

Algorithms for Minimizing Phi-Regret: Nonlinear Deviations and Ellipsoid



15 888 **Computational Game Solving** (Fall 2026)
Ioannis Anagnostides

Today's lecture

- Are the oracles required by Gordon et al. necessary?
 - Necessity of fixed points
 - The role of *mixed strategies*
 - Can we always minimize regret over the set of deviations?
- Ellipsoid against hope
 - The natural LP for CE has *exponentially* many variables
 - Algorithmic maneuver to handle multi-player games
 - Parallels with Gordon et al.

The importance of nonlinear deviations

- Linear deviations suffice in normal-form games
- But are restrictive more generally
- Finding fixed points of nonlinear functions is as hard as finding Nash equilibria!
- But do we really have to compute fixed points?

The XOR problem

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \mapsto \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \mapsto \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix} \mapsto \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \text{ and } \begin{pmatrix} 0 \\ 0 \end{pmatrix} \mapsto \begin{pmatrix} * \\ * \end{pmatrix}$$

$$\mathbf{M} \left(\frac{1}{2} \begin{pmatrix} 0 \\ 1 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = \frac{1}{2} \mathbf{M} \begin{pmatrix} 0 \\ 1 \end{pmatrix} + \frac{1}{2} \mathbf{M} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

$\neq$

$$\mathbf{M} \left(\frac{1}{2} \begin{pmatrix} 0 \\ 1 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = \frac{1}{2} \mathbf{M} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

Are fixed points necessary for minimizing Phi-regret?

Yes

Theorem. (Hazan and Kale, 2007) An algorithm that minimizes Phi-regret can compute a fixed point of *any* function in Phi.

$$u^{(t)} : \mathbf{x} \mapsto \frac{1}{\|\phi(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}\|_2} \langle \phi(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}, \mathbf{x} - \mathbf{x}^{(t)} \rangle$$

$$\Phi\text{Reg}^{(T)} \geq \sum_{t=1}^T u^{(t)}(\phi(\mathbf{x}^{(t)})) - \sum_{t=1}^T u^{(t)}(\mathbf{x}^{(t)}) = \sum_{t=1}^T \|\phi(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}\|_2 \geq \epsilon T.$$

Are fixed points necessary for minimizing Phi-regret?

Yes

Theorem. (Hazan and Kale, 2007) An algorithm that minimizes Phi-regret can compute a fixed point of *any* function in Phi.

$$u^{(t)} : \mathbf{x} \mapsto \frac{1}{\|\phi(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}\|_2} \langle \phi(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}, \mathbf{x} - \mathbf{x}^{(t)} \rangle$$

$$\Phi\text{Reg}^{(T)} \geq \sum_{t=1}^T u^{(t)}(\phi(\mathbf{x}^{(t)})) - \sum_{t=1}^T u^{(t)}(\mathbf{x}^{(t)}) = \sum_{t=1}^T \|\phi(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}\|_2 \geq \epsilon T.$$




No if the learner is using mixed strategies!

(Despite Kuhn's theorem)

Expected fixed points

We define a relaxation of fixed points:


$$\|\mathbb{E}_{x \sim \mu}[\phi(x) - x]\| \leq \epsilon.$$



Expected fixed points

We define a relaxation of fixed points:

$$\|\mathbb{E}_{x \sim \mu}[\phi(x) - x]\| \leq \epsilon.$$



Expected fixed points are easy to compute!

$$\mathbf{x}_k \leftarrow \phi(\mathbf{x}_{k-1})$$

$$\sum_{k=1}^K (\phi(\mathbf{x}_k) - \mathbf{x}_k) = \phi(\mathbf{x}_K) - \mathbf{x}_1 + \sum_{k=1}^{K-1} (\phi(\mathbf{x}_k) - \mathbf{x}_{k+1}) = \phi(\mathbf{x}_K) - \mathbf{x}_1.$$

Theorem. There is a polynomial-time algorithm for computing expected FPs.

Refining the framework of Gordon et al.

Algorithm 1: A refinement of Gordon et al. [2008] using expected fixed points.

```
1 Input: An external regret minimizer  $\mathfrak{R}_\Phi$  for the set  $\Phi$ 
2 NEXTSTRATEGY():
3   Set  $\phi^{(t)} := \mathfrak{R}_\Phi.\text{NEXTSTRATEGY}()$ ;
4   return  $\Delta(\mathcal{X}) \ni \mu^{(t)} := \text{EXPECTEDFP}(\phi^{(t)})$ ;

5 OBSERVEUTILITY( $\mathbf{u}^{(t)}$ ):
6   Set  $\mathbf{u}_\Phi^{(t)} : \phi \mapsto \mathbb{E}_{\mathbf{x}^{(t)} \sim \mu^{(t)}} \langle \phi(\mathbf{x}^{(t)}), \mathbf{u}^{(t)} \rangle$ ;
7    $\mathfrak{R}_\Phi.\text{OBSERVEUTILITY}(\mathbf{u}_\Phi^{(t)})$ ;
```

$$\Phi\text{Reg}^{(T)} := \max_{\phi \in \Phi} \sum_{t=1}^T \mathbb{E}_{\mathbf{x}^{(t)} \sim \mu^{(t)}} \langle \phi(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}, \mathbf{u}^{(t)} \rangle.$$

Refining the framework of Gordon et al.

Theorem. The Phi-regret of the algorithm is roughly equal to the external regret over the set of deviations.

$$\begin{aligned}\Phi\text{Reg}^{(T)} &= \max_{\phi \in \Phi} \sum_{t=1}^T \mathbb{E}_{\mathbf{x}^{(t)} \sim \mu^{(t)}} \langle \phi(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}, \mathbf{u}^{(t)} \rangle \\ &= \max_{\phi \in \Phi} \sum_{t=1}^T \mathbb{E}_{\mathbf{x}^{(t)} \sim \mu^{(t)}} \langle \phi(\mathbf{x}^{(t)}) - \phi^{(t)}(\mathbf{x}^{(t)}), \mathbf{u}^{(t)} \rangle + \sum_{t=1}^T \mathbb{E}_{\mathbf{x}^{(t)} \sim \mu^{(t)}} \langle \phi^{(t)}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}, \mathbf{u}^{(t)} \rangle \\ &\leq \max_{\phi \in \Phi} \sum_{t=1}^T (u_{\Phi}^{(t)}(\phi) - u_{\Phi}^{(t)}(\phi^{(t)})) + \sum_{t=1}^T \left\| \mathbb{E}_{\mathbf{x}^{(t)} \sim \mu^{(t)}} (\phi^{(t)}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}) \right\| \|\mathbf{u}^{(t)}\|_* \\ &\leq \text{Reg}^{(T)} + \epsilon BT.\end{aligned}$$

Computing correlated equilibria using ellipsoid

- Regret minimization gives an algorithm for computing (C)CEs that is polynomial in $1/\epsilon$. But what if ϵ is very close to zero?
- A linear program returns an exact solution, but the number of variables scales *exponentially* with the number of players
- Can we compute an exact (C)CE even in multi-player games?

$$\mathbb{E}_{(x_1, \dots, x_n) \sim \mu} u_i(\mathbf{x}_1, \dots, \mathbf{x}_n) \geq \mathbb{E}_{(x_1, \dots, x_n) \sim \mu} u_i(\phi_i(\mathbf{x}_i), \mathbf{x}_{-i}) - \epsilon.$$

The ellipsoid against hope algorithm

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

- This is a *zero-sum* game!
- One player—the mediator—picks a correlated distribution, and then each player tries to deviate optimally
- The strategy set of the mediator is massive!

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

- The dual has **polynomially** many variables
- The idea is to apply ellipsoid on the dual
- What do we need?

The ellipsoid against hope algorithm

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

This can be solved under two basic assumptions:

- We can **optimize** over the set of deviations



Optimization over a set

Under mild geometric assumptions, the following are equivalent:

1. Deciding **membership** of a point

2. Separating a point from the set

3. Given a utility vector, find a best response within the set



Regret minimization over that set!

- Necessity follows by considering a static utility
- Sufficiency uses **follow the regularized leader**

The ellipsoid against hope algorithm

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

This can be solved under two basic assumptions:

- We can optimize over the set of deviations
- There is a **good-enough-response** (GER) oracle:

$$\forall \mathbf{y} \exists \mathbf{x} \text{ such that } \langle \mathbf{y}, G(\mathbf{x}) \rangle \geq 0$$



This implies that the dual is infeasible!

Good-enough-response oracle for CE

Do we have a good-enough-response oracle in our problem?

$$\mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \mathbf{u}_i(\mathbf{x}_1, \dots, \mathbf{x}_n) \geq \mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \mathbf{u}_i(\phi_i(\mathbf{x}_i), \mathbf{x}_{-i}) - \epsilon.$$

$$\Longleftrightarrow$$


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

because $\mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \langle \mathbf{x}_i - \mathbf{M}_i \mathbf{x}_i, \mathbf{u}_i(\mathbf{x}_{-i}) \rangle = \mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \langle \mathbf{I} - \mathbf{M}_i, \mathbf{u}_i(\mathbf{x}_{-i}) \otimes \mathbf{x}_i \rangle.$

Good-enough-response oracle for CE

Do we have a good-enough-response oracle in our problem?

$$\mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \mathbf{u}_i(\mathbf{x}_1, \dots, \mathbf{x}_n) \geq \mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \mathbf{u}_i(\phi_i(\mathbf{x}_i), \mathbf{x}_{-i}) - \epsilon.$$



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

because $\mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \langle \mathbf{x}_i - \mathbf{M}_i \mathbf{x}_i, \mathbf{u}_i(\mathbf{x}_{-i}) \rangle = \mathbb{E}_{(\mathbf{x}_1, \dots, \mathbf{x}_n) \sim \mu} \langle \mathbf{I} - \mathbf{M}_i, \mathbf{u}_i(\mathbf{x}_{-i}) \otimes \mathbf{x}_i \rangle.$



Fixed points implement a GER!

The ellipsoid against hope algorithm

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

- Even though the dual is infeasible, we *still* run the ellipsoid algorithm—this is why the algorithm is called ellipsoid against hope
- In every step, the ellipsoid starts from a point in the dual, and we use the good-enough-response oracle to refute that point—an entire halfspace
- Eventually the ellipsoid shrinks to have exponentially small volume

The ellipsoid against hope algorithm

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

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

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



The certificate of infeasibility is a correlated equilibrium.



We end up with a much smaller zero-sum game!

Nonlinear deviations

We want to minimize regret with respect to the following set of deviations

Definition 3.1. Given a map $\psi : \mathcal{X} \rightarrow \mathbb{R}^{k'}$, the set of deviations Φ^ψ is the set of all maps $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$ that can be expressed as the matrix-vector product $\mathbf{K}(\phi)\psi(\mathbf{x}) + \mathbf{c}(\phi)$ for some matrix $\mathbf{K} \in \mathbb{R}^{d \times k'}$ and $\mathbf{c} \in \mathbb{R}^d$. We denote by $k = d \times k' + d$ the dimension of $(\mathbf{K}, \mathbf{c})$.

- A canonical example to have in mind is **low-degree polynomials**

Optimizing over the set of deviations is hard

- Minimizing external regret over the set of deviations is hard
- So we cannot use the algorithm of Gordon et al.

Gordon et al.

Separation over the deviations

AND

Computing (expected) fixed points

Optimizing over the set of deviations is hard

- Minimizing external regret over the set of deviations is hard
- So we cannot use the algorithm of Gordon et al.

Gordon et al.

Daskalakis et al.

Separation over the deviations

Separation over the deviations

AND




OR

Computing (expected) fixed points

Computing (expected) fixed points

Semi-separation

To minimize Φ -regret, or running ellipsoid against hope, it's enough to have **semi-separation** oracle:



From Daskalakis et al.

Definition 3.2 (Semi-separation). In the *semi-separation* problem we are given a function $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$ and we have to compute

- *either* an ϵ -expected fixed point $\mu \in \Delta(\mathcal{X})$ of ϕ ,
- *or* a point $x \in \mathcal{X}$ such that $\phi(x) \notin \mathcal{X}$.

Minimizing Phi-regret with nonlinear deviations

Theorem. There is an algorithm that minimizes Phi-regret with respect to all deviations with **polynomial dimension**. Furthermore, it is possible to efficiently run EAH with respect to this set.

- This relies on the *semi-separation oracle*
- Captures as a special case low-degree polynomials
- The dimension of the set represents a fundamental barrier