

Learning Convex Partitions and Computing Game-theoretic Equilibria from Best Response Queries[★]

Paul W. Goldberg¹ and Francisco J. Marmolejo-Cossío¹ ^{★★}

University of Oxford, UK

{paul.goldberg, francisco.marmolejo}@cs.ox.ac.uk

Abstract. Suppose that an m -simplex is partitioned into n convex regions having disjoint interiors and distinct labels, and we may learn the label of any point by querying it. The learning objective is to know, for any point in the simplex, a label that occurs within some distance ε from that point. We present two algorithms for this task: Constant-Dimension Generalised Binary Search (CD-GBS), which for constant m uses $\text{poly}(n, \log(\frac{1}{\varepsilon}))$ queries, and Constant-Region Generalised Binary Search (CR-GBS), which uses CD-GBS as a subroutine and for constant n uses $\text{poly}(m, \log(\frac{1}{\varepsilon}))$ queries. We show via Kakutani’s fixed-point theorem that these algorithms provide bounds on the best-response query complexity of computing approximate well-supported equilibria of bimatrix games in which one of the players has a constant number of pure strategies.

Keywords: Query protocol, equilibrium computation, revealed preferences

1 Introduction

The computation of game-theoretic equilibria is a topic of long-standing interest in the algorithmic and AI communities. This includes computation in the “classical” setting of complete information about a game, as well as settings of partial information, communication-bounded settings, and distributed algorithms (for example, best-response dynamics). A recent line of research has studied computation of equilibria based on query access to players’ payoff functions. That work, along with the notion of revealed preferences in economics, inspires the new setting we study here.

We study algorithms that have query access to the players’ best-response behaviour: an algorithm may query a mixed-strategy profile (i.e. probability distributions constructed by the algorithm, over each player’s pure strategies) and learn the players’ best responses. Our focus is on standard bimatrix games, which is arguably the most natural starting-point for an investigation of this new query model. The solution concept of interest is ε -approximate Nash equilibria (exact equilibria are typically impossible to find using finitely many such queries). A basic challenge is to identify algorithms that achieve this goal with good bounds on their query complexity (and also, ideally, their runtime complexity).

In more detail, we assume an $m \times n$ game G : a row player has m pure strategies and a column player has n pure strategies. G has two unknown $m \times n$ payoff matrices that represent payoffs to the players for all combinations of pure strategy choices they may make. A query consists of a probability distribution over the pure strategies of one of the players, and elicits an answer consisting of a best response for the other player (i.e. a pure strategy that maximises that player’s expected payoff). We seek an ε -well-supported Nash equilibrium (ε -WSNE): a pair of probability distributions over their pure strategies with the property that any strategy of player p whose expected payoff is more than ε below the value of p ’s best response, gets probability zero. The general question of interest is: how many queries are needed, as a function of m, n, ε .

Using Kakutani’s fixed point theorem, we reduce this question to a novel and more geometrical challenge in the design of query protocols. Suppose that the m -simplex Δ^m is partitioned into n convex regions having labels in $[n] = \{1, \dots, n\}$. When we query a point $x \in \Delta^m$ we are told the label of x . How many queries (in terms of m, n, ε) are needed in order to ensure that all points in Δ^m are within ε of a point whose label we know? We show how to achieve this using time and queries polynomial in $\log \varepsilon$ and $\max(m, n)$ provided that $\min(m, n)$ is constant. This leads to a polynomial query complexity algorithm for 2-player games, provided that one of the players has a constant number of strategies.

[★] Full Online Version of Paper: <https://arxiv.org/abs/1807.06170>

^{★★} Supported by the Mexican National Council of Science and Technology (CONACyT)

1.1 Further details

In essence, we consider partitions of the unit m -simplex Δ^m into n convex polytopes, P_1, \dots, P_n , with disjoint interiors, and aim to approximately learn the partition with access to a membership oracle that for a given $x \in \Delta^m$, returns a polytope to which x belongs. The notion of approximation we study is that of ε -close labellings, a collection of empirical polytopes, $\{\hat{P}_i\}_{i=1}^n$, such that $\hat{P}_i \subseteq P_i$ for $i = 1, \dots, n$ and $\cup_{i=1}^n \hat{P}_i$ is an ε -net of $\Delta^m \subset \mathbb{R}^m$ in the ℓ_2 norm.

Note that in one dimension ($m = 1$) we can use binary search to solve this problem using $n \log(1/\varepsilon)$ queries. We generalise to higher dimension, exploiting convexity of the regions to reduce query usage in computing ε -close labellings. We present two algorithms for this task: Constant-Dimension Generalised Binary Search (CD-GBS), which for constant m uses $\text{poly}(n, \log(\frac{1}{\varepsilon}))$ queries, and Constant-Region Generalised Binary Search (CR-GBS), which uses CD-GBS as a subroutine and for constant n uses $\text{poly}(m, \log(\frac{1}{\varepsilon}))$ queries.

This problem derives from the question of how to compute approximate (well-supported) Nash equilibria (ε -WSNE) using only best response information, obtained via queries in which the algorithm selects a mixed strategy profile and a player, and receives a best response for that player to the mixed profile. Via Kakutani's fixed-point theorem [18] we reduce this variant of equilibrium computation to finding ε -close labellings of polytope partitions. For $m \times n$ games where m is constant (or n equivalently, by symmetry), we show that an ε -WSNE can be computed using $\text{poly}(n, \log(\frac{1}{\varepsilon}))$ best response queries.

1.2 Related Work

Earlier work in computational learning theory has studied exact learning of geometrical regions over a discretised domain, where algorithms are sought with query complexity logarithmic in a resolution parameter and binary search is repeatedly applied in a systematic way [6]. Goldberg and Kwek [12] specifically study the learnability of polytopes in this context, deriving query efficient algorithms, and precisely classifying polytopes learnable in this setting. These algorithms can be adapted to approximately learn a single polytope with membership queries, but the obtained notion of approximation is not directly applicable to computing ε -close labellings.

The Nash equilibrium (NE) is a fundamental concept in game theory [21]. They are guaranteed to exist in finite games, yet computational challenges in finding one abound, most notably, the PPAD-completeness of computing an exact equilibrium even for two-player normal form games [7, 9]. For this reason, query complexity has been extensively used as a tool to differentiate hardness of equilibrium concepts in games. For payoff queries, some notable examples include: exponential lower bounds for randomised computation of exact Nash equilibria and exact correlated equilibria via communication complexity lower bounds in multiplayer games [16, 17]; exponential lower bounds for randomised computation of approximate well-supported equilibria and general approximate equilibria for a small enough approximation factor in multiplayer games [1]; upper and lower bounds for equilibrium computation in bimatrix games, congestion games [10] and anonymous games [15]; upper and lower bounds for randomised algorithms computing approximate correlated equilibria [13]. Babichenko et al. have also proved lower bounds in communication complexity for computing ε -WSNE for small enough ε in both bimatrix and multiplayer games [2].

Best response queries are a weaker but natural query model which is powerful enough to implement fictitious play, a dynamic first proposed by Brown [5], and proven to converge by Robinson [22] in two-player zero-sum games to an approximate NE. Fictitious play does not always converge for general games where both players have more than two strategies [11]. Furthermore, Daskalakis and Pan have proven that the rate of convergence of the dynamic is quite slow in the worst case (with arbitrary tie-breaking) [8]. Also, beyond non-convergence, the dynamic can have a poor approximation value for general games [14]. In addition, the relationship between best responses and convex partitions of simplices has been studied by Von Stengel [23] in the context of sequential games where one player has to commit to and announce a strategy before playing.

For a bimatrix game, simple ε -close labellings can be constructed by querying best responses at mixed strategies arising as uniform distributions over sufficiently large multisets of pure strategies. As a consequence of our main theorem, best responses to these multiset distributions contain enough information to compute approximate WSNE. This result is in the spirit of [3] and [20], who aim to quantify specific k such that some approximate equilibrium arises as a uniform mixture over multisets of size k . We note

in our scenario that there is also a guaranteed existence of an approximate equilibrium using sufficiently large multisets, however *verifying* that a *specific* pair of mixed strategies is an approximate WSNE is not straightforward using only best response queries. This is in contrast to the verification of approximate equilibria via utility queries as studied in [3].

Separately, we note that the present paper is possibly relevant to the search for a price equilibrium in certain markets. Baldwin and Klemperer study markets consisting of *strong-substitutes* demand functions for N different goods available in multiple discrete units [4]. These markets are a generalisation of the *product-mix auction* of [19]; a basic task is to identify prices at which some desired bundle of the goods is demanded. Consider the space $(\mathbb{R}^+)^N$ of all price vectors. As analysed in [4], a strong-substitutes demand function partitions this price space into convex polytopes, each of which comprises the prices at which some particular bundle of goods is demanded. So, the present paper relates to a setting where price vectors may be queried, and responses consist of demand bundles. The connection is imperfect, since the main objective in the context of [4] would be to learn a price at which some target bundle is demanded, rather than the entire demand function. The ideas here may be useful for learning the values that the market has for various bundles.

2 Preliminaries and Notation

Our main object of study will be families of polytopes that precisely cover the unit simplex, with the property that any two distinct polytopes from the family are either disjoint, meet at their boundary, or entirely coincide. Throughout, the polytopes we work with are convex.

Definition 1 (((m, n) -Polytope Partition). A (m, n) -polytope partition is a set of n convex polytopes, $\mathcal{P} = \{P_1, \dots, P_n\}$ such that $\bigcup P_i = \Delta^m = \{x \in \mathbb{R}^m \mid \forall i, x_i \geq 0, \sum_i x_i \leq 1\}$ and for each $i \neq j$, either $\text{relint}(P_i) \cap \text{relint}(P_j) = \emptyset$ or $P_i = P_j$ ($\text{relint}(H)$ being the relative interior of H).




Fig. 1. Polytope partition, cross-section and slices.

Definition 2 (Cross-sections and Slices). Let $P \subset \mathbb{R}^m$ be a polytope and $\pi : \mathbb{R}^m \rightarrow \mathbb{R}$ the projection function into the first coordinate. For $x \in \mathbb{R}$, we define the x -cross-section of P as $P^x = \pi^{-1}(x) \cap P$. For any $I = [x, y] \subset \mathbb{R}$ we define the $[x, y]$ -slice of P as $P^I = P^{x,y} = \pi^{-1}([x, y]) \cap P$. Suppose that $\mathcal{P} = \{P_i\}_i$ is an (m, n) -polytope partition. The definitions of cross-sections and slices extend to $\mathcal{P}^x = \{P_i^x\}_i$ and $\mathcal{P}^I = \mathcal{P}^{x,y} = \{P_i^{x,y}\}_i$.

Figure 1 gives a visualisation of these two definitions. Notice that in the same figure, \mathcal{P}^x is essentially a lower-dimensional polytope partition linearly scaled by a factor of $(1 - x)$. This however, is not the case in general. We distinguish between these two scenarios with the following formal definition:

Definition 3 (Non-Degenerate and Degenerate cross-sections). Let \mathcal{P} be an (m, n) -polytope partition. For $x \in [0, 1)$ let $f_x : \mathcal{P}^x \rightarrow \Delta^{m-1}$ be defined as $f_x(v_1, \dots, v_m) = \frac{1}{1-x}(v_2, \dots, v_m)$. If $f_x(\mathcal{P}^x)$ is an $(m-1, n)$ -polytope partition, we say that \mathcal{P}^x is a non-degenerate cross-section. Otherwise we say that \mathcal{P}^x is a degenerate cross-section.

The recursive structure of polytope partitions on non-degenerate cross-sections will be crucial to our constructions. Luckily, Lemma 1 (the full proof can be found in the online version of the paper) shows that for any polytope partition, there are only a finite number of points $x \in [0, 1]$ that give rise to degenerate cross-sections. Before stating the lemma, we define an important discrete subset of $[0, 1]$ given by the projections of vertices of polytopes under π .

Definition 4 (Vertex Critical Coordinates). *For a given polytope $P \subset \mathbb{R}^m$ let $V_P \subset \mathbb{R}^m$ be the vertex set of P . Define the set of vertex critical coordinates as $C_P = \pi(V_P) \subset \mathbb{R}$. If $\mathcal{P} = \{P_i\}_{i=1}^n$ is an (m, n) -polytope partition, then we extend our definition to include $C_{\mathcal{P}} = \bigcup_{i=1}^n C_{P_i} \subset [0, 1]$ as the vertex critical coordinates of \mathcal{P} .*

Lemma 1. *Let \mathcal{P} be an (m, n) -polytope partition and $x \in [0, 1] \setminus C_{\mathcal{P}}$. Then \mathcal{P}^x is non-degenerate.*

2.1 Membership Query Oracles and ε -close Labellings

We study two natural query oracle models for polytope membership in any \mathcal{P} .

Definition 5 (Membership Query Oracles for Polytope Partitions). *Any (m, n) -polytope partition, \mathcal{P} has the following membership query oracles:*

- *Lexicographic query oracle:* $Q_\ell : \Delta^m \rightarrow [n]$, which for a given y returns the smallest index of polytope to which y belongs, namely $Q_\ell(y) = \min\{i \in [n] \mid y \in P_i\}$.
- *Adversarial query oracle(s):* $Q_A : \Delta^m \rightarrow [n]$, which can return any polytope to which y belongs. Namely, Q_A is any function such that $Q_A(y) \in \{i \in [n] \mid y \in P_i\}$ for all $y \in \Delta^m$.

When we wish to refer to an arbitrary oracle from the above models, we use the notation Q . Before continuing, we also clarify that for $A, B \subseteq \mathbb{R}^n$, we denote $\text{Conv}(A, B) \subseteq \mathbb{R}^m$ as the convex combination of the two sets. In addition, if $A_i \subseteq \mathbb{R}^m$ is an indexed family of sets with $i = 1, \dots, r$, we denote $\text{Conv}(A_i \mid i = 1, \dots, r) \subseteq \mathbb{R}^n$ as the convex combination of all A_i . Upon making queries to Q , we can infer labels of $x \in \Delta^m$ by taking convex combinations. We abstract this notion in the following definition.

Definition 6 (Empirical Polytopes and Labellings). *Suppose that \mathcal{P} is an (m, n) -polytope partition and $S \subset \Delta^m$ is a finite set for which queries to Q have been made. Let $\hat{P}_i = \text{Conv}(\{x \in S \mid Q(x) = i\}) \subset P_i$. We say each \hat{P}_i is an empirical polytope of P_i and that $\hat{\mathcal{P}} = \{\hat{P}_i\}$ is an empirical labelling of \mathcal{P} . Furthermore, we use the notation $\hat{P}_\perp = \Delta^m \setminus \bigcup_{i=1}^n \hat{P}_i$ to refer to points in Δ^m unlabelled under $\hat{\mathcal{P}}$.*

An ε -net in the ℓ_2 norm for $\Delta^m \subset \mathbb{R}^m$ is a set $N_\varepsilon^m \subseteq \Delta^m$ with the property that for all $x \in \Delta^m$, there exists a $y \in N_\varepsilon^m$ such that $\|x - y\|_2 \leq \varepsilon$. Our learning goal is to use query access to an oracle, Q , to compute an empirical labelling $\hat{\mathcal{P}}$ such that $\bigcup_{i=1}^n \hat{P}_i$ is an ε -net of Δ^m .

Definition 7 (ε -close Labelling). *Suppose that $\varepsilon \geq 0$ and that $\hat{\mathcal{P}}$ is an empirical labelling for \mathcal{P} . If $\bigcup_{i=1}^n \hat{P}_i$ is an ε -net of $\Delta^m \subset \mathbb{R}^m$ in the ℓ_2 norm, we say that $\hat{\mathcal{P}}$ is an ε -close labelling of \mathcal{P} .*

Although ε -close labellings are defined for polytope partitions, we extend our terminology to also encompass slices of polytope partitions. As such, when we mention computing an ε -close labelling of $\mathcal{P}^{x,y}$, we mean an empirical labelling of $\mathcal{P}^{x,y}$ (in the same vein as Definition 6) with the property that the union of its empirical polytopes forms an ε -net of $(\Delta^m)^{x,y}$.

2.2 Learning in Thickness to Learning in Distance

Definition 8 (Thickness of Sets). *Suppose that $Z \subseteq \mathbb{R}^m$ is a set. We define the thickness of Z as the radius of the largest ℓ_2 ball fully contained in Z and we denote it by $\tau(Z) = \sup\{\delta \geq 0 \mid \exists x \in Z \text{ with } B_\delta(x) \subseteq Z\}$ where $B_\delta(x) = \{y \in \mathbb{R}^m \mid \|x - y\|_2 \leq \delta\}$. In the language of convex geometry, $\tau(Z)$ is the depth of the Chebyshev centre of Z .*

For a polytope partition \mathcal{P} , if $\hat{\mathcal{P}}$ is an ε -close labelling, then $\tau(\hat{P}_\perp) \leq \varepsilon$, but the converse does not hold in general. Even though \hat{P}_\perp may be of small thickness, if it contains vertices of Δ^m , these vertices may be far from labelled points. The following results (with full proofs in the online version of the paper) lead up to Lemma 4, a slightly weaker version of the aforementioned converse. Lemma 4 shows that if we are able to learn an empirical labelling where the set of unlabelled points is of small enough thickness, then we will in fact have succeeded in learning an ε -close labelling, where any unlabelled point is close in distance to a labelled point.

Lemma 2. *Let $P \subset \mathbb{R}^m$ be a full-dimensional polytope with $\text{Diam}(P) = \sup_{x,y \in P} \|x - y\|_2$.*

- *If $A \subsetneq P$ and $\gamma > \left(\frac{\text{Diam}(P)}{\tau(P)}\right) \tau(A)$, then $B_\gamma(x) \cap (P \setminus A) \neq \emptyset$ for all $x \in A$.*
- *If $A \subseteq P$ is such that $\text{int}(P) \setminus A \neq \emptyset$ ($\text{int}(P)$ refers to the interior of P) and $\gamma > \left(\frac{\text{Diam}(P)}{\tau(P)}\right) \tau(A)$, then $B_\gamma(x) \cap (\text{int}(P) \setminus A) \neq \emptyset$ for all $x \in A$.*

Lemma 3. *$\text{Diam}(\Delta^m) = \sqrt{2}$ and $\tau(\Delta^m) \geq \frac{1}{m+\sqrt{m}}$.*

Lemma 4. *Suppose that \mathcal{P} is an (m, n) -polytope partition. Furthermore suppose that $\hat{\mathcal{P}}$ is an empirical labelling with $\tau(\hat{P}_\perp) < \varepsilon$. For any $\gamma > \sqrt{2}(m+\sqrt{m})\varepsilon$, it follows that $\hat{\mathcal{P}}$ is a γ -close labelling. In particular, if $\gamma > 4m\varepsilon$, the claim also holds.*

Proof. From Lemma 3, we know that $\tau(\Delta^m) \geq \frac{1}{m+\sqrt{m}}$ and $\text{Diam}(\Delta^m) = \sqrt{2}$. Suppose that $x \in \hat{P}_\perp$. From Lemma 2 our choice of γ implies $B_\gamma(x) \cap (\Delta_m \setminus \hat{P}_\perp) \neq \emptyset$. This in turn means that $\hat{\mathcal{P}}$ is a γ -close labelling. As for the final claim, this holds since $m \geq 1$.

3 Constant-Dimension Generalised Binary Search for Q_ℓ

We set up important groundwork by focusing on arbitrary polytopes $P \subset \mathbb{R}^m$. We let $\pi : \mathbb{R}^m \rightarrow \mathbb{R}$ be the projection function introduced in Definition 2, and we recall Definition 4 regarding the vertex critical coordinates of P denoted by C_P .

Lemma 5. *Suppose that $x, y \in \mathbb{R}$ are such that $[x, y] \cap C_P = \emptyset$. Then taking convex hulls of cross-sections we get $\text{Conv}(P^x, P^y) = P^{x,y}$.*

Proof. $[x, y] \cap C_P = \emptyset$ implies the vertices of the polytope $P^{x,y}$ lie in \mathcal{P}_x and \mathcal{P}_y . Since the convex hull of the set of all vertices of a bounded polytope is the polytope itself, the claim follows.




Fig. 2. $\text{Conv}(P^a, P^b) \neq P^{a,b}$ and $\text{Conv}(P^c, P^d) = P^{c,d}$

This property of polytopes whereby convex combinations give rise to complete information except when traversing a discrete set of critical points (visualised in Figure 2) is critical to CD-GBS. With query access to polytopes however, we no longer fully recover P_x perfectly, but instead an approximation given by an ε -close labelling, \hat{P}_x . It becomes more subtle to show that by taking convex hulls of \hat{P}_x and \hat{P}_y , we recover the desired information along $[x, y]$.

3.1 Necessary Machinery

We delve into the specifics of CD-GBS by defining important machinery. All proofs for the lemmas of this section can be found in the full online version of the paper. To begin, we recall thickness from Definition 8, and see that it satisfies a sub-additivity property when the sets being considered are convex polytopes:

Lemma 6. *Let $P_1, \dots, P_k \subseteq \mathbb{R}^m$ be convex polytopes. Then $\tau(\cup_i P_i) < \frac{10}{3}(\sum_i \tau(P_i))(m+1)^{3/2}$.*

For a given polytope partition $\mathcal{P} = \{P_i\}_i$, it will be important to establish thickness bounds on P_i at specific cross-sections.

Definition 9 (α -Critical Coordinates). *Let $P \subset \mathbb{R}^m$ be a polytope. For $\alpha > 0$, we define $l_\alpha(P) = \inf\{x \in \mathbb{R} \mid \tau(P^x) \geq \alpha\}$ and $r_\alpha(P) = \sup\{x \in \mathbb{R} \mid \tau(P^x) \geq \alpha\}$ so that $\forall z \in \mathbb{R}, \tau(P^z) \geq \alpha$ if and only if $z \in [l_\alpha(P), r_\alpha(P)]$ (Here thickness is with respect to the natural embedding of P^x in \mathbb{R}^{m-1}). These are called α -critical coordinates for P .*

The previous definition allows us to associate to each polytope P_i a segment of $[0, 1]$ within which cross-sections of P_i are thick above a threshold. By combining this with Definition 4 we get the correct notion of critical coordinates.

Definition 10 (Critical Coordinates of a (m, n) -Polytope Partition). *Suppose that $\mathcal{P} = \{P_1, \dots, P_n\}$ is an (m, n) -polytope partition. For $\alpha > 0$, we let $C_\mathcal{P}^\alpha$ be the union of the sets of all vertex critical coordinates of all P_i as defined in Definition 4, and the set of all α -critical coordinates for all P_i as in Definition 9. Specifically, $C_\mathcal{P}^\alpha = (\cup_i C_{P_i}) \cup (\cup_i \{l_\alpha(P_i), r_\alpha(P_i)\})$.*

CD-GBS clusters queries around critical coordinates (up to a desired tolerance). For this reason it is important to bound the number of critical coordinates in a given (m, n) -Polytope partition.

Lemma 7. *If \mathcal{P} is a (m, n) -polytope partition $|C_\mathcal{P}^\alpha| \leq \binom{n+m}{m} + 2n$.*

With this machinery in hand, we are in a position to prove the main auxiliary lemma necessary to demonstrate correctness of CD-GBS. We show that if $x, y \in [0, 1]$ are such that $[x, y]$ contains no critical coordinates, then computing sufficiently fine empirical labellings of \mathcal{P}^x and \mathcal{P}^y with Q_ℓ will contain enough information to compute an ε -close labelling of $\mathcal{P}^{x,y}$ by simply taking convex combinations of the empirical labellings at both cross-sections.

Lemma 8. *Given $m, n, \varepsilon > 0$ let $\alpha = \frac{\varepsilon}{20nm^{5/2}}$ and $\beta = \frac{\varepsilon^2}{85nm^{5/2}}$. Suppose that \mathcal{P} is an (m, n) -polytope partition and that the following hold:*

- $x, y \in [0, 1]$ are such that $x < y \leq 1 - \frac{\varepsilon}{3}$.
- $[x, y] \cap C_\mathcal{P}^\alpha = \emptyset$.
- $\hat{\mathcal{P}}^x$ and $\hat{\mathcal{P}}^y$ are empirical labellings of \mathcal{P}^x and \mathcal{P}^y computed via Q_ℓ , such that $\cup_i \hat{P}_i^x$ and $\cup_j \hat{P}_j^y$ are β -nets for $(\Delta^m)^x$ and $(\Delta^m)^y$ respectively.

Then $\bigcup_i \text{Conv}(\hat{P}_i^x, \hat{P}_i^y)$ is an ε -net of $(\Delta^m)^{x,y}$.

For the following corollary, suppose that \mathcal{P} is an (m, n) polytope partition and that $0 = t_0 < t_1, \dots, < t_k = 1$ are points in $[0, 1]$. Furthermore suppose that $\beta = \frac{\varepsilon^2}{85nm^{5/2}}$ as in Lemma 8. For each t_i , if $t_i \notin C_\mathcal{P}^\alpha$, let $\hat{\mathcal{P}}^{t_i}$ be a β -close labelling of \mathcal{P}^{t_i} , otherwise let $\hat{\mathcal{P}}^{t_i} = \emptyset$. Let $\hat{\mathcal{P}} = \text{Conv}_i(\hat{\mathcal{P}}^{t_i})$ and for $i = 1, \dots, k$, let $I_j = [t_{j-1}, t_j]$. If $\hat{\mathcal{P}}^{t_{i-1}, t_i}$ is an ε -close labelling of $\mathcal{P}^{t_{i-1}, t_i}$, we say that I_j is covered, otherwise we say I_j is uncovered.

Corollary 1. *For any collection of $\{t_i\}_{i=1}^k$, there are no more than $2C_\mathcal{P}^\alpha$ intervals I_j that are uncovered.*

3.2 Specification of CD-GBS and Query Usage

Terms and Notation: The details of CD-GBS are presented in Algorithm 1. We recall our notation from Definition 3 where for $x \in [0, 1]$ we defined $f_x : (\Delta^m)^x \rightarrow \Delta^{m-1}$ given by $f_x(x, \dots, v_m) = \frac{1}{1-x}(v_2, \dots, v_m)$. We note that this is a bijection between both polytopes, hence it is well-defined to use f_x^{-1} . In addition, we let $\mathcal{D}^k = \{\frac{i}{2^k} \mid 1 \leq i \leq 2^k\}$ be the dyadic fractions of k -th power in the unit interval (excluding 0). For every $x \in \mathcal{D}^k$ we can associate the interval $I_x^k = [x - \frac{1}{2^k}, x]$. For each of these intervals $\text{midpoint}(I_x^k)$ denotes its midpoint. We also use the same language as Corollary 3.1 when we talk about whether I_x^k is covered or not (with respect to the current empirical labelling, $\hat{\mathcal{P}}$, obtained from taking convex hulls of labels in Δ^m). We note that in order to have a well-defined base case of CD-GBS (which is equivalent to binary search), we let $\Delta^0 = \mathbb{R}^0 = \{0\}$. Finally, we say that a point $x \in [0, 1]$ is an uncovered critical point if \mathcal{P}^x is computed via a recursive call to CD-GBS and for $(a, b) = B_{\varepsilon/2}(x) \cap [0, 1]$, it holds that $\hat{\mathcal{P}}^{a,b}$ is not an ε -close labelling of $\mathcal{P}^{a,b}$.

Theorem 1. *If CD-GBS is given access to Q_ℓ for a (m, n) -polytope partition, it computes an ε -close labelling of \mathcal{P} using at most $(\prod_{i=1}^m ((\binom{n+i}{i} + 2n)) 2^{2m^2} \log^m \left(\frac{170nm^{5/2}}{\varepsilon} \right))$ membership queries. For constant m this constitutes $O(n^{m^2} \log^m(\frac{n}{\varepsilon})) = \text{poly}(n, \log(\frac{1}{\varepsilon}))$ queries¹.*

Proof. We first prove that CD-GBS indeed computes an ε -close labelling when given access to a valid Q_ℓ by inducting on m . It is straightforward to see that in the case $m = 1$, if CD-GBS is given access to a valid Q_ℓ for a $(1, n)$ polytope partition (a partition of the unit interval into connected subintervals), then it simply performs binary search on the interval $[0, 1] \cong \Delta^1$.

As for the inductive step, for $k = \lceil \log(2/\varepsilon) \rceil$, any two contiguous points of \mathcal{D}^k are less than $\varepsilon/2$ away from each other. For now suppose that every recursive call to CD-GBS was along a non-degenerate cross section \mathcal{P}^t . From the inductive assumption, this means that CD-GBS computes an $\varepsilon/2$ -close labellings of those cross-sections, using the triangle inequality, we know that $\hat{\mathcal{P}}$ is an ε -close labelling of \mathcal{P} .

We note however that there is no guarantee for what a recursive call to CD-GBS does on a degenerate cross section $\hat{\mathcal{P}}^t$. For this reason, it could be the case that at the end of the loop over \mathcal{D}^i , $\hat{\mathcal{P}}$ is not an ε -close labelling. This can only happen if there is some $t \in C_{\mathcal{P}}^\alpha \cap \mathcal{D}^k$ which is an uncovered critical coordinate.

If t is an uncovered critical coordinate we can rectify the situation. If we find a $z \in B_{\varepsilon/2}$ that is not a critical coordinate, then \mathcal{P}^z is non-degenerate and computing CD-GBS along the cross-section gives us an $\frac{\varepsilon}{2}$ -close labelling of \mathcal{P}^z . Using the triangle inequality, we see that this in turn removes t from the set of uncovered critical coordinates, and we say that t is “fixed”. Thus the final while loop of the algorithm eliminates the set of uncovered critical coordinates so that $\hat{\mathcal{P}}$ is indeed an ε -close labelling.

It thus remains to show that the final while loop terminates. However, there are at most $|C_{\mathcal{P}}^\alpha|$ uncovered critical coordinates, and over the course of fixing all uncovered critical coordinates, there are at most $|C_{\mathcal{P}}^\alpha|$ bad guesses for $z \in B_{\varepsilon/2}(x)$ where \mathcal{P}^z is degenerate. Therefore the final while loop makes at most $2|C_{\mathcal{P}}^\alpha|$ invocations to CD-GBS along cross-sections. This concludes the proof of correctness for CD-GBS.

Let us bound the total query usage of CD-GBS. For all values of k in the first for loop, we know from Corollary 3.1 that since Q_ℓ is a valid lexicographic oracle for \mathcal{P} , that the number of uncovered I_x^k will not exceed $2((\binom{n+m}{m} + 2n))$, and since CD-GBS is called once per uncovered interval, it follows that for each k there at most $2((\binom{n+m}{m} + 2n))$ recursive calls to CD-GBS. Furthermore, since Q_ℓ is a valid lexicographic oracle for \mathcal{P} , it will also never be the case that $\exists i, j \in [n], z \in \Delta^m$ such that $\dim(\hat{P}_i) = m$ and $z \in \text{int}(\hat{P}_i)$.

In the worst case, k loops from 1 to $\lceil \log(2/\varepsilon) \rceil$ and makes an extra $2|C_{\mathcal{P}}^\alpha|$ recursive calls to CD-GBS to fix all uncovered critical coordinates. In total if we let $T(m, n, \varepsilon)$ denote the query cost of running

¹ CD-GBS runs in polynomial time for constant m . The time-intensive operation consists of identifying uncovered intervals, but since the dimension of the ambient simplex is constant, each empirical polytope \hat{P}_i has at most a constant number of bounding hyperplanes. These hyperplanes can each be extruded by ε , and checking whether there exists a point outside all these extrusions can be done in time polynomial in n via brute force. In fact, all other algorithms in this paper have efficient runtimes (in their relevant parameters) due to similar reasoning.

CD-GBS on a valid lexicographic oracle, we get the following recursion:

$$T(m, n, \varepsilon) \leq 2|C_{\mathcal{P}}^\alpha| \log\left(\frac{2}{\varepsilon}\right) T\left(m-1, n, \frac{\varepsilon^2}{85nm^{5/2}}\right) + 2|C_{\mathcal{P}}^\alpha|.$$

In order to make this more amenable, we define $f(m) = \left(\binom{n+m}{m} + 2n\right)$ and use Lemma 7 to bound this expression as follows:

$$T(m, n, \varepsilon) \leq 3f(m) \log\left(\frac{2}{\varepsilon}\right) T\left(m-1, n, \frac{\varepsilon^2}{85nm^{5/2}}\right).$$

Furthermore, from the fact that the base case is binary search, we know $T(1, n, \varepsilon) \leq n \log\left(\frac{2}{\varepsilon}\right)$.

To unpack the recursion. Let us define $\varepsilon_0 = \varepsilon$ and $\varepsilon_{k+1} = \frac{\varepsilon_k^2}{85n(m-k)^{5/2}}$ for $k = 1, \dots, m-1$. With this in hand, we can unroll the recursion to obtain:

$$T(m, n, \varepsilon) \leq \left(3^{m-1} \prod_{i=1}^{m-1} f(i)\right) \left(\prod_{k=1}^{m-1} \log\left(\frac{2}{\varepsilon_k}\right)\right).$$

Since each $\varepsilon_{k+1} < \varepsilon_k$, we can upper bound the right-hand product by bounding each term with ε_{m-1} . If we first solve for this value, we obtain:

$$\varepsilon_{m-1} = \frac{\varepsilon^{2^{m-1}}}{\prod_{j=1}^{m-1} (85nj^{5/2})^{2^j}} \geq \frac{\varepsilon^{2^{m-1}}}{\prod_{j=1}^{m-1} (85nm^{5/2})^{2^j}} \geq \left(\frac{\varepsilon}{85nm^{5/2}}\right)^{2^m}.$$

In the first inequality we bounded the denominator product in the base by $j \leq m$, as for the second inequality, we evaluated the geometric series in 2 for the exponent to bound the exponent by 2^m . With this in hand we obtain the desired bounds:

$$T(m, n, \varepsilon) \leq 3^m 2^{m^2} \prod_{i=1}^m f(i) \log^m\left(\frac{170nm^{5/2}}{\varepsilon}\right) \leq \left(\prod_{i=1}^m \left(\binom{n+i}{i} + 2n\right)\right) 2^{2m^2} \log^m\left(\frac{170nm^{5/2}}{\varepsilon}\right).$$

For large enough n , every term in $\prod_{i=1}^m \left(\binom{n+i}{i} + 2n\right)$ is bounded by $(n+m)^m + 2n$. It follows that this product is $O(n^{m^2})$. For constant m , this constitutes $O(n^{m^2} \log^m(\frac{n}{\varepsilon})) = \text{poly}(n, \log(\frac{1}{\varepsilon}))$ queries.

The previous results show that for constant dimension, m , CD-GBS is query efficient in n and $\frac{1}{\varepsilon}$. In the following section we use this algorithm as a building block to construct a method for computing efficient ε -close labellings when the number of regions, n , is held constant instead.

4 Constant-Region Generalised Binary Search

The intuition behind the algorithm lies in the fact that if m is much greater than n , then any vertex of a given P_i cannot lie in the interior of the ambient simplex Δ^m . This is since a vertex in Δ^m must consist of the intersections of at least m half-spaces, all of which cannot arise from adjacencies between different P_i . Not only do all vertices lie on the boundary of Δ^m , but one can show that they are all contained in faces of the boundary of Δ^m that have dimension $O(n^2)$ which is presumed to be constant. The number of such faces in the boundary of Δ^m is thus polynomial in m , and if we could compute 0-close labellings of these faces we could take convex combinations and recover a 0-close labelling of the entire polytope partition. We will demonstrate that for an appropriate value of ε' , if we compute ε' -close labellings of such faces in the boundary, we can recover an ε -close labelling of the entire polytope partition over all of Δ^m by taking convex combinations. CR-GBS computes the necessary ε' -close labellings of lower dimensional faces by using CD-GBS as a subroutine.

Algorithm 1 CD-GBS(m, n, ε, Q)

Input: $m \geq 0$, $n, \varepsilon > 0$, query access to function $Q : \Delta^m \rightarrow [n]$.

Output: $\hat{\mathcal{P}}$: an ε -close labelling of \mathcal{P} .

```
if  $m = 0$  then
  Query  $Q(0)$ 
else
   $\hat{\mathcal{P}}^0 \leftarrow f_0^{-1} \left( \text{CD-GBS} \left( m-1, n, \frac{\varepsilon^2}{85nm^{5/2}}, Q \circ f_0^{-1} \right) \right)$ ,  $\hat{\mathcal{P}}^1 \leftarrow Q(e_1)$ .
  for  $k = 1$  to  $\lceil \log(2/\varepsilon) \rceil$  do
    if Number of uncovered  $I_x^k$  exceeds  $2 \left( \binom{n+m}{m} + 2n \right)$  then
      Halt
    for  $x \in \mathcal{D}^k$  do
      if  $I_x^k$  is uncovered then
         $t \leftarrow \text{midpoint}(I_x)$ 
         $\hat{\mathcal{P}}^t \leftarrow f_t^{-1} \left( \text{CD-GBS} \left( m-1, n, \frac{\varepsilon^2}{85(1-t)nm^{5/2}}, Q \circ f_t^{-1} \right) \right)$ 
        Recompute  $\hat{\mathcal{P}}$  by taking convex hulls of labels
      if  $\exists i, j \in [n]$  such that  $\text{int}(\hat{P}_i) \cap \hat{P}_j \neq \emptyset$  or  $\hat{\mathcal{P}}$  is an  $\varepsilon$ -close labelling then
        Halt
  while  $\exists x \in [0, 1]$  an uncovered critical point do
     $t \leftarrow z$  for arbitrary  $z \in B_{\varepsilon/2}(x)$ 
     $\hat{\mathcal{P}}^t \leftarrow f_t^{-1} \left( \text{CD-GBS} \left( m-1, n, \frac{\varepsilon^2}{85(1-t)nm^{5/2}}, Q \circ f_t^{-1} \right) \right)$ 
    Recompute  $\hat{\mathcal{P}}$  by taking convex hulls of labels
  return  $\hat{\mathcal{P}}$ 
```

Necessary Machinery for CR-GBS: Suppose that \mathcal{P} is an (m, n) -polytope partition with $m > \binom{n}{2}$. Furthermore, let $k = \binom{n}{2}$ and let $\partial_k(\Delta^m)$ denote all k -dimensional faces of Δ^m . For each face F , let \mathcal{P}_F be the restriction of \mathcal{P} to F . If F contains the origin, then it is an isometric embedding of Δ^k in Δ^m , so we let ϕ_F be a canonical isomorphism from F to Δ^k . If F does not contain the origin, then let $v_F \in F$ be the vertex of lowest index, i.e. $v_F = e_i$ where $i = \arg\min\{j \mid e_j \in F\}$. In this case, we let ϕ_F be a canonical linear isomorphism from F to Δ^k that maps v_F to the origin and other vertices of F to other vertices of Δ^k .

As mentioned previously, computing empirical labellings of every face in $\partial_k(\Delta^m)$ via CD-GBS will be enough to compute an empirical labelling for \mathcal{P} . The only issue however, is that CD-GBS is only guaranteed to return an ε -close labelling if given access to a valid lexicographic membership oracle for a polytope partition. For an arbitrary polytope partition, $\phi_F(\mathcal{P}_F)$ is not necessarily a (k, n) -polytope partition for all $F \in \partial_k(\Delta^m)$. For this reason, we slightly refine our notion of polytope partition.

Definition 11. Suppose that \mathcal{P} is an (m, n) -polytope partition such that for all $0 \leq k \leq m$ and $F \in \partial_k(\Delta^m)$, $\phi_F(\mathcal{P}_F)$ is a (k, n) -polytope partition. Then we say \mathcal{P} is a proper polytope partition.

Specification of CR-GBS and Query Usage: For $F \in \partial_k(\Delta^m)$, we let ϕ_F denote the canonical isomorphism from F to Δ^k . This isomorphism is different depending on whether F is axis-aligned or not. Furthermore, we let $\hat{\mathcal{P}}_F$ empirical labelling returned by CD-GBS on $F \in \partial_k(\Delta^m)$. The full proof of correctness and query usage of Algorithm 2 can be found in the online version of the paper.

Algorithm 2 CR-GBS(m, n, ε, Q)

Input: $m, n, \varepsilon > 0$, query access to membership oracle Q for (m, n) -polytope partition \mathcal{P} .

Output: $\hat{\mathcal{P}}$: an ε -close labelling of \mathcal{P} .

```
 $k \leftarrow \binom{n}{2}$ 
for  $F \in \partial_k(\Delta^m)$  do
   $\hat{\mathcal{P}}^F \leftarrow \phi_F^{-1} \left( \text{CD-GBS} \left( k, n, \frac{3\varepsilon}{100n^2\sqrt{k+1}(m+1)^{5/2}}, Q \circ \phi_F^{-1} \right) \right)$ .
 $\hat{\mathcal{P}} \leftarrow \text{Conv}_F(\hat{\mathcal{P}}_F)$ 
return  $\hat{\mathcal{P}}$ 
```

Theorem 2. Let \mathcal{P} be a proper (m, n) -polytope partition where n is constant and $m > k = \binom{n}{2}$. CR-GBS computes an ε -close labelling of \mathcal{P} and uses $O\left(m^k \log^k\left(\frac{m}{\varepsilon}\right)\right) = \text{poly}(m, \log\left(\frac{1}{\varepsilon}\right))$ queries.

5 Upper Envelope Polytope Partitions

We have focused completely on the lexicographic query oracle Q_ℓ , creating algorithms CD-GBS and CR-GBS that compute ε -close labellings of (m, n) -polytope partitions when given access to Q_ℓ . If these algorithms are given access to an adversarial oracle Q_A however, they may fail.

To see why CD-GBS may fail under Q_A we recall that the algorithm recursively computes ε -close labellings of cross-sections \mathcal{P}^t for different values of $t \in [0, 1]$. If ever CD-GBS is called on a degenerate cross-section \mathcal{P}^t , it has conditions to either tell that it is being called on a degenerate cross-section (when it notices that there exist $i, j \in [n]$ and $z \in \Delta^m$ such that $z \in \text{int}(\hat{P}_i) \cap \hat{P}_j$), or in the worst case, prevent it from exceeding its query balance. In both cases however, the algorithm returns a valid empirical labelling, i.e., $\hat{P} = \{\hat{P}_i\}_{i=1}^n$ such that $\hat{P}_i \subseteq P_i$.

When an adversarial oracle is used however, we may see $i, j \in [n]$ and $z \in \Delta^m$ such that $z \in \text{int}(\hat{P}_i) \cap \hat{P}_j$. Indeed this can occur if $P_i = P_j$ and both are full-dimensional. The natural solution seems to merge P_i and P_j (since the second condition of the definition of polytope partitions tells us that $P_i = P_j$ in this case). The main problem however, is that there is no way of telling when the condition above is an artifice of the adversarial oracle, or simply due to the fact that \mathcal{P}^t is degenerate. If we blindly merge labels, we may in fact be performing an incorrect merge on a degenerate cross-section! This of course may return inconsistent polytope partitions.

Since the key problem is the existence of degenerate cross-sections, we consider a slightly stronger variant of polytope partitions with the key property that cross-sections are never degenerate. Furthermore, this special type of polytope partition is expressive enough for our game theoretic applications, and best of all, it allows us to prove results in the adversarial query oracle model.

Definition 12 (Upper Envelope Polytope Partition). Suppose that $A \in \mathbb{R}^{n \times m}$ is an $n \times m$ real-valued matrix and that $b \in \mathbb{R}^n$. Let $P_i = y \in \Delta^m$ such that $(Ay + b)_i \geq (Ay + b)_j$ for all $j \neq i$. We denote the collection $\mathcal{P}(A, b) = P_1, \dots, P_n$, as the upper envelope polytope partition (UEPP) arising from (A, b) .

It is straightforward to see that for any (A, b) , $\mathcal{P}(A, b)$ is itself an (m, n) -polytope partition. Crucially however, it satisfies more properties than the previous definition of polytope partitions. The full proof of the following lemma can be found in the online version of the paper.

Lemma 9. Suppose that A is an $n \times m$ real valued matrix and that $b \in \mathbb{R}^n$. Then $\mathcal{P}(A, b) = \{P_1, \dots, P_n\}$ has the following properties:

- For any $x \in [0, 1]$ let f_x be the canonical affine transformation that maps $(\Delta^m)^x$ to Δ^{m-1} . There exists an $n \times (m-1)$ real matrix A^x and $b^x \in \mathbb{R}^n$ such that $\mathcal{P}(A^x, b^x) = f_x(\mathcal{P}(A, b)^x)$.
- $\mathcal{P}(A, b)$ is a proper polytope partition (Definition 11).
- If $A_{i,\bullet} = A_{j,\bullet}$ and $b_i = b_j$ then $P_i = P_j$. Conversely if P_i is of full affine dimension and $\text{relint}(P_i) \cap P_j \neq \emptyset$, then $A_{i,\bullet} = A_{j,\bullet}$ and $b_i = b_j$; consequently, $P_i = P_j$.
- Let $a_1, \dots, a_k \in \mathbb{R}$ be such that $\sum_{i=1}^k a_i < 1$ with $k < m$. Let $H = \{(z_1, \dots, z_m) \in \Delta^m \mid z_i = a_i, i = 1, \dots, k\}$ where H has codimension k . If $x_1, \dots, x_{m-k} \in \Delta^m$ are affinely independent points of $P_i \cap H$ and $y \in \text{Conv}(x_1, \dots, x_{m-k})$ belongs to P_j , then P_i and P_j coincide in H .

5.1 Adversarial CD-GBS

Suppose that \mathcal{P} is an UEPP. Since it is also a proper (m, n) -polytope partition, it inherits all the properties from before. Along with Lemma 9 we have the necessary tools to show that Algorithm 3 is a query efficient way of computing ε -close labellings of \mathcal{P} with an adversarial query oracle. In the specification of CD-GBS, we use identical terms and notation from Algorithm 1. The full proof of correctness and query usage of Algorithm 3 can be found in the online version of the paper.

Theorem 3. If Adversarial CD-GBS is given access to an adversarial query oracle Q_A of an (m, n) -polytope partition based on a UEPP, it computes an ε -close labelling of \mathcal{P} using at most

$\left(\prod_{i=1}^m \left(\binom{n+i}{i} + 2n\right)\right) 2^{2m^2} \log^m \left(\frac{170nm^{5/2}}{\varepsilon}\right)$ membership queries. For constant m this constitutes $O(n^{m^2} \log^m(\frac{n}{\varepsilon})) = \text{poly}(n, \log(\frac{1}{\varepsilon}))$ queries.

Algorithm 3 Adversarial CD-GBS(m, n, ε, Q_A)

Input: $m \geq 0$, $n, \varepsilon > 0$, query access to oracle $Q_A : \Delta^m \rightarrow [n]$.

Require: Recursive calls to CD-GBS($m-1, n, \frac{\varepsilon^2}{85(1-x)nm^{5/2}}, Q_A \circ f_x^{-1}$).

Output: ε -close labelling of \mathcal{P} .

```
if  $m = 0$  then
  Query  $Q_A(0)$ 
else
   $\hat{\mathcal{P}}^0 \leftarrow f_0^{-1} \left( \text{CD-GBS} \left( m-1, n, \frac{\varepsilon^2}{85nm^{5/2}}, Q_A \circ f_0^{-1} \right) \right)$ 
   $\hat{\mathcal{P}}^1 \leftarrow Q(e_1)$ 
  for  $k = 1$  to  $\lceil \log(2/\varepsilon) \rceil$  do
    for  $x \in \mathcal{D}^k$  do
      if  $I_x^k$  is uncovered then
         $t \leftarrow \text{midpoint}(I_x)$ 
         $\hat{\mathcal{P}}^t \leftarrow f_t^{-1} \left( \text{CD-GBS} \left( m-1, n, \frac{\varepsilon^2}{85(1-t)nm^{5/2}}, Q_A \circ f_t^{-1} \right) \right)$ 
      Recompute  $\hat{\mathcal{P}}$  by taking convex hulls of labels
    while  $\exists i, j \in [n], z \in \Delta^m$  such that  $\dim(\hat{P}_i) = m$  and  $z \in \text{int}(\hat{P}_i)$  do
      Merge label  $i$  with label  $j$ 
      Recompute  $\hat{\mathcal{P}}$  by taking convex hulls of labels
    if  $\hat{\mathcal{P}}$  is an  $\varepsilon$ -close labelling then
      Break
  return  $\hat{\mathcal{P}}$ 
```

5.2 Adversarial CR-GBS

In this section we formalize an adversarial variant of CR-GBS. We note that most of the notation is identical to lexicographic CR-GBS. As before, the full proof of correctness and query usage of Algorithm 4 can be found in the online version of the paper.

Algorithm 4 CR-GBS($m, n, \varepsilon, \mathcal{P}$)

Input: $m, n, \varepsilon > 0$, query access to Q_A for (m, n) -polytope partition \mathcal{P} .

Output: ε -close labelling of \mathcal{P} .

```
 $k \leftarrow \binom{n}{2}$ 
for  $F \in \partial(\Delta^m)^k$  do
   $\hat{\mathcal{P}}^F \leftarrow \phi_F^{-1} \left( \text{CD-GBS} \left( k, n, \frac{3\varepsilon}{100n^2\sqrt{k+1}(m+1)^{5/2}}, Q \circ \phi_F^{-1} \right) \right)$ 
   $\hat{\mathcal{P}} \leftarrow \text{Conv}_F(\hat{\mathcal{P}}^F)$ 
while  $\exists i, j \in [n], z \in \Delta^m$  such that  $\dim(\hat{P}_i) = m$  and  $z \in \text{int}(\hat{P}_i)$  do
  Merge label  $i$  with label  $j$ 
  Recompute convex hulls of labels
return  $\hat{Q}$ 
```

Theorem 4. Let \mathcal{P} be an (m, n) -polytope partition where n is constant. Furthermore, let $k = \binom{n}{2}$. Adversarial CR-GBS computes an ε -close labelling of \mathcal{P} and uses $O\left(m^k \log^k\left(\frac{m}{\varepsilon}\right)\right) = \text{poly}(m, \log\left(\frac{1}{\varepsilon}\right))$ queries.

6 Applications to Game Theory

Let $G = (A, B)$ be an $m \times n$ bi-matrix game where $A, B \in [0, 1]^{m \times n}$ are row and column player payoff matrices respectively with payoffs normalised to $[0, 1]$. The set of row player pure strategies is $[m] = \{1, \dots, m\}$ and that of the column player pure strategies is $[n] = \{1, \dots, n\}$. The set of all row player mixed strategies can be identified with $\Delta^{m-1} = \{x \in \mathbb{R}^{m-1} \mid \sum_{i=1}^{m-1} x_i \leq 1 \text{ and } x_i \geq 0\}$. Similarly, column player mixed strategies are identified with Δ^{n-1} .

Definition 13 (Utility Functions). Let $u \in \Delta^{m-1}$, and $v \in \Delta^{n-1}$ be row and column player mixed strategies and let $u' = (1 - \sum u_i, u_1, \dots, u_{n-1})$ and $v' = (1 - \sum v_i, v_1, \dots, v_{n-1})$. For strategy profile (u, v) , row player utility is $U_r(u, v) = u^T A v$ and column player utility is $U_c(u, v) = u^T B v$.

It will be useful to have shorthand for the following functions: $U_r^i(y) = U_r(e_i, y)$ as the row player utility for playing pure strategy i , and $E_R(y) = \max_{i \in [m]} U_r^i(y)$ as the maximal utility the row player can achieve against mixed strategies. In an identical fashion we can define U_c^j and E_C as the column player utility in playing strategy j and the maximal column player utility. With this notation in hand, we can define the best response oracles our algorithms will have access to.

Definition 14 (Best Response Query Oracles). Any bimatrix game has the following best response query oracles:

- Strong query oracles: for the column player, $BR_s^C(u) = \{j \in [n] \mid U_c^j(u) = E_C(u)\}$ and for the row player, $BR_s^R(v) = \{i \in [m] \mid U_r^i(v) = E_R(v)\}$
- Lexicographic query oracles: for the column player, $BR_\ell^C(u) = \min BR_s^C(u)$ and for the row player, $BR_\ell^R(v) = \min BR_s^R(v)$
- Adversarial query oracles: for the column player, any function BR_A^C such that $BR_A^C(u) \in BR_s^C(u)$ and for the row player, any function such that $BR_A^R(v) \in BR_s^R(v)$

Definition 15 (Nash Equilibrium). Suppose that u and v are row and column player strategies respectively. We say that the pair (u, v) is a Nash Equilibrium (NE) if for all $u' \in \Delta^{m-1}$ and $v' \in \Delta^{n-1}$: $U_r(u, v) \geq U_r(u', v)$ and $U_c(u, v) \geq U_c(u, v')$.

With utility queries the complexity of an exact Nash equilibrium is finite: we can exhaustively query the game. This is not the case for best response queries. Therefore, the relaxation we study is that of approximate well-supported Nash equilibria. Before proceeding, we say that a row player mixed strategy $u \in \Delta^m$ is an ε best response against a column player mixed strategy $v \in \Delta^n$ if $U_r(u, v) \geq U_r(u', v) - \varepsilon$ for all $u' \in \Delta^m$. An identical notion holds for the column player.

Definition 16 (ε -Well-Supported Nash Equilibrium). Suppose that u and v are row and column player strategies respectively. We say that the pair (u, v) is an ε -well-supported Nash equilibrium (ε -WSNE) if and only if u is supported by ε -best responses to v and vice versa.

Definition 16 mentions approximate best responses, yet we only have access to the best response oracle in our model. To resolve this, we bound how much utilities can deviate between “close” mixed strategy profiles. The full proof of the following lemma can be found in the online version of the paper.

Lemma 10. Fix $\varepsilon > 0$ and let $\delta_C = \frac{\varepsilon}{2\sqrt{m-1}}$. Suppose that $u \in \Delta^{m-1}$ is a row player mixed strategy with $c_j \in BR_s^C(u)$. For any u' such that $\|u - u'\|_2 \leq \delta_C$, if $c_i \in BR_s^C(u')$, then $|U_c^i(u) - U_c^j(u)| \leq \varepsilon$. In other words, c_i is an ε -best response to u . Similarly, let $\delta_R = \frac{\varepsilon}{2\sqrt{n-1}}$. Suppose that $v \in \Delta^{n-1}$ is a column player mixed strategy with $r_j \in BR_s^R(u)$. For any v' such that $\|v - v'\|_2 \leq \delta_R$, if $r_i \in BR_s^R(v')$, then $|U_r^i(v) - U_r^j(v)| \leq \varepsilon$. In other words, r_i is an ε -best response to v .

We now prove the connection between computing ε -close labellings of polytope partitions and computing ε -WSNE for bimatrix games using best response queries.

Definition 17 (Best Response Sets). Let $G = (A, B)$ be a bimatrix game. We define column best response sets as the collection of $C_i = \{x \in \Delta^{m-1} \mid BR_s^C(x) \ni c_i\}$. Similarly we define row player best response sets as the collection of $R_j = \{y \in \Delta^{n-1} \mid BR_s^R(y) \ni r_j\}$. We denote the collections by $\mathcal{C} = \{C_i\}_{i=1}^n$ and $\mathcal{R} = \{R_j\}_{j=1}^m$.

Since utilities are affine functions, it is immediately clear that \mathcal{C} and \mathcal{R} are upper envelope polytope partitions. Best response oracles play the same role as membership oracles, Q , from before. Since adversarial oracles are a valid lexicographic oracle, we focus on using adversarial best response oracles. With our language of empirical labellings we can now define the following important concept, but first we clarify some notation: $d(x, S)$ denotes the infimum distance of a point, x to a set S .

Definition 18 (Voronoi Best Response Sets). Suppose that $\widehat{\mathcal{C}} = \{\widehat{C}_i\}$ and $\widehat{\mathcal{R}} = \{\widehat{R}_j\}$ are empirical labellings of \mathcal{C} and \mathcal{R} as in Definition 6. The Voronoi Best Response Sets of the row and column player are $VR_j = \{y \in \Delta^{n-1} \mid \operatorname{argmin}_j d(y, \widehat{R}_j) = r_j\}$ and $VC_i = \{x \in \Delta^{m-1} \mid \operatorname{argmin}_i d(x, \widehat{C}_i) = c_i\}$, defined for any $j \in [m]$ and $i \in [n]$. Furthermore, we let $V^R(v) = \{i \mid VR_i \ni v\}$ and $V^C(u) = \{j \mid VC_j \ni u\}$ be the row and column player Voronoi Best Responses.

Lemma 11. Suppose that $\widehat{\mathcal{C}}$ is a $\frac{\varepsilon}{2\sqrt{m-1}}$ -close labelling and $\widehat{\mathcal{R}}$ is a $\frac{\varepsilon}{2\sqrt{n-1}}$ -close labelling. Then Voronoi best responses are ε best-responses in G .

Lemma 11 (with a full proof in the online version of the paper), tells us precisely that Voronoi best response sets allow us to extend partial information from empirical labellings to approximate best response information across the entire domains Δ^m and Δ^n . This hints at the fact that Voronoi best response sets hold enough information to compute ε -WSNE. In fact we can prove this in the same way as Nash's theorem: via Kakutani's fixed point theorem. In order to do so, we define a Voronoi best response correspondence (which as we have shown before is an approximate best response correspondence), and show that it satisfies the properties of Kakutani's fixed point theorem. The guaranteed fixed point of this correspondence will in turn be an ε -WSNE.

Definition 19 (Voronoi Approximate Best Response Correspondence). For a given mixed strategy profile of both the row and column player, $(u, v) \in \Delta^{m-1} \times \Delta^{n-1}$, we define $B^*(u, v)$ to be the set of all possible mixtures over Voronoi best response profiles both players may have to the other player's strategy. $B^* : \Delta^{m-1} \times \Delta^{n-1} \rightarrow \mathcal{P}(\Delta^{m-1} \times \Delta^{n-1})$ is defined as follows: $B^*(u, v) = (\operatorname{conv}(V^R(v)), \operatorname{conv}(V^C(u))) \subseteq \Delta^{m-1} \times \Delta^{n-1}$.

Theorem 5 (Kakutani's Fixed Point Theorem [18]). Let A be a non-empty, compact and convex subset of \mathbb{R}^n . Let $f : A \rightarrow \mathcal{P}(A)$ be a set-valued function on A with a closed graph and the property that $f(x)$ is non-empty and convex for all $x \in A$. Then f has a fixed point.

Theorem 6. B^* satisfies all the conditions of Kakutani's fixed point Theorem, and hence there exists a strategy profile (u^*, v^*) such that $(u^*, v^*) \in B^*(u^*, v^*)$. In particular, if the Voronoi best responses for B^* arise from $\widehat{\mathcal{C}}$, a $\frac{\varepsilon}{2\sqrt{m-1}}$ -close labelling and $\widehat{\mathcal{R}}$, a $\frac{\varepsilon}{2\sqrt{n-1}}$ -close labelling, then this in turn implies that (u^*, v^*) is an ε -WSNE of G .

Proof. We need to prove the following conditions for Kakutani's fixed point Theorem:

- B^* has a compact and convex domain.
- $B^*(u, v)$ is non-empty and convex for all $(u, v) \in \Delta^{m-1} \times \Delta^{n-1}$.
- (Graph Closedness) Suppose that $\{\sigma_n\}$ and $\{\sigma'_n\}$ are sequences in $\Delta^{m-1} \times \Delta^{n-1}$ that converge to σ and σ' respectively. Furthermore, suppose that $\sigma'_n \in B^*(\sigma_n)$ for all n . Then $\sigma' \in B^*(\sigma)$.

For the first item, the domain of B^* is $\Delta^{m-1} \times \Delta^{n-1}$ which clearly satisfies the desired condition. As for the second item, from the definition of B^* the image of any (u, v) consists of convex combinations of Voronoi best responses, which are defined for all (u, v) (thus satisfying non-emptiness), and since they are convex combinations, they are convex subsets of $\Delta^{m-1} \times \Delta^{n-1}$.

Finally for the third item, let us consider such a sequence where $\sigma_n = (u_n, v_n)$, $\sigma = (u, v)$, $\sigma'_n = (u'_n, v'_n)$, and $\sigma' = (u', v')$. To show the claim, it suffices to consider the sequences $\{u_n\}$ and $\{v'_n\}$ with respective limits u and v' , and show that $v' \in \operatorname{conv}(V^C(u))$. To show this however, it suffices to use the fact that Voronoi best response sets are closed. Suppose that u has a certain set $S \subset [n]$ of Voronoi best responses. Then there exists a constant $\mu > 0$ such that $B_\mu(u) \cap VC_i \neq \emptyset$ if and only if $i \in S$; namely the μ neighbourhood around u only intersects Voronoi best response sets from u 's Voronoi best responses.

To explicitly construct such a μ , let $D_i = d(u, \widehat{C}_i)$ to be the distance between u and the empirical best response set \widehat{C}_i . This means that $S = \operatorname{argmin}_i D_i$, so let us define $D = \min_i D_i$ and $\mu = \frac{\min_{j \notin S} D_j - D}{3}$ which is positive due to the fact that there are finitely many empirical best response sets. Now suppose that $x \in B_\mu(u)$, then for any $j \notin S$ we have $d(u, \widehat{C}_j) \leq d(u, x) + d(x, \widehat{C}_j)$ by the triangle inequality, which rearranging gives us: $d(x, \widehat{C}_j) \geq d(u, \widehat{C}_j) - d(u, x) \geq (D + 3\mu) - \mu = D + 2\mu$. On the other hand, for any $i \in S$, $d(x, \widehat{C}_i) \leq d(x, u) + d(u, \widehat{C}_i) \leq D + \mu$. It thus follows that x can only have Voronoi best responses from S .

From the fact that $u_n \rightarrow u$, then for some $N > 0$, if $n > N$ then $u_n \in B_\mu(u)$. This in turn means that $v_n \in \text{conv}(S)$ by assumption, which means that $v \in \text{conv}(S)$ as well, which is what we wanted to show. To extend this to σ and σ' , it suffices to repeat the previous argument in each component of the correspondence.

Now that the conditions of Kakutani's fixed point Theorem are satisfied, we know of the existence of an (u^*, v^*) such that $(u^*, v^*) \in B^*(u^*, v^*)$. As in the statement of the Theorem, suppose that Voronoi best responses for B^* arise from an $\frac{\varepsilon}{2\sqrt{m-1}}$ -close labelling of Δ^m and an $\frac{\varepsilon}{2\sqrt{n-1}}$ of Δ^n , then we know that all Voronoi best responses are ε best responses for both players. The conditions of the fixed point amount to saying that both players are playing convex combinations of Voronoi best responses, therefore (u^*, v^*) is an ε -WSNE.

With Theorem 6 in hand and our algorithms for constructing ε -close labellings, we can bound the query complexity of computing an ε -WSNE in general bimatrix games.

Theorem 7. *Suppose that G is an $m \times n$ bimatrix game and let n be constant. We can compute an ε -WSNE using $O(m^{n^2} \log^{n^2}(\frac{m}{\varepsilon})) = \text{poly}(m, \log(\frac{1}{\varepsilon}))$ adversarial best response queries.*

Proof. Suppose that \mathcal{C} and \mathcal{R} are UEPP arising from best-response sets in G . This means that \mathcal{C} is a $(m-1, n)$ -polytope partition and \mathcal{R} is a $(n-1, m)$ -polytope partition. Let $\varepsilon_C = \frac{\varepsilon}{2\sqrt{m-1}}$ and $\varepsilon_R = \frac{\varepsilon}{2\sqrt{n-1}}$. From Theorem 6, we know that computing an ε_C -close labelling of \mathcal{C} and a ε_R -close labelling of \mathcal{R} suffice to compute an ε -WSNE of G . We use adversarial CR-GBS on \mathcal{C} and adversarial CD-GBS on \mathcal{R} .

n is the number of polytopes in the partition \mathcal{C} , which is assumed to be constant. Consequently, Theorem 4 states that computing an ε_C -close labelling of \mathcal{C} using CR-GBS uses $O((m-1)^k \log^k(\frac{m-1}{\varepsilon_C}))$ adversarial queries, where $k = \binom{n}{2}$. Since $k \leq n^2$, this is bounded by $O(m^{n^2} \log^{n^2}(\frac{m}{\varepsilon}))$.

$n-1$ is the dimension of the ambient simplex in the partition \mathcal{R} , which is assumed to be finite. Consequently, Theorem 3 states that computing an ε_R -close labelling of \mathcal{R} using CD-GBS uses $O(m^{(n-1)^2} \log^{n-1}(\frac{1}{\varepsilon}))$ queries. We trivially upper bound this quantity by $O(m^{n^2} \log^{n^2}(\frac{m}{\varepsilon}))$.

The total query usage is thus $O(m^{n^2} \log^{n^2}(\frac{m}{\varepsilon})) = \text{poly}(m, \log(\frac{1}{\varepsilon}))$ as desired.

7 Conclusion and Future Directions

In this paper we introduced the concept of learning ε -close labellings of (m, n) -polytope partitions with membership queries, and derived query efficient algorithms for when either the dimension of the ambient simplex in the polytope partition, m , is held constant, or when the number of polytopes in the partition, n , is held constant. Most importantly, we introduced a novel reduction from computing ε -WSNE with best response queries to this geometric problem, thus allowing us to show that in the best response query model, computing ε -WSNE of a bimatrix game has a finite query complexity. More specifically, for $m \times n$ games with one parameter, say n , constant, the query complexity is polynomial in m and $\log(\frac{1}{\varepsilon})$. Furthermore, in the full online version of the paper, we partially extend our results from bimatrix games to n -player games. Although the underlying geometry in n -player games prevents us from using our results from learning polytope partitions, we are still able to show that querying a fine enough ε -net of the mixed strategy space of all players suffices to compute an ε -WSNE.

As mentioned in the introduction, this geometric framework could be of use in other areas where Lipschitz continuous structures appear over domains with convex partitions. Upon further inspection, it is not difficult to see that polytope partitions do not need to be contained in Δ^m , and in fact our algorithms extend to arbitrary ambient polytopes. Furthermore, it would be of great interest to create algorithms with a better query cost, prove lower bounds with regards to computing ε -close labellings, or simply explore weaker query paradigms, such as noisy membership oracles. In the multiplayer setting, best response sets are no longer polytopes, but rather semi-algebraic sets. It would be of interest to create learning algorithms for ε -close labellings of these more complicated geometric objects, since doing so suffices to compute ε -WSNE.

Bibliography

- [1] Yakov Babichenko. Query complexity of approximate Nash equilibria. *J. ACM*, 63(4):36:1–36:24, 2016. <https://doi.org/10.1145/2908734>. URL <http://doi.acm.org/10.1145/2908734>.
- [2] Yakov Babichenko and Aviad Rubinstein. Communication complexity of approximate nash equilibria. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 878–889. ACM, 2017. ISBN 978-1-4503-4528-6. <https://doi.org/10.1145/3055399.3055407>. URL <http://doi.acm.org/10.1145/3055399.3055407>.
- [3] Yakov Babichenko, Siddharth Barman, and Ron Peretz. Empirical distribution of equilibrium play and its testing application. *Math. Oper. Res.*, 42(1):15–29, 2017. <https://doi.org/10.1287/moor.2016.0794>. URL <https://doi.org/10.1287/moor.2016.0794>.
- [4] Elizabeth Baldwin and Paul Klemperer. Understanding Preferences: “Demand Types”, and the Existence of Equilibrium with Indivisibilities. LSE Research Online Documents on Economics 63198, London School of Economics and Political Science, LSE Library, October 2016. URL <https://ideas.repec.org/p/ehl/lserod/63198.html>.
- [5] George W. Brown. Some notes on computation of game solutions. *RAND corporation report*, page 78, April 1949.
- [6] Nader H. Bshouty, Paul W. Goldberg, Sally A. Goldman, and H. David Mathias. Exact learning of discretized geometric concepts. *SIAM J. Comput.*, 28(2):674–699, 1998. <https://doi.org/10.1137/S0097539794274246>. URL <https://doi.org/10.1137/S0097539794274246>.
- [7] X. Chen, X. Deng, and S-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):14:1–14:57, 2009.
- [8] C. Daskalakis and Q. Pan. A counter-example to Karlin’s strong conjecture for fictitious play. In *Procs. of 55th FOCS*, pages 11–20, 2014.
- [9] C. Daskalakis, P.W. Goldberg, and C.H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
- [10] J. Fearnley, M. Gairing, P.W. Goldberg, and R. Savani. Learning equilibria of games via payoff queries. *J. Mach. Learn. Res.*, 16:1305–1344, 2015.
- [11] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- [12] Paul W. Goldberg and Stephen Kwek. The precision of query points as a resource for learning convex polytopes with membership queries. In Nicolò Cesa-Bianchi and Sally A. Goldman, editors, *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory (COLT 2000)*, pages 225–235. Morgan Kaufmann, 2000. ISBN 1-55860-703-X.
- [13] Paul W. Goldberg and Aaron Roth. Bounds for the query complexity of approximate equilibria. *ACM Trans. Economics and Comput.*, 4(4):24:1–24:25, 2016. <https://doi.org/10.1145/2956582>. URL <http://doi.acm.org/10.1145/2956582>.
- [14] Paul W. Goldberg, Rahul Savani, Troels Bjerre Sørensen, and Carmine Ventre. On the approximation performance of fictitious play in finite games. *Int. J. Game Theory*, 42(4):1059–1083, 2013. <https://doi.org/10.1007/s00182-012-0362-6>. URL <https://doi.org/10.1007/s00182-012-0362-6>.
- [15] P.W. Goldberg and S. Turchetta. Query complexity of approximate equilibria in anonymous games. *Journal of Computer and System Sciences*, 90:80–98, 2017.
- [16] S. Hart and Y. Mansour. How long to equilibrium? the communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior*, 69(1):107–126, 2010.
- [17] Sergiu Hart and Noam Nisan. The query complexity of correlated equilibria. *Games and Economic Behavior*, 2016. ISSN 0899-8256. <https://doi.org/10.1016/j.geb.2016.11.003>. URL <http://www.sciencedirect.com/science/article/pii/S0899825616301452>.
- [18] Shizuo Kakutani. A generalization of Brouwer’s fixed point theorem. *Duke Math. J.*, 8(3):457–459, 09 1941.
- [19] Paul Klemperer. The product-mix auction: a new auction design for differentiated goods. *Journal of the European Economic Association*, 8(2/3):526–536, 2010.

- [20] R.J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Procs. of the 4th ACM-EC*, EC '03, pages 36–41, New York, NY, USA, 2003. ACM.
- [21] John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951. URL <http://www.jstor.org/stable/1969529>.
- [22] Julia Robinson. An iterative method of solving a game. *The Annals of Mathematics*, 54(2):296–301, 1951.
- [23] B. Von Stengel. Leadership with commitment to mixed strategies. *CDAM Research Report*, 2004.