Algorithms for Minimizing Phi-Regret: Nonlinear Deviations and Ellipsoid

15 888 **Computational Game Solving** (Fall 2025)

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.

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

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

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.

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

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

No if the learner is using mixed strategies! (Despite Kuhn’s theorem)

Expected fixed points

We define a relaxation of fixed points:

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

Expected fixed points

We define a relaxation of fixed points:

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

Expected fixed points are easy to compute!

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

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

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.

Input: An external regret minimizer $\mathfrak{R}_\Phi$ for the set Φ

NextStrategy():

  Set $\phi^{(t)} \coloneqq \mathfrak{R}_\Phi$.NextStrategy();

  return $\Delta(\mathcal{X}) \ni \boldsymbol{\mu}^{(t)} \coloneqq$ ExpectedFP$(\phi^{(t)})$;

ObserveUtility$(\boldsymbol{u}^{(t)})$:

  Set $u_\Phi^{(t)} : \phi \mapsto \mathbb{E}_{\boldsymbol{x}^{(t)} \sim \boldsymbol{\mu}^{(t)}} \langle \phi(\boldsymbol{x}^{(t)}), \boldsymbol{u}^{(t)} \rangle$;

  $\mathfrak{R}_\Phi$.ObserveUtility$(u_\Phi^{(t)})$;

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

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.

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

Computing correlated equilibria using ellipsoid

- Regret minimization gives an algorithm for computing (C)CEs that is polynomial in 1/ε. 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?

$$\mathbb{E}_{(\boldsymbol{x}_1,\ldots,\boldsymbol{x}_n)\sim\mu} u_i(\boldsymbol{x}_1,\ldots,\boldsymbol{x}_n) \geq \mathbb{E}_{(\boldsymbol{x}_1,\ldots,\boldsymbol{x}_n)\sim\mu} u_i(\phi_i(\boldsymbol{x}_i),\boldsymbol{x}_{-i}) - \epsilon.$$

The ellipsoid against hope algorithm

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

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

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

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

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:

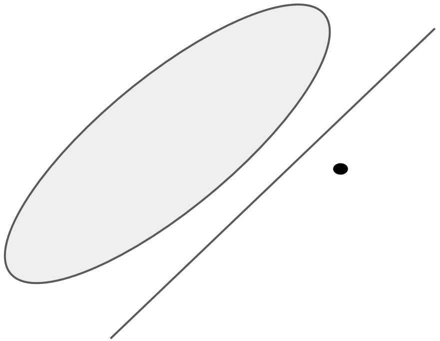

1. Deciding **membership** of a point

$\iff$ Regret minimization over that set!

2. Seperating a point from the set

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

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

The ellipsoid against hope algorithm

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

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

This can be solved under two basic assumptions:

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

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

💡 This implies that the dual is infeasible!

Good-enough-response oracle for CE

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

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

$$\Longleftrightarrow$$

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

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

Good-enough-response oracle for CE

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

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

$$\iff$$

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

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

Fixed points implement a GER!

The ellipsoid against hope algorithm

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

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

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

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

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

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

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

- 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.	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 Phi-regret, or running ellipsoid against hope, it's enough to have **semi-separation** oracle:

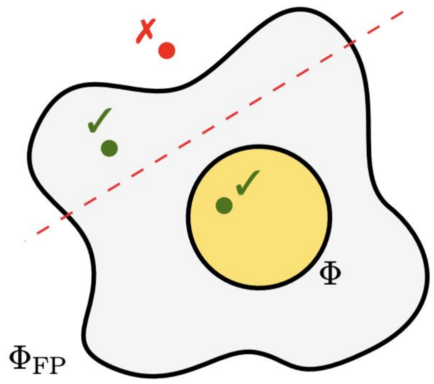

From Daskalakis et al.

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

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

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