guarantees to find a mixed strategy of value 0 for Player 1 which is sufficient for the purpose of the algorithm.

- Uriel Feige, Yishay Mansour, and Robert Schapire. Learning and inference in the presence of corrupted inputs. In *Conference on Learning Theory*, pages 637–657. PMLR, 2015. [20](#)
- Moti Frances and Ami Litman. Optimal mistake bound learning is hard. *Information and Computation*, 144(1): 66–82, 1998. [6](#)
- Ian M. Gemp, Thomas W. Anthony, Yoram Bachrach, Avishkar Bhoopchand, Kalesha Bullard, Jerome T. Connor, Vibhavari Dasagi, Bart De Vylder, Edgar A. Duéñez-Guzmán, Romuald Elie, Richard Everett, Daniel Hennes, Edward Hughes, Mina Khan, Marc Lanctot, Kate Larson, Guy Lever, Siqi Liu, Luke Marris, Kevin R. McKee, Paul Muller, Julien Pérolat, Florian Strub, Andrea Tacchetti, Eugene Tarassov, Zhe Wang, and Karl Tuyls. Developing, evaluating and scaling learning agents in multi-agent environments. *AI Commun.*, 35(4):271–284, 2022. [3](#)
- Steve Hanneke, Roi Livni, and Shay Moran. Online learning with simple predictors and a combinatorial characterization of minimax in 0/1 games. In *Conference on Learning Theory*, pages 2289–2314. PMLR, 2021. [2](#), [3](#), [9](#), [10](#), [18](#)
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000. [9](#)
- Sergiu Hart and David Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14(1):18–25, 1989. [9](#)
- Elad Hazan and Tomer Koren. The computational power of optimization in online learning. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 128–141, 2016. [3](#)
- Wilfrid Hodges et al. *A shorter model theory*. Cambridge university press, 1997. [5](#), [13](#)
- Sham Kakade and Ambuj Tewari. Fat shattering dimension and covering numbers. Lecture 16 notes from CMSC 35900 (Spring 2008) Learning Theory, 2008. URL <https://home.ttic.edu/~tewari/lectures/lecture16.pdf>. [34](#)
- Nick Littlestone. Learning Quickly When Irrelevant Attributes Abound: A New Linear-threshold Algorithm. *Machine Learning*, 2:285–318, 1988. [2](#), [16](#)
- Pasin Manurangsi. Improved inapproximability of vc dimension and littlestone’s dimension via (unbalanced) biclique. *arXiv preprint arXiv:2211.01443*, 2022. [2](#), [6](#)
- Pasin Manurangsi and Aviad Rubinfeld. Inapproximability of VC dimension and Littlestone’s dimension. In *Conference on Learning Theory*, pages 1432–1460. PMLR, 2017. [2](#), [6](#)
- H Brendan McMahan, Geoffrey J Gordon, and Avrim Blum. Planning in the presence of cost functions controlled by an adversary. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 536–543, 2003. [3](#)
- Christos H Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):1–29, 2008. [10](#)
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Random averages, combinatorial parameters, and learnability. *Advances in Neural Information Processing Systems*, 23, 2010. [2](#), [3](#), [10](#), [27](#)
- Mark Rudelson and Roman Vershynin. Combinatorics of random processes and sections of convex bodies. *Annals of Mathematics*, pages 603–648, 2006. [15](#)
- Robert E Schapire and Yoav Freund. Boosting: Foundations and algorithms. *Kybernetes*, 42(1):164–166, 2013. [16](#), [18](#)
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014. [6](#)
- Saharon Shelah. *Classification theory: and the number of non-isomorphic models*. Elsevier, 1990. [5](#), [13](#)

A Preliminaries

A.1 Function Classes and Dimensions

We work with two types of function classes: 0-1 function classes, and the more general real-valued function classes. We define them here as well as their accompanying notions of dimensionality.

A.1.1 0-1 Function classes

Definition A.1 (0-1 function class) Define a 0-1 function class $\mathcal{F}$ to be a set of *concepts*, $f: X \rightarrow \{0, 1\}$, where X is called a *domain set* and $\{0, 1\}$ is a *label set*.

The three important notions of dimensionality of 0-1 function classes are the following.

Definition A.2 (VC Dimension) For a 0-1 function class $\mathcal{F}$, denote its VC Dimension $\text{VC}(\mathcal{F})$ to be the maximal d (possibly infinite) such that there exists a magnitude- d subset $\{x_1, \dots, x_d\} \subseteq X$ satisfying the following. For all binary strings $b \in \{0, 1\}^d$, there exists $f \in \mathcal{F}$ that satisfies

$$f(x_i) = b_j \quad \text{for all } j \in [d]$$

Definition A.3 (Littlestone Dimension) For a 0-1 function class $\mathcal{F}$, denote its Littlestone Dimension $\text{Lit}(\mathcal{F})$ to be the maximum depth d of a complete, binary tree $T = (V, E)$, such that the children of any internal node are ordered by *left* and *right* and such that any internal node $v \in V$ is labeled by $x(v) \in X$ and every leaf is labeled by $f_v \in \mathcal{F}$, and these labeling functions satisfy the following: let $v_1 - v_2 \dots - v_d - v_{d+1} = \ell$ be a root-to-leaf path along the tree, where v_1 is the root and v_{d+1} is a leaf. Then, for any $i \in [d]$:

$$\begin{aligned} f_\ell(x(v_i)) &= 1 && \text{if } v_{i+1} \text{ is a left child of } v_i \\ f_\ell(x(v_i)) &= 0 && \text{if } v_{i+1} \text{ is a right child of } v_i \end{aligned}$$

We also utilize the concept of threshold dimension, a concept similar to VC dimension but only requiring the existence of a hypothesis for each “threshold binary string” of the form $(0, \dots, 0, 1, \dots, 1)$.

Definition A.4 (Threshold Dimension) For a 0-1 function class $\mathcal{F}$, denote its Threshold Dimension $\text{tr}(\mathcal{F})$ to be the maximal d (possibly infinite) such that there exist magnitude- d subsets $\{f_1, \dots, f_d\} \subseteq \mathcal{F}$ and $\{x_1, \dots, x_d\} \subseteq X$ satisfying the following.

$$\begin{aligned} f_i(x_j) &= 1 && \text{for all } i \leq j \in [d] \\ f_i(x_j) &= 0 && \text{for all } i > j \in [d] \end{aligned}$$

These dimensions are related in the following ways.

Lemma A.5 $\text{VC}(\mathcal{F}) \leq \min(\text{Lit}(\mathcal{F}), \text{tr}(\mathcal{F}))$

The Lemma follows from standard arguments.

Lemma A.6 $\log \text{Lit}(\mathcal{F}) \leq \text{tr}(\mathcal{F}) \leq 2^{\text{Lit}(\mathcal{F})}$

For the proof, see [Alon et al., 2019, Theorem 3] and the references therein Hodges et al. [1997], Shelah [1990]

A.1.2 Real-valued function classes

The following definitions are analogous to those for 0-1 function classes.

Definition A.7 (Real-valued function class) Define a real-valued function class $\mathcal{F}$ as a set of concepts $f: X \rightarrow \mathbb{R}$. Similarly, define a $[0, 1]$ -valued function class as a collection of functions $f: X \rightarrow [0, 1]$.

When defining the three analogous concepts of dimensionality in the real-valued setting, we introduce a margin parameter ϵ to ensure hypotheses are sufficiently distinct. Let us start with the real-valued analogue of VC dimension: “ ϵ -fat-shattering dimension”.

Definition A.8 (ϵ -fat-shattering dimension) For a real-valued function class $\mathcal{F}$, denote its ϵ -fat-shattering dimension $\text{fat}(\mathcal{F}, \epsilon)$ to be the maximal d (possibly infinite) such that there exists a magnitude- d subset $\{x_1, \dots, x_d\} \subseteq X$ and witnesses $(\theta_1, \dots, \theta_d) \in \mathbb{R}^d$ satisfying the following. For all binary strings $b \in \{0, 1\}^d$, there exists $f_b \in \mathcal{F}$ with

$$\begin{aligned} f_b(x_j) &\geq \theta_j + \epsilon && \text{for all } j \in [d] \text{ with } b_j = 1 \\ f_b(x_j) &\leq \theta_j && \text{for all } j \in [d] \text{ with } b_j = 0 \end{aligned} \tag{7}$$

Next, the real-valued analogue of Littlestone dimension: “ ϵ -sequential-fat-shattering dimension”.

Definition A.9 (ϵ -sequential-fat-shattering dimension) For a real-valued function class $\mathcal{F}$, denote its ϵ -sequential-fat-shattering dimension $\text{sfat}(\mathcal{F}, \epsilon)$ to be the maximum depth of a complete binary tree $T = (V, E)$, whose internal nodes $v \in V$ are labeled by elements $x(v) \in X$ and are accompanied by some witnesses $\theta(v)$, whose leaves $\ell \in V$ are labeled by $f_\ell \in \mathcal{F}$, such that the following holds: for any root-to-leaf path $v_1, \dots, v_d, v_{d+1} = \ell$ in the tree and for any $i \in [d]$:

$$\begin{aligned} f_\ell(x(v_i)) &\geq \theta(v_i) + \epsilon && \text{if } v_{i+1} \text{ is a left child of } v_i \\ f_\ell(x(v_i)) &\leq \theta(v_i) && \text{if } v_{i+1} \text{ is a right child of } v_i \end{aligned} \quad (8)$$

For the real-valued analogue of threshold dimension (ϵ -fat-threshold dimension), we slightly shift our use of the witness parameters. Rather than have a distinct witness θ_j for each x_j (as in the definition of ϵ -fat-shattering dimension) or for each v (as in the definition of ϵ -sequential-fat-shattering dimension), we use the same θ across. The reasoning behind this stems from our eventual use of representing game matrices as function classes, where it is preferred to have a definition that is symmetric to transposing the matrix, or equivalently, to swapping the roles of the concept class $\mathcal{F}$ and the domain set X . When we are considering bounded real-valued function classes (for example, only taking on values in $[0, 1]$), defining the ϵ -fat-threshold dimension in this way will only lead to a $O(1/\epsilon)$ factor difference in the dimension.

Definition A.10 (ϵ -fat-threshold dimension) For a real-valued function class $\mathcal{F}$, denote its ϵ -fat-threshold dimension $\text{fattr}(\mathcal{F}, \epsilon)$ to be the maximal d (possibly infinite) such that there exist magnitude- d subsets $\{f_1, \dots, f_d\} \subseteq \mathcal{F}$ and $\{x_1, \dots, x_d\} \subseteq X$ and a witness $\theta \in \mathbb{R}$ satisfying the following.

$$\begin{aligned} f_i(x_j) &\geq \theta + \epsilon && \text{for all } i \leq j \in [d] \\ f_i(x_j) &\leq \theta && \text{for all } i > j \in [d] \end{aligned} \quad (9)$$

For real-valued function classes bounded on $[0, 1]$, these dimensions are related in the following ways.

Lemma A.11 For any $[0, 1]$ -valued function class $\mathcal{F}$ and any $\epsilon > 0$,

$$\text{fat}(\mathcal{F}, \epsilon) \leq \min \left(\text{sfat}(\mathcal{F}, \epsilon), \frac{\text{fattr}(\mathcal{F}, \epsilon/2)}{2\epsilon} \right). \quad (10)$$

Further, there exist universal constants $c, C > 0$ such that:

$$\frac{c\epsilon \log(\epsilon \log(\text{sfat}(\mathcal{F}, \epsilon)))/C)}{\log(1/(C\epsilon))} \leq \text{fattr}(\mathcal{F}, \epsilon) \leq 2^{\text{sfat}(\mathcal{F}, \epsilon)+1} \quad (11)$$

We note that the left part of Eq. (11) was proved by [Daskalakis and Golowich, 2022, Lemma 8.4] and the remaining proofs are standard. We include sketches in Appendix G.

Lastly, we define the convex hull of a function class:

Definition A.12 The *convex hull* of a function-class $\mathcal{F}$ over domain X is defined as:

$$\text{conv}(\mathcal{F}) = \left\{ \sum_{j=1}^m \lambda_j f_j : m \in \mathbb{N}, f_1, \dots, f_m \in \mathcal{F}, \lambda_1, \dots, \lambda_m \geq 0, \sum_{i=1}^m \lambda_i = 1 \right\}$$

The *dual convex hull* of a class is defined as a class on the domain $\text{conv}(X)$ as the set of all formal convex combinations of finitely many elements from X , namely,

$$\text{conv}(X) = \left\{ \sum_{i=1}^\ell \lambda_i x_i : \ell \in \mathbb{N}, \lambda_1, \dots, \lambda_\ell \in [0, 1], \sum_{i=1}^\ell \lambda_i = 1 \right\}$$

where $\sum_i \lambda_i x_i$ is a formal sum. Extend each $f \in \mathcal{F}$ to the domain $\text{conv}(X)$ by defining the extended function $\text{dconv}(f)$ as

$$\text{dconv}(f) \left(\sum_i \lambda_i x_i \right) = \sum_i \lambda_i f(x_i).$$

Define the dual convex hull of $\mathcal{F}$ as

$$\text{dconv}(\mathcal{F}) = \{\text{dconv}(f) : f \in \mathcal{F}\}$$

Lastly, define $\text{conv2}(\mathcal{F}) = \text{dconv}(\text{conv}(\mathcal{F}))$.

A.1.3 Uniform convergence

Below, we use the notion of uniform convergence, which enables one to sample-down a distribution and compress it to a distribution over a small number of elements, while changing the expectation of each function by at most ϵ .

Definition A.13 (Uniform convergence) For a concept class $\mathcal{F}$ over a domain set X , define by $c(\mathcal{F}, \epsilon, \delta)$ the smallest number m , such that for any measure μ over X ,

$$\Pr_{x_1, \dots, x_m \sim \mu \text{ (i.i.d)}} \left[\forall f \in \mathcal{F}, \left| \frac{1}{m} \sum_{i=1}^m f(x_i) - \mathbb{E}_{x \sim \mu} [f(x)] \right| \leq \epsilon \right] \geq 1 - \delta.$$

Denote by $c(\mathcal{F}, \epsilon) = \inf_{\delta < 1} c(\mathcal{F}, \epsilon, \delta)$.

Intuitively, the support of any distribution can be compressed down to a size of at most $c(\mathcal{F}, \epsilon)$, while changing expectations of functions in $\mathcal{F}$ by at most ϵ .

We notice that $c(\mathcal{F}, \epsilon, \delta)$ can be bounded in terms of the VC dimension of 0-1 valued classes and in terms of the shattering numbers of real-valued classes: [Rudelson and Vershynin \[2006\]](#).

Lemma A.14 Let $\mathcal{F}$ be a concept class. If $\mathcal{F}$ is 0-1 valued then

$$c(\mathcal{F}, \epsilon, \delta) \leq C \frac{\text{VC}(\mathcal{F}) + \log(1/\delta)}{\epsilon^2},$$

where $C > 0$ is a universal constant. Similarly, if $\mathcal{F}$ is $[0, M]$ -valued then

$$c(\mathcal{F}, \epsilon, \delta) \leq C \frac{\left(\int_0^M \sqrt{\text{fat}(\mathcal{F}, \epsilon)} \right)^2 + \log(1/\delta)}{\epsilon^2}$$

A.2 Online learning and Multiplicative-Weights algorithm

We address the online learning setting in realizable and agnostic settings. Let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class, where $\mathcal{X}$ is the instance space and $\mathcal{Y} = \{0, 1\}$ is the label space. The online learning protocol can be formulated as a game between the learner and an adversary, where at rounds $t = 1, 2, \dots, T$,

1. The adversary chooses $(x_t, y_t) \in \mathcal{X} \times \{0, 1\}$.
2. The learner observes x_t and predicts $\hat{y}_t \in \{0, 1\}$.
3. The learner observes y_t and suffers a loss $|\hat{y}_t - y_t|$.

In the section, we assume realizability which means that all target labeled are generated by a function $h^* \in \mathcal{H}$, that is, $h^*(x_t) = y_t$ for $t \in [T]$. Define the mistake bound for a deterministic algorithm $\mathcal{A} : (\mathcal{X} \times \mathcal{Y})^* \times \mathcal{X} \rightarrow \{0, 1\}$ by

$$\mathcal{M}(\mathcal{H}, \mathcal{A}) = \sup_{h^* \in \mathcal{H}, T, x_{1:T}} \sum_{t=1}^T |\mathcal{A}((x_{1:(t-1)}, y_{1:(t-1)}), x_t) - h^*(x_t)|.$$

We can write the output of the algorithm as a function $\hat{c}_t(x) = \mathcal{A}((x_{1:(t-1)}, y_{1:(t-1)}), x)$. We say that the algorithm is *improper* if the output functions $\hat{c}_t$ do not belong to $\mathcal{H}$.

The learner is allowed to make randomized predictions, where the adversary picks (x_t, y_t) without knowing the random bits of the learner in this round. We analyze the expected loss of the learner. Formally, for a randomized algorithm $\mathcal{A} : (\mathcal{X} \times \mathcal{Y})^* \times \mathcal{X} \rightarrow \Delta(\{0, 1\})$, we define the expected loss of algorithm $\mathcal{A}$ in T rounds

$$\mathcal{L}(\mathcal{H}, \mathcal{A}, T) = \sup_{h^* \in \mathcal{H}, x_{1:T}} \sum_{t=1}^T \mathbb{E} |\mathcal{A}((x_{1:(t-1)}, y_{1:(t-1)}), x_t) - h^*(x_t)|,$$

where the expected loss at round t can be interpreted as the probability of predicting incorrectly at round t . If we write the algorithm's (random) output as a function $\hat{c}_t(x) = \mathcal{A}((x_{1:(t-1)}, y_{1:(t-1)}), x)$, we say that the algorithm is *randomized proper* for a function class $\mathcal{H}$ if the function $\hat{c}_t$ belong to $\mathcal{H}$. In other words, the algorithm draws a function from a distribution that is supported on $\mathcal{H}$.

In the agnostic setting, we define the more general setting of prediction with expert advice, where the loss function is arbitrary and not necessarily the 0-1 loss as in the realizable setting.

Definition A.15 (Prediction with expert advice) There are N experts indexed by $[N] = \{1, \dots, N\}$. In each time step $t = 1, \dots, T$ the learner chooses a probability vector $\mu^t = (\mu_1^t, \dots, \mu_N^t)$ from the simplex $\mathcal{S}_N = \{p \in \mathbb{R}^N : \forall i, p_i \geq 0 \text{ and } \sum_{i=1}^N p_i = 1\}$. Thereafter, a loss vector $\ell_t \in [0, 1]^N$ is revealed to the learner. In the adversarial setting, the loss vectors $\ell_1, \dots, \ell_T$ are entirely arbitrary and may be chosen by an adversary. The goal of the learner is to minimize the regret, given by

$$\mathcal{R}_T := \sum_{t=1}^T \mu^t \cdot \ell_t - \min_{i \in [N]} \sum_{t=1}^T \ell_t(i).$$

Lemma A.16 (Multiplicative weights algorithm) *Multiplicative Weights algorithm with learning rate $\eta = \sqrt{\frac{\log N}{2T}}$ suffers a regret of $O(\sqrt{T \log N})$ where N is the number of experts and T is the horizon length. Further, the time complexity is $O(N)$ per iteration.*

For proof, see [Littlestone \[1988\]](#), [Arora et al. \[2012\]](#), and [\[Schapire and Freund, 2013, Section 6\]](#)

Algorithm 4 Multiplicative Weights

Parameters: Learning rate parameter $\eta > 0$.

1. Initialize a uniform probability distribution $\mu^1 = (\mu_1^1, \dots, \mu_N^1)$ over the N experts, where $\mu_i^1 \leftarrow \frac{1}{N}$ for all $i = 1, \dots, N$.
2. For $t = 1, 2, \dots, T$:
 - (a) Predict μ^t and suffer a loss $\mu^t \cdot \ell_t$.
 - (b) Update μ^{t+1} based on ℓ_t :

$$\forall i, \mu_i^{t+1} \leftarrow \mu_i^t \exp(-\eta \ell_t(i)) / Z^t \text{ where } Z^t = \sum_{j=1}^N \mu_j^t \exp(-\eta \ell_t(j)).$$

A.3 Games

A.3.1 Games and equilibria

Definition A.17 (Multi-player game) A k -player game is a pair $(\mathcal{A}, u)$, where $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_k$ and $u = (u_1, \dots, u_k)$, where each $u_p : \mathcal{A} \rightarrow \mathbb{R}$. We assume that each $\mathcal{A}_i$ is accompanied with some Σ -algebra that u is measurable with respect to. Each $\mathcal{A}_p$ is the set of *actions* available to player p and each u_p is the *utility, or payoff, function* of player p , which maps the set of *action profiles* $\mathcal{A}$ to the reals. Each player's goal is to maximize their own utility. We denote by $\mathcal{A}_{-p}$ the Cartesian product of $\{\mathcal{A}_j\}_{j \neq p}$. Similarly, for any action $a = (a_1, \dots, a_k) \in \mathcal{A}$, denote by a_{-p} the Cartesian product of $\{a_j\}_{j \neq p}$.

Definition A.18 (Zero-sum game) A *zero-sum* game is a two-player game such that $u_1(a, b) = -u_2(a, b)$ for all $a \in \mathcal{A}$ and $b \in \mathcal{B}$. We sometimes compress our notation and represent a zero-sum game as $(\mathcal{A}, \mathcal{B}, u)$ where $u : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ is a single function representing the utility function of player 2. Player 2 aims to maximize this utility while player 1 aims to minimize this utility.

Definition A.19 (ϵ -Nash equilibrium and ϵ -CCE) Let $(\mathcal{A}, u)$ denote a game. An ϵ -approximate Nash equilibrium is a collection of probability measures, $\mu_1, \dots, \mu_k$, over $\mathcal{A}_1, \dots, \mathcal{A}_k$, respectively, such that for any player $p \in [k]$ and any $d_p \in \mathcal{A}_p$,

$$\mathbb{E}_{a \sim \mu_1 \times \dots \times \mu_k} [u_p(d_p, a_{-p})] \leq \mathbb{E}_{a \sim \mu_1 \times \dots \times \mu_k} [u_p(a)] - \epsilon.$$

A *Coarse Correlated Equilibrium* (CCE) is a joint measure μ over $\mathcal{A}$ such that for any player p and any $d_p \in \mathcal{A}_p$,

$$\mathbb{E}_{a \sim \mu} [u_p(d_p, a_{-p})] \leq \mathbb{E}_{a \sim \mu} [u_p(a)] - \epsilon.$$

Definition A.20 (Minimax equilibrium) Given a zero-sum game $G = (\mathcal{A}, \mathcal{B}, u)$ ⁷, we say that G has a *minimax equilibrium* if

$$\inf_{\mu_1 \in \Delta(\mathcal{A})} \sup_{\mu_2 \in \Delta(\mathcal{B})} \mathbb{E}_{a \sim \mu_1, b \sim \mu_2} [u(a, b)] = \sup_{\mu_2 \in \Delta(\mathcal{B})} \inf_{\mu_1 \in \Delta(\mathcal{A})} \mathbb{E}_{a \sim \mu_1, b \sim \mu_2} [u(a, b)] \quad (12)$$

where $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$ denote the set of all probability measures over $\mathcal{A}$ and $\mathcal{B}$ respectively. If Eq. (12) is satisfied, we say the probability measures μ_1, μ_2 optimizing Eq. (12) are a *minimax equilibrium* of G and denote by $\text{Val}(G)$ *value* of the game, which is the value of both sides of Eq. (12).

The following relation holds between an ϵ -Nash and the value of the game:

Lemma A.21 *Let $G(\mathcal{A}, \mathcal{B}, u)$ denote a zero-sum game with a minimax equilibrium. If (μ_1, μ_2) is an ϵ -approximate Nash equilibrium for this game, then*

$$\inf_{a \in \mathcal{A}} \mathbb{E}_{b \sim \mu_2} [u(a, b)] \geq \text{Val}(G) - \epsilon; \quad \sup_{b \in \mathcal{B}} \mathbb{E}_{a \sim \mu_1} [u(a, b)] \leq \text{Val}(G) + \epsilon.$$

We note that a distribution over actions of a particular player p is also termed a *mixed strategy*, and given mixed strategies $\mu_1, \dots, \mu_k$ for the players, we abuse notation and denote for any $p \in [k]$ the utility of player p given mixed strategies $\mu_1, \dots, \mu_k$ as $u_p[\mu_1, \dots, \mu_k] := \mathbb{E}_{a \sim \mu_1 \times \dots \times \mu_k} [u_p(a)]$.

A.3.2 Game dimensions

We now extend these dimensionality definitions from function classes to games. For each player p , her utility $u_p : \mathcal{A}_p \times \mathcal{A}_{-p} \rightarrow \mathbb{R}$ can be thought of as a concept class $\mathcal{F}_p$ over the domain set $X_p = \mathcal{A}_p$, whose concepts $f_{a_{-p}}$ are parameterized by elements $a_{-p} \in \mathcal{A}_{-p}$ and are defined by $f_{a_{-p}}(a_p) := u_p(a_p, a_{-p})$, for each $a_p \in \mathcal{A}_p$. We define the dimension of a game to be the maximal dimension over these utility function classes, where p ranges across all players.

Definition A.22 (Real-valued fat and threshold dimension for a multi-player game) Let $G = (\mathcal{A}, u)$ be a k -player game. We define

$$\begin{aligned} \text{fat}(G, \epsilon) &= \max_p \text{fat}(\mathcal{F}_p, \epsilon) \\ \text{fatr}(G, \epsilon) &= \max_p \text{fatr}(\mathcal{F}_p, \epsilon) \\ \text{sfat}(G, \epsilon) &= \max_p \text{sfat}(\mathcal{F}_p, \epsilon) \\ c(G, \epsilon, \delta) &= \max_p c(\mathcal{F}_p, \epsilon, \delta) \\ I(G) &= \left(\int_0^1 \sqrt{\text{fat}(G, \epsilon)} \right)^2 \end{aligned}$$

Lastly, the *convex hull* of a game

$$G = ((\mathcal{A}_1, \dots, \mathcal{A}_k), u)$$

as the game $\text{conv}(G) = ((\text{conv}(\mathcal{A}_1), \dots, \text{conv}(\mathcal{A}_k)), \tilde{u})$ where $\text{conv}(\mathcal{A}_p)$ is the set of all finitely-supported probability measures over $\mathcal{A}_p$ and $\tilde{u}_p(\mu_1, \dots, \mu_k) := \mathbb{E}_{a_p \stackrel{i.i.d.}{\sim} \mu_p} [u_p(a_1, \dots, a_k)]$.

A.3.3 Best-response oracle

Definition A.23 (Best-response oracle) Let $(\mathcal{A}, u)$ be a multi-player game. An ϵ -*best-response* oracle receives a bounded-support distribution μ over $\mathcal{A}_{-p}$ and outputs an action $\hat{a}_p \in \mathcal{A}_p$ that is an ϵ -best response, namely:

$$\mathbb{E}_{a_{-p} \sim \mu} u_p(\hat{a}_p, a_{-p}) \geq \sup_{a_p \in \mathcal{A}_p} \mathbb{E}_{a_{-p} \sim \mu} u_p(a_p, a_{-p}) - \epsilon.$$

⁷We slightly abuse notation and throughout denote a zero-sum game $(\mathcal{A} \times \mathcal{B}, u)$ as $(\mathcal{A}, \mathcal{B}, u)$

A.4 On learnability and the existence of game equilibria

We say that a concept class $\mathcal{F}$ is *uniformly online learnable* if there is a function $R: \mathbb{N} \rightarrow [0, \infty)$, that satisfies $R(T) = o(T)$ as $t \rightarrow \infty$, such that for any T there exists an online learner that achieves a regret of at most $R(T)$. We notice that a class is uniformly online learnable with regret of $o(T)$ if and only if its ϵ -sequential-fat-shattering dimension is finite for all $\epsilon > 0$. We notice that Algorithm 1 achieves a regret of $o(T)$ for any such class, using only a best response oracle, yet, with suboptimal regret.

Next, we discuss the reduction (see e.g., [Schapire and Freund, 2013, Section 6]) which shows that online learnability of an appropriate function-class implies the existence of an equilibria. First, we define the notion of a *repeated game*: this denotes the iterative setting, where in each iteration $i = 1, \dots, T$, each of the players is playing an action and gains a reward according to their utility. Equivalently, the negation of their reward can be viewed as a loss that they suffer. The next lemma states that if both players in a zero-sum game play a no-regret learning algorithm, then the average-iterate converges to a Nash equilibrium. Similarly, in a general sum game the average iterates converge to a CCE:

Lemma A.24 (Equilibria computation via a repeated game) *Assume a repeated game between two players in a zero-sum game, that is repeated for T iterations, where in each iteration t , player 1 plays an action a_1^t and player 2 plays a_2^t . Assume that each player p suffers a regret of ϵT , for $p \in \{1, 2\}$, namely:*

$$\sum_{t=1}^T u_p(a_p^t, a_{-p}^t) \geq \sup_{a_p \in \mathcal{A}_p} \sum_{t=1}^T u_p(a_p, a_{-p}^t) - \epsilon T.$$

Denote by μ_p the uniform distribution over $a_p^1, \dots, a_p^T$, for $p \in \{1, 2\}$. Then, (μ_1, μ_2) is an ϵ -Nash equilibrium.

Similarly, in a general-sum multi-player game, assume a repeated game such that every player $p \in [k]$ plays action a_p^t in iteration $t \in [T]$ and suffers a regret of at most ϵT . Denote by μ the joint distribution over $\mathcal{A}$ which is the uniform distribution over the multiset $\{(a_1^t, \dots, a_k^t)\}_{t \in [T]}$. Then, μ is an ϵ -CCE.

In particular, this implies that any game which admits no-regret learners has a minimax (for a zero-sum game) or a CCE (for general games). This implies that games with bounded ϵ -sequential-fat-shattering dimension for all $\epsilon > 0$ admit such equilibria. Equivalently, games with bounded ϵ -fat-threshold dimension $\epsilon > 0$ attain such equilibria.

The converse is not completely true, yet, for zero-sum games, there are known lower bounds that use similar notions of a dimension. For example, in a 0-1 valued game, assume that the game contains a subgame which is an infinitely large threshold game, namely, there exist actions $\{a_1^n\}_{n \in \mathbb{N}}$ and $\{a_2^n\}_{n \in \mathbb{N}}$ such that $u(a_1^n, a_2^m) = 1$ if $m \geq n$ and 0 otherwise. Then, as observed by Hanneke et al. [2021], this subgame does not contain a Nash Equilibrium. There is a gap between the above described upper and lower bound, which is the setting where the threshold dimension is infinite, namely, there are arbitrarily large threshold games, yet, there is no infinitely large threshold subgame. In this particular setting, Hanneke et al. [2021] showed that if additional the VC dimension is finite then the game admits a Nash Equilibrium.

B Deferred proofs for Online Learning

Lemma B.1 *Let $0 < \epsilon \leq 1/2$, let $d \geq 1$, let $\alpha = \epsilon/(d+1)$ and let J such that*

$$1 \leq J \leq \left(\frac{1}{\epsilon - \alpha} \right)^d.$$

Then,

$$\frac{J \log J}{\alpha^2} \leq C d^3 \log(1/\epsilon) \left(\frac{1}{\epsilon} \right)^d$$

where $C > 0$ is a universal constant.

Proof First,

$$\frac{1}{\epsilon - \alpha} = \frac{1}{\epsilon} \frac{1}{1 - 1/(d+1)} = \frac{1}{\epsilon} \left(1 + \frac{1}{d} \right).$$

Consequently,

$$\begin{aligned} \frac{J \log J}{\alpha^2} &\leq \frac{d \log(1/(\epsilon - \alpha))}{\alpha^2} \left(\frac{1}{\epsilon - \alpha} \right)^d = d^3 (\log(1/\epsilon) + \log(1 + 1/d)) \left(\frac{1}{\epsilon} \right)^d \left(1 + \frac{1}{d} \right)^d \\ &\leq C d^3 \log(1/\epsilon) \left(\frac{1}{\epsilon} \right)^d. \end{aligned}$$

■

Lemma B.2 *The improper variant of Algorithm 1 makes at most $C^{\text{Lit}(\mathcal{F})}$ mistakes, for a universal constant $C > 0$.*

Proof Substitute $\epsilon = 1/2$ and $\alpha = 1/4$. It has argued, in the main proof body, that $\text{fattr}(\text{dconv}(\mathcal{H} \oplus h^*), \epsilon - \alpha) \geq J$, under the assumption that the algorithm runs for more than J phases. By Theorem 2.3,

$$\begin{aligned} J &\leq \text{dconv}(\mathcal{H} \oplus h^*, \epsilon - \alpha) = e^{C \text{sfat}(\mathcal{H} \oplus h^*, (\epsilon - \alpha)/C) / (\epsilon - \alpha)^2} = e^{C \text{Lit}(\mathcal{H} \oplus h^*) / (\epsilon - \alpha)^2} \\ &= e^{C' \text{Lit}(\mathcal{H} \oplus h^*)} \end{aligned}$$

for universal constants $C, C' > 0$. It is well known and follows from definition that $\text{Lit}(\mathcal{H}) = \text{Lit}(\mathcal{H} \oplus h^*)$, which yields that

$$J \leq e^{C' \text{Lit}(\mathcal{H})}.$$

Since the number of times the distribution μ^t of the algorithm makes an ϵ -mistakes in each round j is bounded by $\mathcal{O}(\log(j)/\alpha^2) \leq \mathcal{O}(\log(J))$ (recall $\alpha = 1/4$), the total number of times of such an ϵ -mistake across all phases is bounded, up to constants, by

$$J \log J \leq C \text{Lit}(\mathcal{H}) e^{C' \text{Lit}(\mathcal{H})},$$

where C is a universal constant. Recall that in the improper setting, the algorithm predicts according to the majority. Consequently, it makes a mistake if μ^t makes an $\epsilon = 1/2$ mistake. Hence, the number of mistakes of the algorithm is bounded by $\mathcal{O}(e^{C' \text{Lit}(\mathcal{H})})$. ■

B.1 Agnostic Online Learning

Proof Proof of Theorem 3.5 We start with the bound for improper learners. We use the reduction of Ben-David et al. [2009] to reduce from agnostic to realizable. In their proof, they instantiate $\binom{T}{M}$ algorithms for the realizable setting, where M is the mistake bound of the realizable algorithm (i.e. the regret bound). Each of these instantiations is fed with the original sequence of $x_1, \dots, x_T$, however, the labels $y_1, \dots, y_T$ are different in each instantiation. A multiplicative weights algorithm (Algorithm 4) is used to choose between these experts. In their paper, they proved that the regret of this algorithm is bounded by $\mathcal{O}(\sqrt{TM \log T})$ ⁸, which translates to the bound of Theorem 3.5, when one substitutes M with mistake bound (i.e. the regret bound) of the improper learner from Theorem 3.3. In the randomized proper setting, instead of a bound on the number of mistakes, the proof of Theorem 3.3 yields a bound $M(\epsilon)$ on the number of ϵ -mistakes. While translating it into the reduction of Ben-David et al. [2009], one obtains a total regret of $\mathcal{O}(\sqrt{TM(\epsilon) \log T} + T\epsilon)$, where $T\epsilon$ accounts for the additional loss caused by the fact that each of the experts possibly suffers an additional regret of $T\epsilon$, accounting for less-than- ϵ -mistakes. Substituting the bound of $M(\epsilon) \leq \tilde{\mathcal{O}}(\epsilon^{-2 \text{tr}(\mathcal{H}) - 2})$ from Theorem 3.3 and substituting $\epsilon = T^{-1/(2 \text{tr}(\mathcal{H}) + 4)}$ yields a total regret of $T^{\frac{2 \text{tr}(\mathcal{H}) + 3}{2 \text{tr}(\mathcal{H}) + 4}}$.

Number of oracle calls and Runtime: Since the algorithm selects from $\binom{T}{M}$ using multiplicative weight update, the runtime is linear in the number of experts and equals $\mathcal{O}\left(\binom{T}{M}\right)$. In the improper setting, $M = \tilde{\mathcal{O}}(4^{\text{tr}(\mathcal{H})})$, which translates to a bound of $\binom{T}{M} \leq T^{\tilde{\mathcal{O}}(4^{\text{tr}(\mathcal{H})})}$ and in the proper setting, one a bound on $M(\epsilon) \leq \tilde{\mathcal{O}}(\epsilon^{-2 \text{tr}(\mathcal{H}) - 2})$, which yields a regret of $T^{\tilde{\mathcal{O}}(\epsilon^{-2 \text{tr}(\mathcal{H}) - 2})}$. For the specific choice of $\epsilon = T^{-1/(2 \text{tr}(\mathcal{H}) + 4)}$, the runtime per iteration is at most $T^{\tilde{\mathcal{O}}(T^{(2 \text{tr}(\mathcal{H}) + 3)/(2 \text{tr}(\mathcal{H}) + 4)})}$.

⁸In their paper, Ben-David et al. [2009] provided an algorithm for the agnostic setting with a mistake bound of $\mathcal{O}(\sqrt{T \text{Lit}(\mathcal{H}) \log T})$, however, the only property they used of the Littlestone's dimension is the existence of an algorithm with a mistake bound $\text{Lit}(\mathcal{H})$ for the realizable setting. Hence, we can replace $\text{Lit}(\mathcal{H})$ with the mistake bound of any algorithm and the proof would follow.

The number of oracle calls can be trivially bounded by the runtime. Though, one can obtain a better bound. There are at most t^d distinct functions if we restrict our hypothesis class to a domain of size t . In particular, there are at most t^d maximal realizable partial labelings of $\{x_1, \dots, x_t\}$. It suffices to call the consistent oracle only on these maximal labelings (and the collection maximal realizable labelings of $x_1, \dots, x_{t+1}$ can be efficiently constructed from those over $x_1, \dots, x_t$). This yields at bound of at most T^d oracle calls per iteration. ■

C Minmax in zero-sum games

This section is dedicated to the proof of Theorem 4.2. We denote the game as $G = (\mathcal{A}, \mathcal{B}, u)$, where $\mathcal{A}$ and $\mathcal{B}$ are the sets of actions of Players 1 and 2, respectively. In Section C.1 we will describe an algorithm that solves the minimax of a game where one player has infinitely many actions at hand, and she is equipped with a best-response oracle, and the other player has finitely many actions, to which he has random access. Next, in Section C.2 we describe an algorithm, that iteratively uses the algorithm for the half-infinite game to compute a Nash equilibrium for a game where both players have infinitely many actions.

C.1 Nash for a half-infinite game

In this section, we will describe an algorithm to find a Nash equilibrium for a game where one of the players has a finite set of actions and the other player has an infinite set of actions. Assume without loss of generality that player 1 has the finite action set. A similar approach to computing an equilibria in half-infinite games appeared in the context of robust PAC learning [Feige et al., 2015, Attias et al., 2022].

In order to compute the Nash equilibrium, recall from Lemma A.24 that if the players play a repeated game and if both players play a no-regret learning algorithm, then their average-iterate converge to a Nash equilibrium. In particular, multiplicative weights (Algorithm A.2) and ϵ -best response are both no-regret learning algorithms, therefore, the following lemma is a (well known) consequence of Lemma A.16:

Lemma C.1 (Multiplicative weights vs. best response) *Assume a repeated zero-sum game for T iterations, between player 1 who plays over n actions and player 2 whose number of actions is unbounded and possibly infinite. Assume that Player 1 plays the exponential weights algorithm to choose a mixed strategy μ^t over their actions in each iteration $t \in [T]$, and assume that player 2 reacts with an ϵ -best-response b_t to the uniform mixture of $\mu^1, \dots, \mu^t$, namely,*

$$\frac{1}{t} \sum_{i=1}^t u(\mu^i, b_t) \geq \sup_{b \in \mathcal{B}} \frac{1}{t} \sum_{i=1}^t u(\mu^i, b) - \epsilon.$$

Denote by μ the uniform mixture of $\mu^1, \dots, \mu^T$ and by ξ the uniform distribution over $b_1, \dots, b_T$. Then, (μ, ξ) is an $O\left(\sqrt{\log(n)/T} + \epsilon\right)$ -Nash equilibrium.

Following Lemma C.1, we propose Algorithm 5, where Player 1 plays exponential weights over her actions and Player 2 plays a best-response, using his best-response oracle.

Algorithm 5 ϵ -Nash Equilibrium for a half-infinite zero-sum game

Input: A game $(\mathcal{A} = \{a_1, \dots, a_n\}, \mathcal{B}, u)$, a value $\epsilon > 0$.

Subroutines:

- **BestResponse** oracle: receives a mixed strategy over actions from $\mathcal{A}$ and an $\epsilon > 0$ and outputs an ϵ -best response from $\mathcal{B}$ (see Definition A.23).
- The utility function $u(a, b)$ that receives a pair of actions and outputs its utility.

1. $T \leftarrow \left\lceil \frac{C \log |\mathcal{A}|}{\epsilon^2} \right\rceil$; $\eta \leftarrow \sqrt{\frac{\log |\mathcal{A}|}{2T}}$.

2. $\mu^1 = (\mu_1^1, \dots, \mu_n^1)$ denotes a uniform probability distribution over $\mathcal{A}$ where $\mu_i^1 \leftarrow \frac{1}{n}$ for all $i = 1, \dots, n$.

3. $b_1 \leftarrow \text{BestResponse}(\mu^1, \epsilon)$.

4. **For** $t = 2, \dots, T$

(a) **For** $i = 1, \dots, n$,

$$\mu_i^t \leftarrow \mu_i^{t-1} \exp(\eta u(a_{t-1}, b_i)) / Z^t \text{ where } Z^t = \sum_{j=1}^n \mu_j^{t-1} \exp(\eta u(a_{t-1}, b_j))$$

(b) $b_t \leftarrow \text{BestResponse}(\mu^t, \epsilon)$

5. **Return** $(\bar{\mu}, \bar{\xi})$, where $\bar{\mu}$ is the uniform mixture of $\mu^1, \dots, \mu^T$, namely, $\bar{\mu}_i = \frac{1}{T} \sum_{t=1}^T \mu_i^t$ and $\bar{\xi}$ is the uniform distribution over $b_1, \dots, b_T$

We obtain the following statement:

Lemma C.2 (Nash for the half-infinite game) *Let $G = (\mathcal{A} = \{a_1, \dots, a_n\}, \mathcal{B}, u)$ be a zero-sum game where $|\mathcal{A}| = n$ and $\mathcal{B}$ is possibly infinite, and let $\epsilon > 0$. Then, Algorithm 5, executed with the parameter ϵ , finds an $O(\epsilon)$ -Nash equilibrium, after $T = O(\log n / \epsilon^2)$ iterations.*

Proof Algorithm 5 implements exponential-weights vs best response, therefore, the proof follows directly from Lemma C.1. ■

C.2 The algorithm for the fully-infinite game

We first argue that the output of the algorithm is an $O(\epsilon)$ -approximate Nash equilibrium.

Lemma C.3 *Assume that Algorithm 3 stops. Then, the returned strategies constitute a 5ϵ -Nash for the original game $(\mathcal{A}, \mathcal{B}, u)$.*

Proof Since the game ends, then either $\text{Val}(A_t, B_{t-1}) \geq \text{Val}(A_{t-1}, B_{t-1}) - \epsilon$ or $\text{Val}(A_t, B_t) \leq \text{Val}(A_t, B_{t-1}) + \epsilon$. Assume first that the latter holds.

Notice that for any t ,

$$\text{Val}(A_t, B_{t-1}) \leq \text{Val}(\mathcal{A}, B_{t-1}) + \epsilon. \quad (13)$$

This is due to the fact that in the game (A_t, B_{t-1}, u) , player 1 has a strategy that guarantees her a value of at least $\text{Val}(\mathcal{A}, B_{t-1})$: indeed, her strategy $\xi^{t,1}$ satisfies this property, because $\xi^{t,1}$ is a strategy for player 1 in the game $(\mathcal{A}, B_{t-1}, \epsilon)$. Similarly,

$$\text{Val}(A_t, B_t) \geq \text{Val}(A_t, \mathcal{B}) - \epsilon. \quad (14)$$

Recall that we assumed that $\text{Val}(A_t, B_t) \leq \text{Val}(A_t, B_{t-1}) + \epsilon$. This with the equations above yields:

$$\text{Val}(A_t, \mathcal{B}) - \epsilon \leq \text{Val}(A_t, B_t) \leq \text{Val}(A_t, B_{t-1}) + \epsilon \leq \text{Val}(\mathcal{A}, B_{t-1}) + 2\epsilon.$$

Recall that $\mu^{t,2}$ is the strategy for player 2 in an ϵ -Nash in the game $(\mathcal{A}, B_{t-1}, u)$. This implies that for any $a \in \mathcal{A}$,

$$u(a, \mu^{t,2}) \geq \text{Val}(\mathcal{A}, B_{t-1}) - \epsilon \geq \text{Val}(A_t, B_t) - 3\epsilon. \quad (15)$$

Similarly, since $\xi^{t,1}$ is the strategy for player 1 in an ϵ -Nash for the game $(A_t, \mathcal{B}, u)$, then for every $b \in \mathcal{B}$:

$$u(\xi^{t,1}, b) \leq \text{Val}(A_t, \mathcal{B}) + \epsilon \leq \text{Val}(A_t, B_t) + 2\epsilon.$$

Combining the equations above, we obtain that

$$\text{Val}(A_t, B_t) - 3\epsilon \leq u(\xi^{t,1}, \mu^{t,2}) \leq \text{Val}(A_t, B_t) + 2\epsilon.$$

Consequently,

$$u(\xi^{t,1}, \mu^{t,2}) - \inf_{a \in \mathcal{A}} u(a, \mu^{t,2}) \leq \text{Val}(A_t, B_t) + 2\epsilon - (\text{Val}(A_t, B_t) - 3\epsilon) = 5\epsilon,$$

and similarly,

$$\sup_{b \in \mathcal{B}} u(\xi^{t,1}, b) - u(\xi^{t,1}, \mu^{t,2}) \leq \text{Val}(A_t, B_t) + 2\epsilon - (\text{Val}(A_t, B_t) - 3\epsilon) = 5\epsilon.$$

This concludes that $(\xi^{t,1}, \mu^{t,2})$ is a 5ϵ -Nash, and recall that we assumed that $\text{Val}(A_t, B_t) \leq \text{Val}(A_t, B_{t-1}) + \epsilon$. This was one of the stopping conditions. However, we have to consider the second stopping condition, namely, that $\text{Val}(A_t, B_t) \geq \text{Val}(A_{t-1}, B_t) - \epsilon$. Recall that Eq. (13) states that

$$\text{Val}(A_t, B_{t-1}) \leq \text{Val}(\mathcal{A}, B_{t-1}) + \epsilon$$

and we substitute $t - 1$ instead of t in Eq. (14) to obtain that

$$\text{Val}(A_{t-1}, B_{t-1}) \geq \text{Val}(A_{t-1}, \mathcal{B}) - \epsilon.$$

Combining with our assumption that

$$\text{Val}(A_t, B_{t-1}) \geq \text{Val}(A_{t-1}, B_{t-1}) - \epsilon,$$

we obtain that

$$\text{Val}(A_{t-1}, \mathcal{B}) - 2\epsilon \leq \text{Val}(A_{t-1}, B_{t-1}) - \epsilon \leq \text{Val}(A_t, B_{t-1}) \leq \text{Val}(\mathcal{A}, B_{t-1}) + \epsilon.$$

Since $\mu^{t,2}$ is the strategy of player 2 in an ϵ -Nash for the game $(\mathcal{A}, B_{t-1}, u)$, we obtain that for any $a \in \mathcal{A}$,

$$u(a, \mu^{t,2}) \geq \text{Val}(\mathcal{A}, B_{t-1}) - \epsilon \geq \text{Val}(A_{t-1}, B_{t-1}) - 3\epsilon.$$

Further, since $\xi^{t,1}$ is the strategy of player 1 in the ϵ -Nash for the game $(A_t, \mathcal{B}, u)$ and since the Nash equilibrium is monotone under addition of actions to a single player, for any $b \in \mathcal{B}$,

$$u(\xi^{t,1}, b) \leq \text{Val}(A_t, \mathcal{B}) + \epsilon \leq \text{Val}(A_{t-1}, \mathcal{B}) + \epsilon \leq \text{Val}(A_{t-1}, B_{t-1}) + 2\epsilon.$$

The proof concludes similarly to how it ended in the first of the two cases that we analyze. ■

Next, we are bounding the stopping time of the algorithm.

Lemma C.4 *Algorithm 3 stops after $O(\text{fattn}(\text{conv}(G), \epsilon)/\epsilon)$ iterations.*

Proof of lemma C.4 Suppose the game runs for more than T iterations. First, we would like to find indices $1 \leq i_1 < i_2 < \dots < i_q < T$ for which the following hold:

$$\begin{aligned} \text{Val}(A_{i_j}, B_{i_k}) &\leq \theta & \text{if } j > k \\ \text{Val}(A_{i_j}, B_{i_k}) &\geq \theta + \epsilon/2 & \text{if } j \leq k. \end{aligned} \tag{16}$$

We claim that we can always guarantee $q = O(T\epsilon)$ such indices.

Define a 'crossing' of an interval $[a, b]$ where $0 \leq a < b \leq 1$ to be a pair of numbers c, d such that $c \leq a < b \leq d$. We will prove that we can find q pairs of the form $(\text{Val}(A_{i_j}, B_{i_{j-1}}), \text{Val}(A_{i_j}, B_{i_j}))$ that cross the same interval. Indeed, if we take the numbers $\text{Val}(A_{i_j}, B_{i_{j-1}})$, for $1 \leq i_j \leq T$, by Lemma G.1 we have that at least $\epsilon T/2$ of them lying in the interval $U = [\theta - \frac{\epsilon}{2}, \theta]$, for some θ , and denote them by $i_1 < \dots < i_q$. Since the algorithm has not stopped at any of these iterations, by the stopping condition of the algorithm it holds that $\text{Val}(A_{i_j}, B_{i_j}) \geq \text{Val}(A_{i_j}, B_{i_{j-1}}) + \epsilon \geq \theta + \epsilon/2$. This concludes that $(\text{Val}(A_{i_j}, B_{i_{j-1}}), \text{Val}(A_{i_j}, B_{i_j}))$ crosses the interval $[\theta, \theta + \epsilon/2]$. This concludes Eq. (16): indeed, for all $j > k$, since $i_j > i_k$, consequently, $i_k \leq i_j - 1$

then, by the definition of the algorithm, $B_{i_k} \subseteq B_{i_{j-1}}$, which implies, by the monotonicity of the value of the game, that

$$\text{Val}(A_{i_j}, B_{i_k}) \leq \text{Val}(A_{i_j}, B_{i_{j-1}}) \leq \theta.$$

Similarly, one can deduce that $\text{Val}(A_{i_j}, B_{i_k}) \geq \theta + \epsilon/2$ for $j \leq k$, which concludes Eq. (16).

Recall that taking the average of $c(G, \frac{\epsilon}{12}, \delta)$ (defined in Definition A.22) actions can guarantee us an $\frac{\epsilon}{12}$ -approximate minmax strategy α_{i_j} for player 1 in the game $(A_{i_j}, B_{i_{j-1}}, u)$, due to uniform convergence (A.13). Similarly we can get a $\frac{\epsilon}{12}$ -approximate minmax strategy β_{i_j} for player 2 in the game (A_{i_j}, B_{i_j}, u) . For these strategies, notice that for all $j > k$, it holds that

$$u(A_{i_j}, B_{i_k}) \leq \theta + \frac{\epsilon}{6}$$

and for all $j \leq k$, it holds that

$$u(A_{i_j}, B_{i_k}) \geq \theta + \frac{\epsilon}{3}.$$

Now, look at the matrix where the rows are parameterized by the strategies of player 1, $\{\alpha_{i_j} : 1 \leq j \leq q\}$, the columns by the strategies of player 2, $\{\beta_{i_k} : 1 \leq k \leq q\}$, and the value of the (j, k) entry is taken to be $u(\alpha_{i_j}, \beta_{i_k})$. This constitutes a matrix where all entries above and at the diagonal are at least $\theta + \frac{\epsilon}{3}$ and below the diagonal at most $\theta + \frac{\epsilon}{6}$. This concludes that the $\frac{\epsilon}{6}$ -fat threshold dimension of the mixed-strategy game is at least $\Omega(T\epsilon)$, namely $\text{fatrr}(\text{conv}(G), \epsilon) \geq \Omega(T\epsilon)$ which gives $T \leq O(\frac{\text{fatrr}(\text{conv}(G), \epsilon)}{\epsilon})$, as required. ■

In order to bound the number of oracle calls, we add the following lemma:

Lemma C.5 *Assume that Algorithm 3 runs for T iterations. Then, the number of oracle calls is bounded by $O(T/\epsilon^2 \cdot \log(T/\epsilon^2))$.*

Proof First, we would like to bound the sizes of A_t and B_t by $O(t \log t/\epsilon^2)$. In order to show that, notice that A_t is obtained from A_{t-1} by adding the support of an ϵ -approximate Nash for the half-infinite game $(\mathcal{A}, B_{t-1}, u)$. We would like to bound the size of the support of the strategy of Player 1 in such an approximate Nash. This approximate Nash is computed in Algorithm 5, and the size of the support equals the number of iterations of this algorithm, which is bounded by $O(\log |B_{t-1}|/\epsilon^2)$, by Lemma C.2. This implies that $|A_t| \leq |A_{t-1}| + C(\log |B_{t-1}| + C)/\epsilon^2$ for a universal constant $C > 0$ and similarly, since B_t is obtained from B_{t-1} by adding the support of an approximate Nash for the game $(A_t, \mathcal{B}, u)$, it holds that $|B_t| \leq |B_{t-1}| + C(\log |A_t| + C)/\epsilon^2$. It can be proven by induction that $|A_t|, |B_t| \leq O((t/\epsilon^2) \log(t/\epsilon^2))$. We notice that the number of oracle calls by iteration t equals exactly $|A_t| + |B_t|$, which concludes the proof. ■

We are now ready to prove our main theorem.

Theorem 4.2 *Let $G = (\mathcal{A}, \mathcal{B}, u)$ be a zero-sum two-player game, where $u : \mathcal{A} \times \mathcal{B} \rightarrow [0, 1]$ and let $\epsilon > 0$. There is an algorithm to find an $O(\epsilon)$ -Nash for this game using the following number of ϵ -best response oracle calls (assuming this number is finite):*

$$O\left(\min\left(e^{C \text{sfat}(G, \epsilon/C)/\epsilon^2}, (1/\epsilon)^{CI(G)^2 \text{fatrr}(G, \epsilon/C)/\epsilon^5}\right)\right); I(G) = \left(\int_0^1 \sqrt{\text{fat}(G, \epsilon)}\right)^2$$

Proof Proof of Theorem 4.2 We notice that Lemma C.3 implies that the output of Algorithm 3 is an $O(\epsilon)$ -Nash. Furthermore, Lemma C.4 bounds the number of iterations by $T \leq O(\text{fatrr}(\text{conv}(G), \epsilon)/\epsilon)$ and Lemma C.5 implies that the number of oracle calls is bounded by $\tilde{O}(\text{fatrr}(\text{conv}(G), \epsilon)/\epsilon^3)$. The proof of Theorem 4.2 follows by substituting $\text{fatrr}(\text{conv}(G), \epsilon)$ with its bound in terms of the various dimensions of the game G , according to Theorem 2.3:

$$\text{fatrr}(\text{conv}(G), \epsilon) \leq e^{C \text{sfat}(G, \epsilon/C)/\epsilon^2}.$$

The total bound on the number of oracle calls is then

$$\tilde{O}(\text{fatrr}(\text{conv}(G), \epsilon)/\epsilon^3) \leq e^{C \text{sfat}(G, \epsilon/C)/\epsilon^2}/\epsilon^3.$$

Notice that the fact of $1/\epsilon^3$ can be omitted by changing the constant in the exponent, and this yields the desired bound on the number of oracle calls and concludes Theorem 4.2. ■

D Bounding the threshold dimension of the mixed-strategy game

In this Sections [D.1](#) and [D.2](#) we prove Theorem [2.3](#), which consists of two upper bounds on $\text{fattr}(\text{conv}(\mathcal{F}), \epsilon)$ and $\text{fattr}(\text{conv}(\mathcal{G}), \epsilon)$ for a zero-sum game G and a concept class $\mathcal{F}$. We further present lower bounds in Section [D.3](#).

D.1 A bound using Ramsey's theory

In this section, we will prove the following lemma:

Lemma D.1 *Let G be a zero-sum game, then,*

$$\text{fattr}(\text{conv}(G), \epsilon) \leq O \left(\left(\frac{2}{\epsilon} \right)^{c(G, \epsilon/4)^2 \text{fattr}(G, \epsilon/4)/\epsilon} \right).$$

We notice that substituting $c(G, \epsilon/4)$ with $\left(\int_0^1 \sqrt{\text{fat}(G, \epsilon)} d\epsilon \right)^2 / \epsilon^2$ using Lemma [A.14](#), yields one of the two bounds of Theorem [2.3](#).

In the proof, we will use the following variants of Ramsey numbers, as defined below:

Definition D.2 (Multi-colored Ramsey number) • $R(n, Q)$ is defined as the maximal size of a complete graph that contains no monochromatic n -clique, where each edge can be colored by one of Q colors.

- $R(n, Q, \ell)$ is defined as the maximal size of a complete graph that contains no n -monochromatic clique, if each edge is colored by ℓ colors out of Q possible colors.
- $R\left((n_1, n_2, \dots, n_Q), Q, \ell\right)$ is defined as the maximal size of a complete graph that contains no n_i -monochromatic clique of color i , if each edge is colored by ℓ colors out of Q possible colors. Note that $R\left((n, n, \dots, n), Q, \ell\right) = R(n, Q, \ell)$

The following bound holds:

Proposition D.3 (A bound on $R(m, Q)$, from [Balaji et al. \[2021\] Corollary 3.9](#)) *For $n, Q \in \mathbb{N}$, $R(n, Q)$ can be bounded as follows*

$$R(n, Q) \leq \frac{3 + e}{2} \frac{(Q(n-2))!}{((n-2)!)^Q}$$

Using the bound $m! = O(m^m)$, we can write the above as:

$$R(n, Q) \leq O(Q^{Q^n})$$

Here, we show how to improve the bound if we know that each edge is colored with ℓ colors:

Proposition D.4 (A bound on $R(n, Q, \ell)$) *For $n, Q, \ell \in \mathbb{N}, \ell < Q$ it holds that:*

$$R\left(n, Q, \ell\right) \leq 2 \left(\frac{Q}{\ell} \right)^{(n-2)Q}$$

An immediate corollary from the above proposition is the following:

Corollary D.5 (A bound on $R(n, Q, \epsilon Q)$) *We have that for $n, Q \in \mathbb{N}$ and $\epsilon \in (0, 1)$*

$$R(n, Q, \epsilon Q) \leq O \left(\left(\frac{1}{\epsilon} \right)^{Q^n} \right)$$

Now to prove Proposition [D.4](#), we use the following:

Proposition D.6 (Proposition 2.1, [Balaji et al. \[2021\]](#)) *Suppose $\ell, Q, n_1, \dots, n_r \in \mathbb{N}$, and suppose $A_1, \dots, A_Q$ are sets with $|A_1 \cup \dots \cup A_Q| = n$. If $n > n_1 + \dots + n_Q - Q$, then $|A_i| \geq n_i$ for some $1 \leq i \leq Q$*

Lemma D.7 *For some $Q, \ell, n_1, n_2, \dots, n_Q \in \mathbb{N}$, we have that:*

$$\begin{aligned} & \ell \left(R\left((n_1, n_2, \dots, n_r), Q, \ell\right) - 1 \right) \\ & \leq R\left((n_1 - 1, n_2, \dots, n_Q), Q, \ell\right) + \dots + R\left((n_1, n_2, \dots, n_Q - 1), Q, \ell\right) - (Q - 1) \end{aligned}$$

Proof of Lemma D.7 Assume $N = R((n_1, n_2, \dots, n_r), Q, \ell)$. Let v be an arbitrary node out of the N nodes of the K_N graph and denote it as the *parent node*. Allocate the $N - 1$ remaining vertices into r sets, namely $A_1, A_2, \dots, A_Q$, where each node can appear in multiple sets. A node u will be assigned to set A_c , where $c = 1, 2, \dots, Q$ if one of the colors of the edge that connects u and v is c . Now, note that

$$|A_1 \cup A_2 \cup \dots \cup A_Q| = (N - 1)\ell$$

since each edge has ℓ colors. If we assume towards contradiction that:

$$(N - 1)\ell > R((n_1 - 1, n_2, \dots, n_r), Q, \ell) + \dots + R((n_1, n_2, \dots, n_r - 1), Q, \ell) - (Q - 1)$$

then there exists an i such that $|A_i| \geq R((n_1, \dots, n_i - 1, \dots, n_r), Q, \ell)$. If such an i exists, that means that we either have a K_{n_j} clique for $i \neq j$ in A_i , or we have a $K_{n_i - 1}$ clique in A_i , and therefore we can create a K_{n_i} clique by connecting the parent node v with the $n_i - 1$ nodes in A_i that form that clique. That means that we have a color i for which we have a monochromatic clique of size n_i , contradicting the definition of the multi-colored Ramsey number $R((n_1, n_2, \dots, n_r), Q, \ell)$. Thus it must be that

$$(N - 1)\ell > R((n_1 - 1, n_2, \dots, n_r), Q, \ell) + \dots + R((n_1, n_2, \dots, n_r - 1), Q, \ell) - (Q - 1)$$

as required. ■

Proof of Proposition D.4 Let us define the following quantity

$$S(N, Q, \ell) = \max_{(n_1, n_2, \dots, n_Q) | n_i \geq 2, \sum_i n_i = N} R((n_1, \dots, n_Q), Q, \ell)$$

Note that $S(2Q, Q, \ell) = 2$, as we have that $\sum_i n_i = 2Q$ which means either $n_i = 2 \forall i$ in which case the multi-colored Ramsey number is 2, or there exists i such that $n_i = 1$ in which case the number is 1. We can use Lemma D.7 as follows:

$$\begin{aligned} S(N, Q, \ell) &= R((n_1, \dots, n_Q), Q, \ell) \\ &\leq \frac{R((n_1 - 1, \dots, n_Q), Q, \ell) + \dots + R((n_1, \dots, n_Q - 1), Q, \ell)}{\ell} + 1 - \frac{Q - 1}{\ell} \\ &\leq \frac{R((n_1 - 1, \dots, n_Q), Q, \ell) + \dots + R((n_1, \dots, n_Q - 1), Q, \ell)}{\ell} \leq \frac{Q}{\ell} S(N - 1, Q, \ell) \end{aligned}$$

Using the above we can conclude that:

$$S(N, Q, \ell) \leq \left(\frac{Q}{\ell}\right)^{N-2Q} S(2Q, Q, \ell) = 2 \left(\frac{Q}{\ell}\right)^{N-2Q}$$

Thus for the Multi colored Ramsey number $R((n_1, \dots, n_Q), Q, \ell)$ we have

$$R(n, Q, \ell) \leq S(nQ, Q, \ell) \leq 2 \left(\frac{Q}{\ell}\right)^{(n-2)Q}$$

■

We proceed with the proof of Lemma D.1 as well as Corollary D.8 for the binary setting. Denote $m = \text{fatr}(\text{conv}(G), \epsilon)$ and $n = \text{fatr}(G, \epsilon/2)$. Since the ϵ -fat threshold dimension of the mixed-strategy game is m , that means that there exist strategies $\alpha_i, \beta_i, i \in [m]$ and a threshold $\theta \in [0, 1]$ for which $u(\alpha_i, \beta_j) \leq \theta$ if $i > j$ and $u(\alpha_i, \beta_j) \geq \theta + \epsilon$ if $i \leq j$. These strategies' supports may be unbounded. Towards applying the bound on the multi-colored Ramsey numbers, we would like to replace these strategies with strategies that have the above properties, with perhaps a gap smaller than ϵ , however with bounded support. We use the uniform-convergence parameter of the game from Definition A.22, $c(G, \epsilon/8)$, to argue that for any $i \in [m]$, there exists a strategy α'_i for Player 1, that is a uniform distribution on $c(G, \epsilon/8)$ actions (possibly with repetitions), which approximates the strategy α_i up to an error of $\epsilon/8$. More precisely, for any strategy β of player 2 and any $i \in [m]$,

$$|u(\alpha'_i, \beta) - u(\alpha_i, \beta)| \leq \frac{\epsilon}{8}$$

Similarly, there exist strategies $\beta'_i, i \in [m]$ supported on $c(G, \epsilon/8)$ actions, such that for any strategy α of player 1,

$$|u(\alpha, \beta'_i) - u(\alpha, \beta_i)| \leq \frac{\epsilon}{8}.$$

Thus, for strategies α'_i, β'_j we will have $u(\alpha'_i, \beta'_j) \leq \theta + \frac{\epsilon}{4}$ if $i > j$ and $u(\alpha'_i, \beta'_j) \geq \theta + \frac{3\epsilon}{4}$ if $i \leq j$.

We want to bound the ϵ -fat threshold dimension of the pure strategy game. To do so, let us take the graph K_m . Suppose $Q = c(G, \epsilon/8)^2$. We determine how to color an edge (i, j) , where $i > j$, in the following manner: create two matrices of dimensions $Q \times Q$ each, denoted by $A^{i,j}$ and $B^{i,j}$. The matrix $A^{i,j}$ is the utility matrix of the subgame where Player 1 is restricted to play from $\text{support}(\alpha'_i)$ and Player 2 from $\text{support}(\beta'_j)$. In particular, denoting $\text{support}(\alpha'_i) = \{a_{i,1}, \dots, a_{i,c(G, \epsilon/8)}\}$ and $\text{support}(\beta'_j) = \{b_{j,1}, \dots, b_{j,c(G, \epsilon/8)}\}$, we define $A^{i,j}_{k,\ell} = u(a_{i,k}, b_{j,\ell})$. Similarly, the matrix $B^{i,j}$ corresponds to the game where Player 1 plays from $\text{support}(\alpha_j)$ and Player 2 from $\text{support}(\beta_i)$, where $B^{i,j}_{k,\ell} = u(a_{j,k}, b_{i,\ell})$.

Since α'_i and β'_j are each uniform distributions over their supports, then for all $i > j$, $u(\alpha'_i, \beta'_j) = \frac{1}{Q^2} \sum_{k,\ell} A^{i,j}_{k,\ell}$ and similarly $u(\alpha'_j, \beta'_i) = \frac{1}{Q^2} \sum_{k,\ell} B^{i,j}_{k,\ell}$. Since for $i > j$, $u(\alpha'_i, \beta'_j) \leq \theta + \frac{\epsilon}{4}$ and $u(\alpha'_j, \beta'_i) \geq \theta + \frac{3\epsilon}{4}$, it holds that $\frac{1}{Q^2} \sum_{k,\ell} (B^{i,j}_{k,\ell} - A^{i,j}_{k,\ell}) \geq \frac{\epsilon}{2}$.

Suppose out of the Q^2 pairs (k, l) , δQ^2 are such that $B^{i,j}_{k,\ell} - A^{i,j}_{k,\ell} \geq \frac{\epsilon}{4}$ and for the rest we have that $B^{i,j}_{k,\ell} - A^{i,j}_{k,\ell} < \frac{\epsilon}{4}$. That implies that

$$\frac{\epsilon}{2} \leq \frac{1}{Q^2} \sum_{k,\ell} (B^{i,j}_{k,\ell} - A^{i,j}_{k,\ell}) \leq \frac{(1-\delta)Q^2 \frac{\epsilon}{4} + \delta Q^2}{Q^2} = (1-\delta) \frac{\epsilon}{4} + \delta \implies \delta \geq \frac{\frac{\epsilon}{4}}{1 - \frac{\epsilon}{4}} \geq \frac{\epsilon}{4}$$

Consider now all the pairs (k, l) for which $B^{i,j}_{k,\ell} - A^{i,j}_{k,\ell} \geq \frac{\epsilon}{4}$ holds. Divide the interval of $[0, 1]$ into $\lceil \frac{8}{\epsilon} \rceil$ intervals of size $\frac{\epsilon}{8}$ ($I_1 = [0, \frac{\epsilon}{8}]$, $I_2 = [\frac{\epsilon}{8}, \frac{2\epsilon}{8}]$, $\dots$) and assign the pair (k, l) in the interval which $A^{i,j}_{k,\ell}$ belongs to. If the pair belongs to interval I_r , we color the edge (i, j) with color (k, l, r) . Note that in total there are $Q^2 \lceil \frac{8}{\epsilon} \rceil$ colors.

Suppose we can form a monochromatic clique of size t in the graph K_m we constructed, with vertices $u_1, u_2, \dots, u_t$. That means that there exist actions $a_1, a_2, \dots, a_t \in \mathcal{A}$ and $b_1, b_2, \dots, b_t \in \mathcal{B}$ such that $u(a_i, b_j) \leq \theta'$ if $i > j$ and $u(a_i, b_j) \geq \theta' + \frac{\epsilon}{8}$ if $i \leq j$. That implies that the ϵ -fat threshold dimension of the game is at least t . Above, we constructed a complete graph, where each edge is multicolored with $\frac{\epsilon Q^2}{4}$ colors out of the available $Q^2 \lceil \frac{8}{\epsilon} \rceil$ colors. By Corollary D.5, we have that a monochromatic clique in K_m of size $n+1$ will exist if $m \geq O\left(\left(\frac{1}{\epsilon}\right)^{\frac{1}{\epsilon} Q^2 n}\right) \geq R(n+1, Q^2 \lceil \frac{8}{\epsilon} \rceil, \frac{\epsilon Q^2}{4})$. As we assumed the ϵ -fat threshold dimension of the pure strategy game is exactly n , we need to have $m \leq O\left(\left(\frac{1}{\epsilon}\right)^{\frac{1}{\epsilon} Q^2 n}\right) = O\left(\left(\frac{1}{\epsilon}\right)^{\frac{1}{\epsilon} c(G, \epsilon/8)^2 n}\right)$ as otherwise we would have a threshold of at least $n+1$.

With a slight modification to the previous argument, we attain the following corollary for binary-valued games,

Corollary D.8 *Let G be a $\{0, 1\}$ -valued zero-sum game, then,*

$$\text{tr}(\text{conv}(G), \epsilon) \leq O\left((1/\epsilon)^{C \text{VC}(G)^2 \text{tr}(G) \log(\text{VC}(G))}\right)$$

Proof As in the general case of real valued games, for all $i > j \in [m]$, we can construct the $Q \times Q$ matrix $A^{i,j}$ with entries $A^{i,j}_{k,\ell} = u(a_{i,k}, b_{j,\ell})$ for $a_{i,k} \in \text{support}(\alpha'_i)$ and $b_{j,\ell} \in \text{support}(\beta'_j)$ and the $Q \times Q$ matrix $B^{i,j}$ with entries $B^{i,j}_{k,\ell} = u(a_{j,k}, b_{i,\ell})$ for $a_{j,k} \in \text{support}(\alpha'_j)$ and $b_{i,\ell} \in \text{support}(\beta'_i)$. Again, we have $\frac{1}{Q^2} \sum_{k,\ell} (B^{i,j}_{k,\ell} - A^{i,j}_{k,\ell}) \geq \frac{\epsilon}{2}$.

Now, importantly though, the matrices have values in $\{0, 1\}$. This means there must exist at least 1 pair (k, l) for each (i, j) with $B^{i,j}_{k,\ell} - A^{i,j}_{k,\ell} = 1$. So, rather than discretizing the $[0, 1]$ interval and coloring each edge i, j using $Q^2 \lceil \frac{8}{\epsilon} \rceil$ colors, we can use simply Q^2 colors, and a monochromatic clique will imply a threshold matrix in the pure strategy game. Even though each (i, j) pair receives one color here, the reduction in number of colors gives us the improved $m \leq O\left((Q^2)^{Q^2 n}\right)$, as desired. ■

D.2 A bound based on the sequential fat-shattering dimension of the original game

In this section, we prove one of the two bounds from Theorem 2.3, stated as the following lemma:

Lemma D.9 Let $\mathcal{F}$ be a $[0, 1]$ -valued concept class. Then, there exist universal constants $C_1, C_2 > 0$ such that

$$\text{fattr}(\text{conv}(\mathcal{F}), \epsilon) \leq 2^{C_1 \text{sfat}(\mathcal{F}, \epsilon/C_2)/\epsilon^2}.$$

Further, the same holds when $\text{conv}(\mathcal{F})$ is replaced with either $\text{dconv}(\mathcal{F})$ and $\text{conv2}(\mathcal{F})$, and when $\mathcal{F}$ is replaced with a zero-sum game G .

We use the definition of the sequential Rademacher complexity of a class $\mathcal{F}$ [Rakhlin et al., 2010]:

Definition D.10 Given a concept class $\mathcal{F}$ over a domain X , given horizon length $T > 0$, and given functions $Z_0, \dots, Z_{T-1}$ where $Z_t: \{-1, 1\}^t \rightarrow X$, define the *sequential Rademacher complexity* of $\mathcal{F}$ with respect to $Z = (Z_0, \dots, Z_{T-1})$ as

$$\text{sRad}(\mathcal{F}, T, Z) = \mathbb{E}_{\epsilon_1, \dots, \epsilon_T} \left[\sup_{f \in \mathcal{F}} \sum_{t=1}^T \epsilon_t f(Z_{t-1}(\epsilon_1, \dots, \epsilon_{t-1})) \right],$$

where $\epsilon_1, \dots, \epsilon_T$ are sampled uniformly and independently from $\{-1, 1\}$. Define $\text{sRad}(\mathcal{F}, T) = \sup_Z \text{sRad}(\mathcal{F}, T, Z)$.

The following are upper and lower bounds on the sequential Rademacher complexity given the sequential fat-shattering dimension were initially proved by [Rakhlin et al., 2010, Proposition 9] and the upper bound was improved by Block et al. [2021].

Lemma D.11 Let $\mathcal{F}$ be a concept class, and let $\epsilon > 0$. Then, for any $T \geq \text{sfat}(\mathcal{F}, \epsilon)$, it holds that

$$\text{sRad}(\mathcal{F}, T) \geq c\epsilon \sqrt{T \text{sfat}(\mathcal{F}, \epsilon)},$$

where $c > 0$ is a universal constant. Further, if $\mathcal{F}$ is $[0, 1]$ -valued, then for any $T \geq 1$,

$$\text{sRad}(\mathcal{F}, T) \leq C \left(T\epsilon + \int_{\epsilon}^1 \sqrt{T \text{sfat}(\mathcal{F}, r)} dr \right),$$

where $C > 0$ is a universal constant.

Further, we use the following lemma [Rakhlin et al., 2010, Lemma 3]:

Lemma D.12 Let $\mathcal{F}$ be a concept class. Then, $\text{sRad}(\mathcal{F}) = \text{sRad}(\text{conv}(\mathcal{F}))$.

Consequently, we obtain the following bound:

Lemma D.13 Let $\mathcal{F}$ be a $[0, 1]$ -valued concept class and let $\epsilon \in (0, 1]$. Then,

$$\text{sfat}(\text{conv}(\mathcal{F}), \epsilon) \leq \frac{C_1 \text{sfat}(\mathcal{F}, \epsilon/C_2)}{\epsilon^2}$$

where $C_1, C_2 > 0$ are universal constants.

Proof Let $T = \text{sfat}(\text{conv}(\mathcal{F}), \epsilon)$, denote by c, C the constants of Lemma D.11, let $\epsilon' > 0$. Apply Lemma D.11, Lemma D.12 and the fact that $\text{sfat}(\mathcal{F}, r)$ is monotonic decreasing in r to obtain:

$$\begin{aligned} c\epsilon \sqrt{T \text{sfat}(\text{conv}(\mathcal{F}), \epsilon)} &\leq \text{sRad}(\text{conv}(\mathcal{F}), T) = \text{sRad}(\mathcal{F}, T) \\ &\leq C \left(T\epsilon' + \int_{\epsilon'}^1 \sqrt{T \text{sfat}(\mathcal{F}, r)} dr \right) \leq CT\epsilon' + C\sqrt{T \text{sfat}(\mathcal{F}, \epsilon')} \\ &= C\epsilon' \sqrt{T \text{sfat}(\text{conv}(\mathcal{F}), \epsilon)} + C\sqrt{T \text{sfat}(\mathcal{F}, \epsilon')}. \end{aligned}$$

Setting $\epsilon' = c/(2C)$, one obtains that

$$\frac{c\epsilon}{2} \sqrt{T \text{sfat}(\text{conv}(\mathcal{F}), \epsilon)} \leq C\sqrt{T \text{sfat}(\mathcal{F}, \epsilon')}.$$

Consequently,

$$\text{sfat}(\text{conv}(\mathcal{F}), \epsilon) \leq \left(\frac{2C}{c} \right)^2 \frac{\text{sfat}(\mathcal{F}, \epsilon')}{\epsilon^2}.$$

This concludes the proof. ■

We use that $\text{fattr}(\mathcal{F}, \epsilon) \leq 2^{\lceil \text{sfat}(\mathcal{F}, \epsilon) \rceil}$ to obtain the desired bound. In order to obtain the bounds where $\text{conv}(\mathcal{F})$ is replaced with either $\text{dconv}(\mathcal{F})$ or $\text{conv2}(\mathcal{F})$, we notice that the proof follows from the same arguments, replacing Lemma D.12 with $\text{sRad}(\mathcal{F}, T) = \text{sRad}(\text{conv}(\mathcal{F}), T) = \text{sRad}(\text{dconv}(\mathcal{F}), T) = \text{sRad}(\text{conv2}(\mathcal{F}), T)$, whose proof follows similar arguments as the proof of Lemma D.12. Lastly, for a game G , the bound follows directly from the bound for concept classes $\mathcal{F}$, since the various dimensions of G are obtained by considering the relevant concept classes.

D.3 A lower bound

We prove two lower bounds on $\text{fat}_{\text{tr}}(\text{conv}(\mathcal{F}), \epsilon)$: one in terms of $\text{tr}(\mathcal{F})$ and another one in terms of $\text{Lit}(\mathcal{F})$, for some 0-1 classes $\mathcal{F}$. This is stated in the following two lemmas.

Lemma D.14 *For any $n \in \mathbb{N}$, $n \geq 2$ and any $\epsilon < 1$, there exists a 0-1 valued concept class $\mathcal{F}$ with $\text{fat}_{\text{tr}}(\text{conv}(\mathcal{F}), \epsilon) = n$ whereas $\text{tr}(\mathcal{F}) \leq \mathcal{O}(\log n)$ (where $\mathcal{O}$ hides constants that depend only on ϵ).*

Proof We will prove for $\epsilon = 0.2$ however the proof for any ϵ follows the same arguments with the numerical constants changed. We will define a distribution over concept classes and show that with probability greater than 0, a random concept class from this distribution satisfies the desired properties. Define the random concept class $\mathcal{F}$ as a union of classes $\mathcal{F}_1, \dots, \mathcal{F}_n$ over the domain $X = [n]$. Each $\mathcal{F}_i$ contains ℓ different elements, $f_{i,1}, \dots, f_{i,\ell}$ where $\ell = \Theta(\log n)$ is to be determined exactly later. We independently, for each $i \in [n]$, $k \in [\ell]$ and $j \in [n]$, define

$$f_{i,k}(j) = \begin{cases} 1 \text{ w.p } 0.7 \text{ and } 0 \text{ w.p } 0.3 & \text{if } i \leq j \\ 1 \text{ w.p } 0.3 \text{ and } 0 \text{ w.p } 0.7 & \text{if } i > j. \end{cases}$$

We notice that for each $i \leq j \in [n]$, from Chernoff-Hoeffding bound,

$$\Pr \left[\frac{1}{\ell} \sum_{k=1}^{\ell} f_{i,k}(j) \geq 0.6 \right] \geq 1 - e^{-c\ell},$$

where $c > 0$ is a universal constant, and similarly for any $i > j \in [n]$,

$$\Pr \left[\frac{1}{\ell} \sum_{k=1}^{\ell} f_{i,k}(j) \leq 0.4 \right] \geq 1 - e^{-c\ell}.$$

We set $\ell = \log(3n^2)/c$ and notice that $e^{-c\ell} \leq 1/(3n^2)$. By a union bound over $i, j \in [n]$, we get that with probability at least $2/3$, for all $i, j \in [n]$:

$$\begin{aligned} \sum_{k=1}^{\ell} f_{i,k} &\geq 0.6 & \text{if } i \leq j \\ \sum_{k=1}^{\ell} f_{i,k} &\leq 0.4 & \text{if } i > j \end{aligned}$$

If this holds, then $\text{fat}_{\text{tr}}(\text{conv}(\mathcal{F}), 0.2) = n$. This is obtained by taking the convex combinations $\{\frac{1}{\ell} \sum_{k=1}^{\ell} f_{i,k}\}_{i \in [n]}$ as the functions f_i in Definition A.10.

Lastly, we will show that with probability at least $1/3$, $\text{tr}(\mathcal{F}) \leq O(\log n)$. Let $M > 0$ and we will see that if $M \geq \Omega(\log n)$ (with a sufficiently large constant), there are no functions $f^1, \dots, f^M \in \mathcal{F}$ and elements $x_1, \dots, x_M \in X$ such that for all $i, j \in [M]$, $f^i(x_j) = \mathcal{I}(i \leq j)$. Fix some functions $f^1, \dots, f^M \in \mathcal{F}$ and $x_1, \dots, x_M \in X$, and notice that

$$\Pr[\forall i, j \in [n], f^i(x_j) = \mathcal{I}(i \leq j)] \leq 0.7^{M^2},$$

since $f^i(x_j)$ are independently chosen and each value for $f^i(x_j)$ can be taken with probability at most 0.7 and there are M^2 values to be satisfied. We note that there are at most $(\ell n)^M n^M \leq O(n^2 \log n)^M$ choices for $\{f^1, \dots, f^M, x_1, \dots, x_M\}$, and by a union bound over all choices, we have

$$\begin{aligned} &\Pr[\text{tr}(\mathcal{F}) \geq M] \\ &= \Pr[\exists f^1, \dots, f^M \in \mathcal{F}, \exists x_1, \dots, x_M \in X, \text{ s.t. } \forall i, j \in [M], f^i(x_j) = \mathcal{I}(i \leq j)] \\ &\leq (Cn^2 \log n)^M 0.7^{M^2} \leq e^{C_1 \log(n)M - C_2 M^2}, \end{aligned}$$

where C, C_1, C_2 are universal constants. If we set $M = C_3 \log(n)$ for a sufficiently large constant C_3 , then the probability above is bounded by $1/3$. We obtain that with probability $1/3$, $\text{fat}_{\text{tr}}(\text{conv}(\mathcal{F}), 0.2) = n$ and $\text{tr}(\mathcal{F}) \leq O(\log(n))$ as required. $\blacksquare$

Lastly, we notice the following lower bound $\text{fat}_{\text{tr}}(\text{conv}(\mathcal{F}), \epsilon)$ in terms of Littlestone's dimension of the class:

Lemma D.15 *For any n there exists a 0-1 valued class $\mathcal{F}$ such that $\text{Lit}(\mathcal{F}) \leq \log n$ while $\text{fat}(\text{conv}(\mathcal{F}), \epsilon) = n$, for all $\epsilon < 1$. Similarly, there a class $\mathcal{F}$ such that $\text{Lit}(\mathcal{F}) \leq \log n$ while $\text{fat}(\text{dconv}(\mathcal{F}), \epsilon) = n$.*

Proof For the first part of the lemma that involves $\text{conv}(\mathcal{F})$, this follows from the fact that $\text{fat}(\text{conv}(\mathcal{F}), \epsilon) \leq \text{Lit}(\mathcal{F})$, that $\text{Lit}(\mathcal{F}) \leq \log |\mathcal{F}|$, and that for n there exists a class $\mathcal{F}$ of n elements with $\text{tr}(\mathcal{F}) = n$: Indeed, consider the class $\mathcal{F} = \{f_1, \dots, f_n\}$ where $f_i: [n] \rightarrow \{0, 1\}$ defined by $f_i(j) = \mathcal{I}(i \leq j)$. This class has threshold dimension n by definition. The second part of the lemma, that involves $\text{dconv}(\mathcal{F})$, is proved similarly. ■

E CCE

Recall our definition of the CCE-matrix of a game

Definition 4.4 (The CCE-matrix of a game) *For a game $G = (\mathcal{A} = \prod_{p=1}^k \mathcal{A}_p, u = (u_1, \dots, u_k))$, the CCE-matrix $\mathbf{M}_{\text{CCE}}^G: \mathcal{A} \times \left(\bigcup_p \mathcal{A}_p\right)$ is defined, for $a \in \mathcal{A}$ and $d_p \in \mathcal{A}_p$ as*

$$\mathbf{M}_{\text{CCE}}^G[a, (p, d_p)] = u_p(d_p, a_{-p}) - u_p(a_p, a_{-p})$$

As stated in Section 4.2, in order to compute a CCE of G , we would like to run something analogous to Algorithm 3 on the zero-sum matrix game on $\mathbf{M}_{\text{CCE}}^G$. To do so, we need two things:

1. Bounds on the dimension parameters of $\mathbf{M}_{\text{CCE}}^G$
2. Best-response oracles for the two players of the game

For item 1, viewing $\mathbf{M}_{\text{CCE}}^G$ as a $[-1, 1]$ -concept class $\{f_a | a \in \mathcal{A}\}$ over the domain set $X = \bigcup_p \mathcal{A}_p$, we prove the following.

Lemma E.1 *Let $G = (\mathcal{A}, u)$ be a k -player game with bounded utilities $u_p: \mathcal{A} \rightarrow [0, 1]$ for all p . Then, the combinatorial dimensions of the game G bound those of the $\mathbf{M}_{\text{CCE}}^G$ concept class as follows*

$$\begin{aligned} \text{fat}(\mathbf{M}_{\text{CCE}}^G, 16\epsilon) &\leq O(k \text{fat}(G, \epsilon) \log \text{fat}(G, \epsilon) \log^2(1/\epsilon)) \\ \text{sfat}(\mathbf{M}_{\text{CCE}}^G, 2\epsilon) &\leq O((k/\epsilon) \text{sfat}(G, \epsilon)) \\ \text{fat}(\mathbf{M}_{\text{CCE}}^G, 2\epsilon) &\leq (k/\epsilon) \text{fat}(\mathcal{A}, \epsilon) \end{aligned}$$

Proof We introduce the notation $\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p]$ to denote the submatrix of $\mathbf{M}_{\text{CCE}}^G$ containing only the columns in $\mathcal{A}_p$. We can express $\mathbf{M}_{\text{CCE}}^G$ as the horizontal concatenation of matrices:

$$\mathbf{M}_{\text{CCE}}^G = [\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_1], \dots, \mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_k]]$$

We will bound the dimensions of $\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p]$ for all p and use that to bound the dimensions of $\mathbf{M}_{\text{CCE}}^G$ using the horizontal concatenation lemmas of Appendix F. We partition the rows of each matrix $\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p]$ as follows. For each integer $z \in [0, 1/\epsilon]$, define

$$S_{p,z,\epsilon} = \{a \in \mathcal{A} | u_p(a_p, a_{-p}) \in [z\epsilon, (z+1)\epsilon)\}$$

We can express $\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p]$ as the vertical concatenation of matrices:

$$\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p] = \begin{bmatrix} \mathbf{M}_{\text{CCE}}^G[S_{p,0,\epsilon}, \mathcal{A}_p] \\ \mathbf{M}_{\text{CCE}}^G[S_{p,1,\epsilon}, \mathcal{A}_p] \\ \vdots \end{bmatrix}$$

We will bound the dimensions of $\mathbf{M}_{\text{CCE}}^G[S_{p,z,\epsilon}, \mathcal{A}_p]$ for all z and use that to bound the dimensions of $\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p]$ using the vertical concatenation lemmas of Appendix F. Recall the utility concept class of player p . $\mathcal{F}_p$ is defined over the domain set $X_p = \mathcal{A}_p$, and has concepts $f_{a_{-p}}$ parameterized by elements $a_{-p} \in \mathcal{A}_{-p}$

defined as $f_{a_{-p}}(d_p) := u_p(d_p, a_{-p})$, for each $d_p \in \mathcal{A}_p$. Note, for all $(a_p, a_{-p}) \in S_{p,z,\epsilon}$ we have $u_p(a_p, a_{-p}) \in [z\epsilon, (z+1)\epsilon]$. Therefore, for all $(a_p, a_{-p}) \in S_{p,z,\epsilon}$ and $d_p \in \mathcal{A}_p$,

$$\mathbf{M}_{\text{CCE}}^G[(a_p, a_{-p}), (p, d_p)] - (z+1/2)\epsilon \in [f_{a_{-p}}(d_p) - \epsilon/2, f_{a_{-p}}(d_p) + \epsilon/2]$$

Since the $(z+1/2)\epsilon$ -shifted concepts of $\mathbf{M}_{\text{CCE}}^G[S_{p,z,\epsilon}, \mathcal{A}_p]$ correspond to concepts of $\mathcal{F}_p$ up to an additive factor of $\epsilon/2$, any shattering structure present in $\mathbf{M}_{\text{CCE}}^G[S_{p,z,\epsilon}, \mathcal{A}_p]$ with a margin of 2ϵ must exist in $\mathcal{F}_p$ with a margin of at least ϵ . This is due to the fact that the $(z+1/2)\epsilon$ shift can be incorporated in the witness parameters of the structure θ . Thus,

$$\begin{aligned} \text{fat}(\mathbf{M}_V^G[S_{p,z,\epsilon}, \mathcal{A}_p], 2\epsilon) &= \text{fat}(\mathcal{F}_p, \epsilon) && \leq \text{fat}(G, \epsilon) \\ \text{sfat}(\mathbf{M}_V^G[S_{p,z,\epsilon}, \mathcal{A}_p], 2\epsilon) &= \text{sfat}(\mathcal{F}_p, \epsilon) && \leq \text{sfat}(G, \epsilon) \\ \text{fattr}(\mathbf{M}_V^G[S_{p,z,\epsilon}, \mathcal{A}_p], 2\epsilon) &= \text{fattr}(\mathcal{F}_p, \epsilon) && \leq \text{fattr}(G, \epsilon) \end{aligned} \quad (17)$$

Using the vertical concatenation lemmas of Appendix F, as well as equation (17), we conclude

$$\text{fat}(\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p], 16\epsilon) \leq O(\text{fat}(G, \epsilon) \log \text{fat}(G, \epsilon) \log^2(1/\epsilon)) \quad (18)$$

$$\text{sfat}(\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p], 2\epsilon) \leq O((1/\epsilon) \text{sfat}(G, \epsilon)) \quad (19)$$

$$\text{fattr}(\mathbf{M}_{\text{CCE}}^G[\mathcal{A}, \mathcal{A}_p], 2\epsilon) \leq (1/\epsilon) \text{fattr}(G, \epsilon) \quad (20)$$

where (18) follows from Lemma F.1, (19) follows from Lemma F.3, and (20) follows from Lemma F.7. Then,

$$\text{fat}(\mathbf{M}_{\text{CCE}}^G, 16\epsilon) \leq O(k \text{fat}(G, \epsilon) \log \text{fat}(G, \epsilon) \log^2(1/\epsilon)) \quad (21)$$

$$\text{sfat}(\mathbf{M}_{\text{CCE}}^G, 2\epsilon) \leq O((k/\epsilon) \text{sfat}(G, \epsilon)) \quad (22)$$

$$\text{fattr}(\mathbf{M}_{\text{CCE}}^G, 2\epsilon) \leq (k/\epsilon) \text{fattr}(G, \epsilon) \quad (23)$$

where (21) follows from Lemma F.2, (22) follows from Lemma F.5, and (23) follows from Lemma F.8. $\blacksquare$

For item 2, we need to be able to compute two things. We have a zero-sum matrix game on $\mathbf{M}_{\text{CCE}}^G : \mathcal{A} \times (\bigcup_p \mathcal{A}_p) \rightarrow [-1, 1]$ between a minimizing player selecting distributions over action profiles $\mu \in \Delta\{\mathcal{A}\}$ and a maximizing player selecting distributions over deviations $\xi \in \bigcup_p \mathcal{A}_p$. The corresponding “ ϵ -best-response” oracles would return, for all $\xi \in \bigcup_p \mathcal{A}_p$, $\hat{a} = \text{BestResponse}(\xi, \epsilon) \in \mathcal{A}$ satisfying

$$\mathbb{1}[\hat{a}]^T \mathbf{M}_{\text{CCE}}^G \xi \leq \inf_{a \in \mathcal{A}} \mathbb{1}[a]^T \mathbf{M}_{\text{CCE}}^G \xi + \epsilon \quad (24)$$

and, for all $\mu \in \Delta(\mathcal{A})$, $(\hat{p}, d_{\hat{p}}) = \text{BestResponse}(\mu, \epsilon) \in \bigcup_p \mathcal{A}_p$ satisfying

$$\mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(\hat{p}, d_{\hat{p}})] \geq \sup_{(p, a_p) \in \bigcup_p \mathcal{A}_p} \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(p, a_p)] - \epsilon \quad (25)$$

As stated in Section 4.2, (24) is not necessary. It is used at time step t in Algorithm 3 when the minimizing player is increasing her set of actions $A_{t-1} \rightarrow A_t$ in order to improve her value in the game versus the maximizing player’s current set of actions B_{t-1} . For the specific case of the matrix game $\mathbf{M}_{\text{CCE}}^G$, we have the following construction of A_t in terms of B_{t-1} .

Algorithm 6 Finite-CCE

Input: A finite game $((B^1, \dots, B^k), u)$, an $\epsilon > 0$

Subroutines:

- Multiplicative-weight update: a no-regret learning algorithm that operates on a finite set of actions (Algorithm A.2)
 - 1. Simulate a repeated game between the k players, where each player plays according to multiplicative weight update (Algorithm A.2), for $T = \Theta(\log(\max_p |B_p|) / \epsilon^2)$ iterations
 - 2. **Return** the uniform distribution over $\{a^1, \dots, a^T\}$ where a^t is the strategy profile played by the players at iteration t .
-

From Lemma A.24, we have that the output of Algorithm 6 will constitute an ϵ -CCE for the subgame $G' = (\prod_p B_{t-1,p}, u)$. Therefore, we have

$$\text{Val}_{\mathbf{M}_{\text{CCE}}^G}(\text{Support}(\text{Finite-CCE}(B_{t-1})), B_{t-1}) \leq \epsilon \quad (26)$$

By the definition of ϵ -CCE, there is no profitable deviation for any player to any strategy within the subgame G' . Therefore, the outputted distribution μ must satisfy $\sup_{\xi \in \Delta(B_{t-1})} \mu^T \mathbf{M}_{\text{CCE}}^G \xi \leq \epsilon$, and Player 1 can force the game to have value $\leq \epsilon$ using a strategy supported on $\text{Support}(\text{Finite-CCE}(B_{t-1}))$. This suffices for the purposes of our algorithm. To achieve (25), we introduce the following subroutine

Algorithm 7 ϵ -best deviation for an action profile distribution

Input: A bounded-support distribution over action profiles $\mu \in \Delta(\mathcal{A})$, a value $\epsilon > 0$

Subroutines:

- **BestResponse_p** oracle: for a player p , receives a bounded-support distribution over adversary actions $\mu_{-p} \in \Delta(\mathcal{A}_{-p})$ and an $\epsilon > 0$ and outputs an ϵ -best response from $\mathcal{A}_p$ (see Definition A.23).
- Utility functions $u_p(a_p, a_{-p})$ for all players p that receive an action profile and outputs player p 's utility

1. **For** $p = 1, 2, \dots, k$

- (a) **Marginalize:** $\mu_{-p}[a_{-p}] \leftarrow \sum_{a_p \in \text{Support}(\mu)_p} \mu[(a_p, a_{-p})]$ **for** $a_{-p} \in \text{Support}(\mu)_{-p}$
- (b) **Best deviation for p :** $d_p \leftarrow \text{BestResponse}_p(\mu_{-p}, \epsilon)$
- (c) **Value of not deviating for p :** $v_p \leftarrow \sum_{a \in \text{Support}(\mu)} u_p(a_p, a_{-p}) \mu[a]$
- (d) **Value of best deviation for p :** $v'_p \leftarrow \sum_{a_{-p} \in \text{Support}(\mu)_{-p}} u_p(d_p, a_{-p}) \mu_{-p}[a_{-p}]$

2. **Player with greatest value increase if she deviates:** $\hat{p} \leftarrow \arg \max_{p \in [k]} v'_p - v_p$

3. **Return** $(\hat{p}, d_{\hat{p}})$

The following lemma demonstrates that this subroutine gives the desired ϵ -best-response of (25).

Lemma E.2 (Best deviation for an action profile distribution) *Consider a game*

$G = (\mathcal{A} = \prod_{p=1}^k \mathcal{A}_p, u = (u_1, \dots, u_k))$, *bounded-support distribution over action profiles* $\mu \in \Delta(\mathcal{A})$, *and a value* $\epsilon > 0$. *Then, Algorithm 7 executed with parameters* μ, ϵ *outputs a* $(\hat{p}, d_{\hat{p}})$ *satisfying*

$$\mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(\hat{p}, d_{\hat{p}})] \geq \sup_{(p, a_p) \in \bigcup_p \mathcal{A}_p} \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(p, a_p)] - \epsilon$$

Proof Recalling Definition A.23, our ϵ -best-response oracle for player p receives a bounded-support distribution $\mu_{-p} \in \Delta(\mathcal{A}_{-p})$ and outputs an action $d_p \in \mathcal{A}_p$ satisfying

$$\sum_{a_{-p} \in \mathcal{A}_{-p}} u_p(d_p, a_{-p}) \mu_{-p}[a_{-p}] \geq \sup_{a_p \in \mathcal{A}_p} \left(\sum_{a_{-p} \in \mathcal{A}_{-p}} u_p(a_p, a_{-p}) \mu_{-p}[a_{-p}] \right) - \epsilon$$

Therefore, for all p , the values v_p, v'_p in Algorithm 7 satisfy

$$\begin{aligned} v'_p - v_p &= \sum_{a_{-p} \in \mathcal{A}_{-p}} \sum_{a_p \in \mathcal{A}_p} (u_p(d_p, a_{-p}) - u_p(a_p, a_{-p})) \mu[(a_p, a_{-p})] \\ &\geq \sup_{a_p \in \mathcal{A}_p} \left(\sum_{a_{-p} \in \mathcal{A}_{-p}} \sum_{a_p \in \mathcal{A}_p} (u_p(a_p, a_{-p}) - u_p(d_p, a_{-p})) \mu[(a_p, a_{-p})] \right) - \epsilon \end{aligned}$$

or equivalently

$$v'_p - v_p = \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(p, d_p)] \geq \sup_{a_p \in \mathcal{A}_p} \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(p, a_p)] - \epsilon$$

Therefore, defining p^* such that $\widehat{(p, a_p)} \in \mathcal{A}_{p^*}$ where

$$\widehat{(p, a_p)} = \arg \sup_{(p, a_p) \in \bigcup_p \mathcal{A}_p} \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(p, a_p)]$$

we have

$$\begin{aligned} \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(\hat{p}, d_{\hat{p}})] &\geq \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(p^*, d_{p^*})] \\ &\geq \sup_{a_{p^*} \in \mathcal{A}_{p^*}} \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(p, a_p)] - \epsilon \\ &= \sup_{(p, a_p) \in \bigcup_p \mathcal{A}_p} \mu^T \mathbf{M}_{\text{CCE}}^G \mathbb{1}[(p, a_p)] - \epsilon \end{aligned}$$

as desired. ■

We are now ready to state our main algorithm and prove our main theorem.

Algorithm 8 $O(\epsilon)$ -approximate CCE for a general-sum game

Input: A general-sum game $G = (\mathcal{A} = \prod_{p=1}^k \mathcal{A}_p, u = (u_1, \dots, u_k))$, a parameter $\epsilon > 0$

Subroutines:

- **BestDeviation:** Receives a bounded-support distribution over action profiles $\mu \in \Delta(\mathcal{A})$, and a value $\epsilon > 0$. Returns an ϵ -best-response from $\bigcup_p \mathcal{A}_p$ for the matrix game on $\mathbf{M}_{\text{CCE}}^G$
- **Nash:** Receives finite set of action profiles $A \subset \mathcal{A}$, and an $\epsilon > 0$. Returns an ϵ -Nash for the two-player zero-sum subgame $(A, \bigcup_p \mathcal{A}_p, \mathbf{M}_{\text{CCE}}^G)$ using **BestDeviation** as its **BestResponse** subroutine (Algorithm 5)
- **FiniteCCE:** Receives a finite game $((B_1, \dots, B_k), u)$ and returns an ϵ -CCE
- **Val:** Receives finite sets of action profiles $A \subset \mathcal{A}$ and of deviations $B \subset \bigcup_p \mathcal{A}_p$. Returns the value of this finite subgame on $\mathbf{M}_{\text{CCE}}^G$.

1. $A_0 \leftarrow \emptyset, B_0 \leftarrow \{(1, a_1), \dots, (k, a_k)\}$, where $(a_1, \dots, a_k) \in \mathcal{A}$ are arbitrary actions

2. **For** $t = 1, 2, \dots$

- (a) $A_t \leftarrow A_{t-1} \cup \text{Support}(\text{FiniteCCE}((B_{t-1,1}, \dots, B_{t-1,k}), u))$ where $B_{t-1,p} = B_{t-1} \cap \mathcal{A}_p$
 - (b) $(\xi^{t,1}, \xi^{t,2}) \leftarrow \text{Nash}(A_t, \bigcup_p \mathcal{A}_p, \epsilon)$
 - (c) $B_t \leftarrow B_{t-1} \cup \text{Support}(\xi^{t,2})$
 - (d) **if** $\text{Val}(A_t, B_t) \leq 3\epsilon$
 - i. **Return** $\xi^{t,1}$
-

As stated previously, this algorithm is analogous to Algorithm 3 run on the game matrix $\mathbf{M}_{\text{CCE}}^G$, with a modification to how the minimizing player adds actions to her support (using **FiniteCCE**). Analogously, we have the following theorem.

In a similar fashion with section C.2, we will use three main results to prove Theorem 4.3, bounding the total number of oracle calls made by Algorithm 8. First, we show that the output of Algorithm 8, given that it stops, constitutes an $O(\epsilon)$ -CCE (Lemma E.3). Then, we show that Algorithm 8 necessarily stops after $\text{fatr}(\text{conv}(\mathbf{M}_{\text{CCE}}^G), 2\epsilon)$ iterations (Lemma E.4). Last, we show that Algorithm 8 running for T iterations makes only $O(kT/\epsilon^2 \cdot \log(T/\epsilon^2))$ oracle calls (Lemma E.5).

Lemma E.3 Assume that Algorithm 8 stops. Then, the returned strategies constitute a 5ϵ -CCE for the original game $G = (\mathcal{A}, u)$.

Proof The returned $\xi^{t,1}$ is the strategy for player 1 in an ϵ -Nash for the game $(A_t, \bigcup_p \mathcal{A}_p, \mathbf{M}_{\text{CCE}}^G)$. So, for every deviation $(p, d_p) \in \bigcup_p \mathcal{A}_p$:

$$\mathbf{M}_{\text{CCE}}^G[\xi^{t,1}, (p, d_p)] \leq \text{Val}(A_t, \bigcup_p \mathcal{A}_p) + \epsilon$$

Since $\xi^{t,2}$ is the strategy for player 2 in an ϵ -Nash for the game $(A_t, \bigcup_p \mathcal{A}_p, \mathbf{M}_{\text{CCE}}^G)$, we have:

$$\text{Val}(A_t, \bigcup_p \mathcal{A}_p) \leq \text{Val}(A_t, B_t) + \epsilon.$$

Lastly, by the stopping condition of Algorithm 8

$$\text{Val}(A_t, B_t) \leq 3\epsilon.$$

Combining these 3 equations gives the desired $\mathbf{M}_{\text{CCE}}^G[\xi^{t,1}, (p, d_p)] \leq 5\epsilon$ for all $(p, d_p) \in \bigcup_p \mathcal{A}_p$. $\blacksquare$

Lemma E.4 *Algorithm 8 terminates after $\text{fatr}(\text{conv}(\mathbf{M}_{\text{CCE}}^G), 2\epsilon)$ iterations.*

Proof Suppose the algorithm runs for $\geq T$ iterations. Then, we will have support sets $A_1 \subseteq A_2 \subseteq \dots \subseteq A_T$ and $B_1 \subseteq B_2 \subseteq \dots \subseteq B_T$ satisfying the following. For all $1 \leq i, j \leq T$,

$$\begin{aligned} \text{Val}_{\mathbf{M}_{\text{CCE}}^G}(A_i, B_j) &\leq \epsilon & \text{if } i > j \\ \text{Val}_{\mathbf{M}_{\text{CCE}}^G}(A_i, B_j) &\geq 3\epsilon & \text{if } i \leq j \end{aligned}$$

The first holds due to (26) and the fact that $\text{Support}(\text{Finite-CCE}(B_{t-1})) \subseteq A_t$. The second holds because the continuation of the algorithm implies the stopping condition is not met. Defining α_i to be the minmax strategy of Player 1 in the subgame $(A_i, B_{i-1}, \mathbf{M}_{\text{CCE}}^G)$ and β_i to be the minmax strategy of Player 2 in the subgame $(A_i, B_i, \mathbf{M}_{\text{CCE}}^G)$ for all i , we have

$$\begin{aligned} \mathbf{M}_{\text{CCE}}^G[\alpha_i, \beta_j] &\leq \epsilon & \text{if } i > j \\ \mathbf{M}_{\text{CCE}}^G[\alpha_i, \beta_j] &\geq 3\epsilon & \text{if } i \leq j \end{aligned}$$

which constitutes a (2ϵ) -fat-thresholding matrix in $\mathbf{M}_{\text{CCE}}^G$. Therefore, $\text{fatr}(\text{conv}(\mathbf{M}_{\text{CCE}}^G), 2\epsilon) \geq T$ as desired. $\blacksquare$

In order to bound the number of oracle calls, we add the following lemma:

Lemma E.5 *Assume that Algorithm 8 runs for T iterations. Then, the number of oracle calls is bounded by $\mathcal{O}(kT/\epsilon^2 \cdot \log(T/\epsilon^2))$.*

Proof First, we would like to bound the sizes of A_t and B_t by $\mathcal{O}(t/\epsilon^2 \cdot \log(t/\epsilon^2))$. In order to show that, notice that B_t is obtained from B_{t-1} by adding the support of an ϵ -approximate Nash for the half-infinite game $(A_t, \bigcup_p \mathcal{A}_p, \mathbf{M}_{\text{CCE}}^G)$. We would like to bound the size of the support of the strategy of Player 2 in such an approximate Nash. This approximate Nash is computed in Algorithm 5, and the size of the support equals the number of iterations of this algorithm, which is bounded by $\mathcal{O}(\log |A_t|/\epsilon^2)$, by Lemma C.2. Further, we would like to argue that $|A_t| \leq |A_{t-1}| + C(\log |B_{t-1}| + C)/\epsilon^2$ for a universal constants $C > 0$. Indeed, this is true since A_t is obtained from A_{t-1} by adding the support of a CCE computed by Algorithm 6, given an action-set taken from B_{t-1} , and the support size is bounded by $\mathcal{O}(\log(|B_{t-1}|)/\epsilon^2)$. By an inductive argument, it is easy to show that these two recursive equations for A_t and B_t imply that $|A_t|, |B_t| \leq \mathcal{O}(t/\epsilon^2 \cdot \log(t/\epsilon^2))$.

Lastly, it remains to bound the number of oracle calls. Notice that any addition of an action to the support of Player 2 involves computing a best deviation (Algorithm 7), which is being used as a subroutine in this instance of the half-infinite equilibrium computation (Algorithm 5). Algorithm 7 makes exactly k calls to the best response oracle. Thus, the total number of calls is bounded by $\mathcal{O}(kt/\epsilon^2 \cdot \log(t/\epsilon^2))$. $\blacksquare$

We are now ready to prove our main theorem. Recall

Theorem 4.3 *Let $G = (\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_k, u = (u_1, \dots, u_k))$ be a multi-player game. Assume that utilities are bounded $u_p : \mathcal{A} \rightarrow [0, 1]$, and let $\epsilon > 0$. Then, Algorithm 8 executed with parameters G, ϵ will compute an $O(\epsilon) - \text{CCE}$ for the game using the following number of ϵ -best response oracle calls:*

$$\mathcal{O}\left(e^{C(k/\epsilon^3) \text{sfat}(G, \epsilon/C)}\right)$$

Proof Proof of Theorem 4.3 We notice that Lemma E.3 implies that the output of Algorithm 8 is an $O(\epsilon)$ -CCE. Furthermore, Lemma E.4 bounds the number of iterations by $T \leq \text{fatr}(\text{conv}(\mathbf{M}_{\text{CCE}}^G), 2\epsilon)$ and Lemma E.5 implies that the number of oracle calls is bounded by $\tilde{\mathcal{O}}((k/\epsilon^2) \text{fatr}(\text{conv}(\mathbf{M}_{\text{CCE}}^G), 2\epsilon))$. The proof of Theorem 4.3 follows by substituting $\text{fatr}(\text{conv}(\mathbf{M}_{\text{CCE}}^G), 2\epsilon)$ according to Theorem 2.3:

$$\text{fatr}(\text{conv}(\mathbf{M}_{\text{CCE}}^G), 2\epsilon) \leq e^{C \text{sfat}(\mathbf{M}_{\text{CCE}}^G, 2\epsilon/C)/\epsilon^2}.$$

The total bound on the number of oracle calls is then

$$\begin{aligned} & \tilde{O}\left((k/\epsilon^2) \text{fattr}(\text{conv}(\mathbf{M}_{\text{CCE}}^G), 2\epsilon)\right) \\ & \leq (k/\epsilon^2) e^{C \text{sfat}(\mathbf{M}_{\text{CCE}}^G, 2\epsilon/C)/\epsilon^2} \end{aligned}$$

Notice that the fact of k/ϵ^2 can be omitted by changing the constant in the exponent. Lastly, plugging in our bounds on the dimensions of $\mathbf{M}_{\text{CCE}}^G$ in terms of those of G from Lemma E.1, we get that the number of oracle calls made by Algorithm 8 is

$$\leq e^{C(k/\epsilon^3) \text{sfat}(G, \epsilon/C)}$$

as desired. ■

F Concatenation Lemmas

Lemma F.1 (Vertical ϵ -fat-shattering concatenation tool) *For each $p \in [k]$, say we have a real-valued function class $\mathcal{F}_p$ defined on a domain set X with $\text{fat}(\mathcal{F}_p, \epsilon) \leq d$. Then, the “vertically-concatenated” real-valued function class $\bigcup_p \mathcal{F}_p$ has $\text{fat}(\bigcup_p \mathcal{F}_p, 8\epsilon) = O(\log(k) + d \log(d) \log^2(1/\epsilon))$.*

Proof For a real-valued function class $\mathcal{F}$, define the ϵ -covering growth function

$$\mathcal{N}_{\mathcal{F}, \epsilon}(n) = \max_{\substack{S \subseteq X \\ |S|=n}} \min\{|V| : V \text{ is an } \epsilon\text{-cover, under the } \infty\text{-norm, of } \mathcal{F} \text{ on } S\} \quad (27)$$

That is, we want to find the smallest $V \subseteq \mathbb{R}^S$ such that, for all $f \in \mathcal{F}$, there exists $v \in V$ such that, for all $x \in S$, $|f(x) - v(x)| \leq \epsilon$. For each $p \in [k]$, if V_p is an ϵ -cover of $\mathcal{F}_p$, then $\bigcup_p V_p$ ϵ -covers $\bigcup_p \mathcal{F}_p$. Therefore,

$$\mathcal{N}_{\bigcup_p \mathcal{F}_p, \epsilon}(n) \leq \sum_p \mathcal{N}_{\mathcal{F}_p, \epsilon}(n)$$

From Theorem 1.5 of Kakade and Tewari [2008], if $\text{fat}(\mathcal{F}, \epsilon) \leq d$, then

$$\mathcal{N}_{\mathcal{F}, 4\epsilon}(n) \leq O\left(\frac{n}{\epsilon^2}\right)^{\lceil d \log(\frac{en}{d\epsilon}) \rceil} \quad (28)$$

Given this, we can argue the following about the vertically-concatenated real-valued function class. Since a union of covers of the $\mathcal{F}_p$ for each p will cover $\bigcup_p \mathcal{F}_p$,

$$\mathcal{N}_{\bigcup_p \mathcal{F}_p, 4\epsilon}(n) \leq \sum_p \mathcal{N}_{\mathcal{F}_p, 4\epsilon}(n) \leq k O\left(\frac{n}{\epsilon^2}\right)^{\lceil d \log(\frac{en}{d\epsilon}) \rceil}$$

and for $n = \Omega(\log(k) + d \log(d) \log^2(1/\epsilon))$

$$\log\left(k O\left(\frac{n}{\epsilon^2}\right)^{\lceil d \log(\frac{en}{d\epsilon}) \rceil}\right) = \log(k) + d(1 + \log(n/d) + \log(1/\epsilon))(\log(n) + 2\log(1/\epsilon)) = o(n)$$

Therefore, for $n = C(\log(k) + d \log(d) \log^2(1/\epsilon))$ with sufficiently large constant C ,

$$\mathcal{N}_{\bigcup_p \mathcal{F}_p, 4\epsilon}(n) < 2^n$$

Therefore, for all $S = \{x_1, \dots, x_n\} \subseteq X$ and all $\theta \in \mathbb{R}^n$, there exists $b \in \{0, 1\}^n$ such that there is no $f \in \bigcup_p \mathcal{F}_p$ with

$$\begin{aligned} f(x_j) &> \theta_j + 8\epsilon && \text{for all } j \in [n] \text{ with } b_j = 1 \\ f(x_j) &\leq \theta_j && \text{for all } j \in [n] \text{ with } b_j = 0 \end{aligned} \quad (29)$$

If there existed such an $f_b \in \bigcup_p \mathcal{F}_p$ for every b , each would have to be covered by a distinct v in the minimal (4ϵ) -cover V . This would force the size of the cover to be at least 2^n , a contradiction. Therefore, $\text{fat}(\bigcup_p \mathcal{F}_p, 8\epsilon) = O(\log(k) + d \log(d) \log^2(1/\epsilon))$, as desired. ■

Lemma F.2 (Horizontal ϵ -fat-shattering concatenation tool) Say we have real-valued function classes $\mathcal{F}_p$ for all $p \in [k]$ defined on mutually-disjoint domain sets X_p with $\text{fat}(\mathcal{F}_p, \epsilon) \leq d$. Let's say these classes all have the same magnitude, and their elements are enumerated by a set $\mathcal{A}$. That is, for all p , there is an $f_{p,a} \in \mathcal{F}_p$ for each $a \in \mathcal{A}$. We define the “horizontal concatenation” function class $\mathcal{F} = \{f_a | a \in \mathcal{A}\}$ where each $f_a : \bigcup_p X_p \rightarrow \mathbb{R}$ is defined

$$f_a(x) = f_{p,a}(x) \quad \text{for all } x \in X_p \text{ for all } p \quad (30)$$

Then, the horizontal concatenation function class $\mathcal{F}$ has $\text{fat}(\mathcal{F}, \epsilon) \leq kd$.

Proof Assume for the sake of contradiction that there exists $S = \{x_1, \dots, x_{kd+1}\} \subseteq \bigcup_p X_p$ and witnesses $\theta_1, \dots, \theta_{kd+1}$ such that for every $b \in \{0, 1\}^{kd+1}$, there exists $a \in \mathcal{A}$ with

$$\begin{aligned} f_a(x_j) &\geq \theta_j + \epsilon && \text{for all } j \in [kd+1] \text{ with } b_j = 1 \\ f_a(x_j) &\leq \theta_j && \text{for all } j \in [kd+1] \text{ with } b_j = 0 \end{aligned} \quad (31)$$

By the pigeon hole principle, there must exist a $p \in [k]$ such that $|S \cap X_p| \geq d+1$. That would imply, for every $b \in \{0, 1\}^{|S \cap X_p|}$, there exists $a \in \mathcal{A}$ with

$$\begin{aligned} f_{p,a}(x_j) &\geq \theta_j + \epsilon && \text{for all } j \text{ with } x_j \in S \cap X_p \text{ and } b_j = 1 \\ f_{p,a}(x_j) &\leq \theta_j && \text{for all } j \text{ with } x_j \in S \cap X_p \text{ and } b_j = 0 \end{aligned} \quad (32)$$

contradicting our assumption that $\text{fat}(\mathcal{F}_p, \epsilon) \leq d$, as desired. $\blacksquare$

Lemma F.3 (Vertical ϵ -sequential-fat-shattering concatenation tool) For each $p \in [k]$, say we have a real-valued function class $\mathcal{F}_p$ defined on a domain set X with $\text{sfat}(\mathcal{F}_p, \epsilon) \leq d$. Then, the “vertically-concatenated” real-valued function class $\bigcup_p \mathcal{F}_p$ has $\text{sfat}(\bigcup_p \mathcal{F}_p, \epsilon) \leq k(d+1) - 1$.

Proof Assume for the sake of contradiction there exists a complete binary tree $T = (V, E)$ of depth $k(d+1)$, whose internal nodes $v \in V$ are labeled by elements $x(v) \in X$ and have witnesses $\theta(v)$, and whose leaves $\ell \in V$ are labeled by $f_\ell \in \bigcup_p \mathcal{F}_p$, such that the following holds: for any root-to-leaf path $v_1, \dots, v_{k(d+1)+1}, v_{k(d+1)+1} = \ell$ in the tree and for any $i \in [k(d+1)]$:

$$\begin{aligned} f_\ell(x(v_i)) &\geq \theta(v_i) + \epsilon && \text{if } v_{i+1} \text{ is a left child of } v_i \\ f_\ell(x(v_i)) &\leq \theta(v_i) && \text{if } v_{i+1} \text{ is a right child of } v_i \end{aligned}$$

We color the leaf nodes of the tree with k colors where $c(\ell) = p$ iff $f_\ell \in \mathcal{F}_p$. We will demonstrate that there exists a complete binary “subtree” of depth $d+1$ that is leaf-monochromatic. Here, we define a subtree to be a tree $T' = (V', E')$ on a subset of nodes $V' \subseteq V$ where every internal node $v \in V'$ satisfies:

$$\begin{aligned} v_{\text{left}} \text{ is the left child of } v \text{ in } T' &\text{ implies } v_{\text{left}} \text{ is a left descendant of } v \text{ in } T \\ v_{\text{right}} \text{ is the right child of } v \text{ in } T' &\text{ implies } v_{\text{right}} \text{ is a right descendant of } v \text{ in } T \end{aligned}$$

We also ensure that leaves in T' are leaves in T . The existence of this depth- $(d+1)$ complete leaf-monochromatic subtree would give the desired contradiction, implying that for some p , $\text{sfat}(\mathcal{F}_p) \geq d+1$. To prove the existence of this subtree, we will use the following lemma.

Lemma F.4 Define $R_{\text{leaf}}(C_1, \dots, C_k)$ to be the minimum integer such that any complete binary tree of depth $R_{\text{leaf}}(C_1, \dots, C_k)$ with k -colored leaves necessarily has, for some $p \in [k]$, a complete binary subtree of depth C_p with all leaves colored p . Then,

$$R_{\text{leaf}}(C_1, \dots, C_k) \leq \sum_p C_p$$

From the lemma, $R_{\text{leaf}}(d+1, \dots, d+1) \leq k(d+1)$ and therefore the desired leaf-monochromatic depth- $(d+1)$ subtree exists. $\blacksquare$

Proof Proof of Lemma F.4 We prove by induction on $\sum_p C_p$. For a base case, note that $R_{\text{leaf}}(0, \dots, 0) = 0$. For a depth-0 binary tree (consisting of 1 leaf node), whatever color we select for the leaf, the leaf itself will constitute the desired depth-0 leaf-monochromatic subtree.

Assume the claim holds for all $R_{\text{leaf}}(C'_1, \dots, C'_k)$ with $\sum_p C'_p \leq d$. Let's consider a k -leaf-colored depth- $(d+1)$ tree and $\sum_p C_p = d+1$. We note that the two child trees of the root have depth $d \geq R_{\text{leaf}}(C_1, \dots, C_p -$

$1, \dots, C_k)$ for an arbitrarily selected color p . If either child tree contains a q -leaf-monochromatic subtree of depth C_q for some $q \neq p$, we are done. So, from the inductive hypothesis, assume both child trees contain p -leaf-monochromatic subtrees of depth $C_p - 1$. These two trees together with the root give the desired p -leaf-monochromatic subtree of depth C_p . ■

Lemma F.5 (Horizontal ϵ -sequential-fat-shattering concatenation tool) *Say we have real-valued function classes $\mathcal{F}_p$ for all $p \in [k]$ defined on mutually-disjoint domain sets X_p with $\text{sfat}(\mathcal{F}_p, \epsilon) \leq d$. Let's say these classes all have the same magnitude, and their elements are enumerated by a set $\mathcal{A}$. That is, for all p , there is an $f_{p,a} \in \mathcal{F}_p$ for each $a \in \mathcal{A}$. We define the "horizontal concatenation" function class $\mathcal{F} = \{f_a | a \in \mathcal{A}\}$ where each $f_a : \bigcup_p X_p \rightarrow \mathbb{R}$ is defined*

$$f_a(x) = f_{p,a}(x) \quad \text{for all } x \in X_p \text{ for all } p$$

Then, the horizontal concatenation function class $\mathcal{F}$ has $\text{sfat}(\mathcal{F}, \epsilon) \leq kd$.

Proof Assume for the sake of contradiction there exists a complete binary tree $T = (V, E)$ of depth $kd + 1$, whose internal nodes $v \in V$ are labeled by elements $x(v) \in \bigcup_p X_p$ and have witnesses $\theta(v)$, whose leaves $\ell \in V$ are labeled by $f_\ell \in \mathcal{F}$, such that the following holds: for any root-to-leaf path $v_1, \dots, v_{kd+1}, v_{kd+2} = \ell$ in the tree and for any $i \in [kd + 1]$:

$$\begin{aligned} f_\ell(x(v_i)) &\geq \theta(v_i) + \epsilon && \text{if } v_{i+1} \text{ is a left child of } v_i \\ f_\ell(x(v_i)) &\leq \theta(v_i) && \text{if } v_{i+1} \text{ is a right child of } v_i \end{aligned}$$

We color the internal nodes of the tree with k colors where $c(v) = p$ iff $x(v) \in X_p$. We will demonstrate that there exists a complete binary subtree of depth $d + 1$ that is internally-monochromatic. The existence of this complete internally-monochromatic subtree would give the desired contradiction, implying that for some p , $\text{sfat}(\mathcal{F}_p) \geq d + 1$. To prove the existence of this subtree, we will use the following lemma.

Lemma F.6 *Define $R_{\text{int}}(C_1, \dots, C_k)$ to be the minimum integer such that any complete binary tree of depth $R_{\text{int}}(C_1, \dots, C_k)$ with k -colored internal nodes necessarily has, for some $p \in [k]$, a complete binary subtree of depth C_p with all internal nodes colored p . Then,*

$$R_{\text{int}}(C_1, \dots, C_k) \leq \sum_p C_p - k + 1$$

From the lemma, $R_{\text{int}}(d+1, \dots, d+1) \leq kd+1$ and therefore the desired internally-monochromatic depth- $(d+1)$ subtree exists. ■

Proof Proof of Lemma F.6 We prove by induction on $\sum_p C_p$. For a base case, note that $R_{\text{int}}(1, \dots, 1) = 1$. For a depth-1 binary tree (consisting of 1 internal node and 2 leaves), whatever color we select for the internal node, the tree itself will constitute the desired depth-1 internally-monochromatic subtree (leaves are not colored).

Assuming the claim holds for all $R_{\text{int}}(C'_1, \dots, C'_k)$ with $\sum_p C'_p \leq d$, let's consider a k -colored depth- $(d+1)$ tree. Assume without loss of generality that the root has color p . We note that its two child trees have depth $d \geq R_{\text{int}}(C_1, \dots, C_p - 1, \dots, C_k)$. If either child tree contains a q -internally-monochromatic subtree of depth C_q for some $q \neq p$, we are done. So, from the inductive hypothesis, assume both child trees contain p -internally-monochromatic subtrees of depth $C_p - 1$. These two trees together with the root of color p give the desired p -internally-monochromatic subtree of depth C_p . ■

Lemma F.7 (Vertical ϵ -fat-threshold concatenation tool) *For each $p \in [k]$, say we have a real-valued function class $\mathcal{F}_p$ defined on a domain set X with $\text{fatr}(\mathcal{F}_p, \epsilon) \leq d$. Then, the "vertically-concatenated" real-valued function class $\bigcup_p \mathcal{F}_p$ has $\text{fatr}(\bigcup_p \mathcal{F}_p, \epsilon) \leq kd$.*

Proof Assume for the sake of contradiction that there exists $F = \{f_1, \dots, f_{kd+1}\} \subseteq \bigcup_p \mathcal{F}_p$, $S = \{x_1, \dots, x_{kd+1}\} \subseteq X$, and witness θ such that for every $i, j \in [kd + 1]$

$$\begin{aligned} f_i(x_j) &\geq \theta + \epsilon && \text{for all } i \leq j \\ f_i(x_j) &\leq \theta && \text{for all } i > j \end{aligned} \tag{33}$$

By the pigeon hole principle, there must exist a $p \in [k]$ such that $|F \cap \mathcal{F}_p| \geq d + 1$. Therefore, for every i, j with $f_i, f_j \in F \cap \mathcal{F}_p$

$$\begin{aligned} f_i(x_j) &\geq \theta + \epsilon && \text{for all } i \leq j \\ f_i(x_j) &\leq \theta && \text{for all } i > j \end{aligned} \quad (34)$$

contradicting our assumption that $\text{fattr}(\mathcal{F}_p, \epsilon) \leq d$, as desired. $\blacksquare$

Lemma F.8 (Horizontal ϵ -fat-threshold concatenation tool) *Say we have real-valued function classes $\mathcal{F}_p$ for all $p \in [k]$ defined on mutually-disjoint domain sets X_p with ϵ -fat-threshold dimension $\leq d$. Let's say these classes all have the same magnitude, and their elements are enumerated by a set $\mathcal{A}$. That is, for all p , there is an $f_{p,a} \in \mathcal{F}_p$ for each $a \in \mathcal{A}$. We define the "horizontal concatenation" function class $\mathcal{F} = \{f_a | a \in \mathcal{A}\}$ where each $f_a : \bigcup_p X_p \rightarrow \mathbb{R}$ is defined*

$$f_a(x) = f_{p,a}(x) \quad \text{for all } x \in X_p \text{ for all } p \quad (35)$$

Then, the horizontal concatenation function class $\mathcal{F}$ has ϵ -fat-threshold dimension $\leq kd$.

Proof Assume for the sake of contradiction that there exists $A = \{f_1, \dots, f_{kd+1}\} \subseteq \mathcal{A}$, $S = \{x_1, \dots, x_{kd+1}\} \subseteq \bigcup_p X_p$, and witness θ such that for every $i, j \in [kd+1]$

$$\begin{aligned} f_i(x_j) &\geq \theta + \epsilon && \text{for all } i \leq j \\ f_i(x_j) &\leq \theta && \text{for all } i > j \end{aligned} \quad (36)$$

By the pigeon hole principle, there must exist a $p \in [k]$ such that $|S \cap X_p| \geq d+1$. Therefore, for every i, j with $x_i, x_j \in S \cap X_p$

$$\begin{aligned} f_{p,i}(x_j) &\geq \theta + \epsilon && \text{for all } i \leq j \\ f_{p,i}(x_j) &\leq \theta && \text{for all } i > j \end{aligned} \quad (37)$$

contradicting our assumption that $\text{fattr}(\mathcal{F}_p, \epsilon) \leq d$, as desired. $\blacksquare$

Lemma F.9 *Let $\mathcal{F}$ be a 0-1 valued function class. Define the XOR of 2 functions as follows: $(f \oplus g)(x) = \mathbb{I}(f(x) \neq g(x))$. We also define the XOR of a function class $\mathcal{F}$ and a function g as follows: $\mathcal{F} \oplus g = \{f \oplus g : f \in \mathcal{F}\}$. Then we have that:*

$$\text{tr}(\mathcal{F} \oplus g) \leq 2 \text{tr}(\mathcal{F}) + 1$$

Proof Consider the function class in which all the entries from $\mathcal{F}$ are flipped and name it $\bar{\mathcal{F}}$. Now let us concatenate the two function classes $\mathcal{F}$ and $\bar{\mathcal{F}}$ vertically, as the functions are the rows of the function class. Name this new function class $\mathcal{F}_{conc}$. Note that the threshold dimension of $\mathcal{F}_{conc}$ is going to be at most the sum of the threshold dimensions of $\mathcal{F}$ and $\bar{\mathcal{F}}$, i.e.

$$\text{tr}(\mathcal{F}_{conc}) \leq \text{tr}(\mathcal{F}) + \text{tr}(\bar{\mathcal{F}})$$

Moreover, one can notice that if $\mathcal{F}$ has threshold dimension $\text{tr}(\mathcal{F})$, then the threshold dimension of $\bar{\mathcal{F}}$ has to be at least $\text{tr}(\mathcal{F}) - 1$. Similarly, the threshold dimension of $\mathcal{F}$ has to be at least $\text{tr}(\bar{\mathcal{F}}) - 1$. So we have $\text{tr}(\mathcal{F}) \geq \text{tr}(\bar{\mathcal{F}}) - 1$. Thus we have:

$$\text{tr}(\mathcal{F}_{conc}) \leq \text{tr}(\mathcal{F}) + \text{tr}(\bar{\mathcal{F}}) \leq 2 \text{tr}(\mathcal{F}) + 1$$

Finally we will prove that $\text{tr}(\mathcal{F} \oplus g) \leq \text{tr}(\mathcal{F}_{conc})$, which combined with the above result will give us the required bound.

Recall that the rows of a function class are the functions. Let us take a look at what happens to the column j of $\mathcal{F}$ when we XOR it with the entry of the function g , g_j , at that column j . If $g_j = 0$ then the column is not swapped, otherwise if $g_j = 1$ the column is swapped. Therefore each column of $\mathcal{F} \oplus g$ is either a column of $\mathcal{F}$ or a column of $\bar{\mathcal{F}}$. That means that any threshold in $\mathcal{F} \oplus g$ will be contained in $\mathcal{F}_{conc}$, thus concluding our proof. $\blacksquare$

G Various inequalities

We first prove the following auxiliary lemma:

Lemma G.1 *Let $\theta_1, \dots, \theta_d$ be numbers in $[0, 1 - \epsilon]$ for some $\epsilon > 0$. Then, there exists some $\theta \in [0, 1]$ such that*

$$|\{i \in [d] : \theta_i \in [\theta - \epsilon/2, \theta]\}| \geq \frac{\epsilon d}{2}$$

Proof For each $i \in [d]$, if θ is drawn uniformly at random from $[0, 1]$, then with probability $\epsilon/2$ it holds that $\theta_i \in [\theta - \epsilon/2, \theta]$. Consequently, taking expectation over θ , the expected number of elements i such that $\theta_i \in [\theta - \epsilon/2, \theta]$ is $\epsilon d/2$. There exists some $\theta \in [0, 1]$ that realizes this expectation, namely, that there are at least $\epsilon d/2$ elements i such that $\theta_i \in [\theta - \epsilon/2, \theta]$, as required. ■

Using Lemma G.1, we proceed to providing a sketch of Lemma A.11:

Proof Proof sketch of Lemma A.11 let $d = \text{fat}(\mathcal{F}, \epsilon)$. By definition of the fat-shattering dimension, there exists a set $\{x_1, \dots, x_d\}$ and witnesses $(\theta_1, \dots, \theta_d)$ that satisfy Eq. (7). To prove $\text{fat}(\mathcal{F}, \epsilon) \leq \text{sfat}(\mathcal{F}, \epsilon)$, notice that we can construct a complete binary tree of depth d , such that all internal nodes of depth i are labelled by x_{i+1} , for $i = 0, \dots, d-1$. Further, Eq. (7) imply that one could label the leaves with appropriate functions $f \in \mathcal{F}$ such that Eq. (8) holds, which implies that $\text{sfat}(\mathcal{F}, \epsilon) \geq d$ as required.

For the inequality $\text{fat}(\mathcal{F}, \epsilon) \leq \frac{\text{fattr}(\mathcal{F}, \epsilon/2)}{2\epsilon}$, we use the same notation of $\{x_1, \dots, x_d\}$ and witnesses $(\theta_1, \dots, \theta_d)$ that satisfy Eq. (7), where $d = \text{fat}(\mathcal{F}, \epsilon)$. By Lemma G.1 there is some $\theta \in [0, 1]$ such that there exist at least $\epsilon d/2$ elements $i \in [d]$ such that $\theta_i \in [\theta, \theta + \epsilon/2]$. Let $i_1, \dots, i_m$ denote the set of indices of these θ_i variables, where $m \geq \epsilon d/2$. By Eq. (7), for each $j \in [m]$ there exists $f_j \in \mathcal{F}$ such that $f(x_{i_\ell}) \geq \theta_{i_j} + \epsilon \geq \theta + \epsilon$ for all $\ell \in \{j, j+1, \dots, m\}$ and such that $f(x_{i_\ell}) \leq \theta_{i_j} \leq \theta$ for all $\ell \in [j-1]$. This, by definition, implies that $\text{fattr}(\mathcal{F}, \epsilon/2) \geq m \geq \epsilon d/2 = \epsilon \text{fat}(\mathcal{F}, \epsilon)/2$, as required.

For the inequality $\text{fattr}(\mathcal{F}, \epsilon) \leq 2^{\text{sfat}(\mathcal{F}, \epsilon)+1}$, notice that it is equivalent to $\text{sfat}(\mathcal{F}, \epsilon) \geq \log \text{fattr}(\mathcal{F}, \epsilon) - 1$. It is sufficient to prove that $\text{sfat}(\mathcal{F}, \epsilon) \geq \lfloor \log \text{fattr}(\mathcal{F}, \epsilon) \rfloor$. To obtain that, denote $m = 2^{\lfloor \text{fattr}(\mathcal{F}, \epsilon) \rfloor}$ and notice that given functions $f_1, \dots, f_m \in \mathcal{F}$ elements $x_1, \dots, x_m \in \mathcal{X}$ and $\theta \in [0, 1]$ that satisfy Eq. (9), one could construct a tree whose internal nodes are labelled by $x_2, \dots, x_m$ and whose leaves are labelled by $f_1, \dots, f_m$, that satisfies Eq. (8).

The left part of Eq. (11) was proved by [Daskalakis and Golowich, 2022, Lemma 8.4]. ■