

Chapter 3

Ellipsoid Against Hope

Ioannis Anagnostides, Gabriele Farina, and Brian Hu Zhang*

The techniques we have covered so far based on regret minimization lead to a running time complexity growing polynomially in $1/\epsilon$, where $\epsilon > 0$ is the approximation quality of the equilibrium—be it a Nash equilibrium in zero-sum games or a (coarse) correlated equilibrium in general-sum games. Specifically, if we employ algorithms with (Φ) -regret bounded by \sqrt{T} as a function of the time horizon, we need $1/\epsilon^2$ iterations to reach an equilibrium. Can we do better in the regime where ϵ is, say, exponentially small? The other main approach we have introduced for equilibrium computation is based on linear programming, which has the advantage of finding an *exact* equilibrium. But, as we have already seen, the LP describing (C)CEs has a number of variables that scales *exponentially* with the number of players.

3.1 Ellipsoid against hope

We will now introduce *ellipsoid against hope* (EAH), a polynomial-time algorithm for computing Φ -equilibria even in games with many players and complex strategy sets. It was first introduced by Papadimitriou and Roughgarden [PR08]. In what follows, we mostly follow the generalized version of the algorithm due to Farina and Pipis [FP24] and Zhang, Anagnostides, Tewolde, Berker, Farina, Conitzer and Sandholm [ZAT+25]. To be clear, this algorithm works in the centralized model, and it's not compatible with the framework of online learning, although as we shall see similarities do exist between the two approaches.

Our goal is to compute an ϵ - Φ -*equilibrium* of a multilinear game, that is, a correlated distribution $\mu \in \Delta(\mathcal{X}_1 \times \dots \times \mathcal{X}_n)$ such that for any player $i \in [n]$ and deviation function $\varphi_i \in \Phi_i$ mapping $\mathcal{X}_i \rightarrow \mathcal{X}_i$,

$$\mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} u_i(\mathbf{x}_1, \dots, \mathbf{x}_n) \geq \mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} u_i(\varphi_i(\mathbf{x}_i), \mathbf{x}_{-i}) - \epsilon. \quad (1)$$

From a more abstract point of view, EAH deals with optimization problems of the form¹

$$\text{find } \mu \in \Delta(\mathcal{X}) \text{ such that } \mathbb{E}_{\mathbf{x} \sim \mu} \langle \mathbf{y}, G(\mathbf{x}) \rangle \geq 0 \quad \forall \mathbf{y} \in \mathcal{Y}, \quad (2)$$

*✉ ianagnos@cs.cmu.edu, {gfarina,zhangbh}@mit.edu. These tutorial notes have not undergone formal peer review. We are grateful for any feedback or reports of typos.

¹We caution that \mathcal{X} here corresponds to $\mathcal{X}_1 \times \dots \times \mathcal{X}_n$ in the context of (1), even though we sometimes take \mathcal{X} to be the strategy set of a single player in the context of regret minimization.

where $\mathcal{X} \subseteq \mathbb{R}^d$, $\mathcal{Y} \subseteq \mathbb{R}^k$, and $G : \mathcal{X} \rightarrow \mathbb{R}^k$ is a function that can be evaluated efficiently. The crux in the optimization problem (2) lies in the fact that μ resides in a high-dimensional space; a canonical case to think about is when $\mathcal{X} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ in a normal-form game, so that even describing a distribution μ could require specifying $\prod_{i=1}^n |\mathcal{A}_i| - 1$ coordinates. We assume that \mathcal{Y} , which corresponds to the set of deviations, admits a *separation oracle*; under mild geometric assumptions, it's equivalent to posit merely a *membership oracle*, which returns whether a point $\mathbf{y} \in \mathbb{R}^k$ belongs to \mathcal{Y} or not. A special case is when \mathcal{Y} is a polytope described with a polynomial number of constraints, as is the case for swap regret in normal-form games—each player's set of deviations amounts to the set of stochastic matrices.

The key assumption in the EAH framework [FP24] is the existence of a *good-enough-response (GER)* oracle, which, given any $\mathbf{y} \in \mathcal{Y}$, returns a point $\mathbf{x} \in \mathcal{X}$ such that $\langle \mathbf{y}, G(\mathbf{x}) \rangle \geq 0$.

The EAH algorithm enables solving (2) with just a separation oracle for \mathcal{Y} and a GER oracle. The basic idea is to consider an ε -approximate version of the dual of (2),

$$\text{find } \mathbf{y} \in \mathcal{Y} \text{ such that } \langle \mathbf{y}, G(\mathbf{x}) \rangle \leq -\varepsilon \quad \forall \mathbf{x} \in \mathcal{X}. \quad (3)$$

On account of the fact that a GER oracle exists, (3) is guaranteed to be infeasible. Even so, EAH proceeds by executing the ellipsoid algorithm on that infeasible program—this is where the apt name “ellipsoid against hope” comes from. In every step t of the algorithm, we have a candidate $\mathbf{y}^{(t)} \in \mathcal{Y}$ and use the GER oracle to produce a point $\mathbf{x}^{(t)} \in \mathcal{X}$ that refutes $\mathbf{y}^{(t)}$; in fact, an entire halfspace in \mathbb{R}^k . This goes on until the volume of the ellipsoid has shrunk to a small enough amount. It then follows that

$$\forall \mathbf{y} \in \mathcal{Y} \exists t \in [T] \text{ such that } \langle \mathbf{y}, G(\mathbf{x}^{(t)}) \rangle > -\varepsilon.$$

Thus,

$$\min_{\mathbf{y} \in \mathcal{Y}} \max_{\mu \in \Delta([T])} \sum_{t=1}^T \mu^{(t)} \langle \mathbf{y}, G(\mathbf{x}^{(t)}) \rangle > -\varepsilon.$$

By the minimax theorem, we conclude that

$$\max_{\mu \in \Delta([T])} \min_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T \mu^{(t)} \langle \mathbf{y}, G(\mathbf{x}^{(t)}) \rangle > -\varepsilon. \quad (4)$$

In other words, there is a convex combination of $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}$ that refutes any possible $\mathbf{y} \in \mathcal{Y}$. That convex combination, which is a *certificate of dual infeasibility*, will thus be an approximate solution to (2). The key point is that the resulting zero-sum game in (4) is much smaller than the one we started with, and can be solved with standard techniques.

Theorem 3.1 ([FP24]). Assuming the existence of a separation oracle for \mathcal{Y} and a GER oracle, EAH runs in time $\text{poly}(d, k, \log(1/\varepsilon))$ and returns an ε -approximate solution to (2).

3.2 Application to computing correlated equilibria

Let's now see to apply this algorithm to solve (1). As before, we assume that each Φ_i contains linear functions of the form $\mathbf{x}_i \mapsto \mathbf{M}_i q_i(\mathbf{x}_i) \in \mathcal{X}_i$ for some feature map $q_i : \mathcal{X}_i \rightarrow \mathbb{R}^{k_i}$. For these notes, we will assume that the set of valid matrices \mathbf{M}_i form a convex, compact set $\mathcal{Y}_i \subset \mathbb{R}^{d_i \times k_i}$; this assumption can be relaxed using a similar idea to the semi-separation oracle [ZAT+25]. We will later briefly explain how to relax this assumption. Then a Φ -equilibrium is a distribution μ such that

$$\sum_{i=1}^n \mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \langle \mathbf{I}_i - \mathbf{M}_i, u_i(\mathbf{x}_{-i}) \otimes q(\mathbf{x}_i) \rangle \geq -\epsilon \quad \forall i \in [n], \mathbf{M}_i \in \mathcal{Y}_i, \quad (5)$$

where \mathbf{I}_i is the matrix for which $\varphi_{\mathbf{I}_i}$ is the identity map. This indeed adheres to (2).

What's left is to prove that (5) admits an efficient GER oracle. To do so, let's consider any $i \in [n]$ and deviation $\mathbf{M}_i \in \mathcal{Y}_i$. Taking $\mu_i \in \Delta(\mathcal{X}_i)$ to be an ϵ -approximate expected fixed point of $\varphi_{\mathbf{M}_i}$, we have

$$\mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \langle \mathbf{x}_i - \mathbf{M}_i q(\mathbf{x}_i), u_i(\mathbf{x}_{-i}) \rangle \geq -\epsilon \cdot \|u_i(\mathbf{x}_{-i})\|_2,$$

But this is not immediately good enough. Recall that our algorithm for expected fixed points, from Chapter 2, Theorem 2.6, has $\text{poly}(1/\epsilon)$ runtime. Since our goal in this section is to construct algorithms with $\log(1/\epsilon)$ runtime, we need to do better. Fortunately, there *is* a $\text{poly}(d, \log(1/\epsilon))$ -time algorithm for semi-separation with expected fixed points:

Theorem 3.2 ([ZAT+25]). There is a $\text{poly}(d, \log(1/\epsilon))$ -time algorithm for semi-separation.

Proof Sketch. We use EAH. Given $\varphi : \mathcal{X} \rightarrow \mathbb{R}^d$, consider the EAH problem with \mathcal{Y} set to the unit ℓ_2 -ball, and $G(\mathbf{x}) = \varphi(\mathbf{x}) - \mathbf{x}$. Then, by Theorem 3.1, it suffices to implement a GER oracle. That is, given $\mathbf{y} \in \mathbb{R}^d$, we need to find \mathbf{x} with $\langle \mathbf{y}, \varphi(\mathbf{x}) - \mathbf{x} \rangle \geq 0$. But this is easy: we can simply take $\mathbf{x} \in \arg \min_{\hat{\mathbf{x}} \in \mathcal{X}} \langle \mathbf{y}, \hat{\mathbf{x}} \rangle$. Then either $\varphi(\mathbf{x}) \notin \mathcal{X}$, in which case we output \mathbf{x} as a certificate of infeasibility of φ , or $\varphi(\mathbf{x}) \in \mathcal{X}$, in which case $\langle \mathbf{y}, \varphi(\mathbf{x}) - \mathbf{x} \rangle \geq 0$ by definition of \mathbf{x} . \square

Finally, if $\mu_i \in \Delta(\mathcal{X}_i)$ is an expected fixed point of φ_i for every i , then the product distribution $\mu_1 \times \dots \times \mu_n$ is a GER solution satisfying (5). We therefore have:

Theorem 3.3 ([FP24; ZAT+25]). If for each player $i \in [n]$ in a multilinear game the set of deviations Φ_i admits a separation oracle, there is an algorithm polynomial in $\log(1/\epsilon)$, n , d_1, \dots, d_n , and k_1, \dots, k_n that outputs an ϵ - Φ -equilibrium.

We conclude this chapter with several remarks about the previous theorem.

1. *Nested EAH.* The algorithm that we derive for Theorem 3.3 is a doubly-nested ellipsoid against hope algorithm: the outer EAH algorithm computes the equilibrium itself, and on every iteration, EAH is invoked for each player i to compute an expected fixed point for i .

2. *Necessity of linearity.* The results presented in this document assume, and use, the fact that utilities are multilinear. The multilinearity assumption is used twice:

- a. in the GER oracle, to calculate expected utilities under a product distribution $\mu = \mu_1 \times \dots \times \mu_n$ (which has support that increases exponentially with the number of players n); and
- b. to ensure that an expected fixed point solution actually guarantees separation: for example, if instead we had arbitrary concave utilities $u(\cdot, \mathbf{x}_{-i})$, then we would need the condition

$$\mathbb{E}_{\mathbf{x} \sim \mu} \langle \nabla u(\mathbf{x}), \varphi(\mathbf{x}) - \mathbf{x} \rangle \geq -\epsilon$$

for every convex function u . Expected fixed points do not guarantee this.

The first issue is not circumventable, and is generally a barrier to efficient equilibrium computation in concave games with many players. The second issue is not circumventable in general, as it would imply the ability to compute fixed points of contraction maps. However, it *is* circumventable in the special case where $u(\cdot)$ is *quadratic* concave for each player i ; therefore, it is possible to compute Φ -equilibria in quadratic games with a constant number of players. The required techniques, however, are beyond the scope of this document. We refer the interested reader to our recent preprint on this topic [ADF+26].

3. *Non-separable sets Φ .* So far in this section, we have assumed that the sets \mathcal{Y}_i admit efficient separation oracles. However, this is not a necessary assumption. As it turns out, like we did for regret minimization in Chapter 2, Section 2.3, it is sufficient to take \mathcal{Y}_i to be a *superset* of the valid deviations and rely on the semi-separation oracle to declare when a given matrix $\mathbf{M}_i \in \mathcal{Y}_i$ is invalid. Therefore, Theorem 3.3 also generalizes to the case where \mathcal{Y}_i only admits an efficient *semi*-separation oracle, such as the sets Φ^q discussed there. For more details, we refer the interested reader to our paper [ZAT+25].

Bibliography for this chapter

- [PR08] C. H. Papadimitriou and T. Roughgarden. (2008). Computing correlated equilibria in multi-player games. *Journal of the ACM*.
- [FP24] G. Farina and C. Papis. (2024). Polynomial-Time Computation of Exact Φ -Equilibria in Polyhedral Games. *Neural Information Processing Systems*.
- [ZAT+25] B. H. Zhang, I. Anagnostides, E. Tewolde, R. E. Berker, G. Farina, V. Conitzer and T. Sandholm. (2025). Learning and Computation of Φ -Equilibria at the Frontier of Tractability. *Economics and Computation*.
- [ADF+26] I. Anagnostides, C. Daskalakis, G. Farina, N. Golowich, T. Sandholm and B. H. Zhang. (2026). On the Complexity of Correlated Equilibria Beyond Normal-Form Games. *arXiv preprint arXiv:2605.17665*.