The Complexity of Symmetric Equilibria in Min-Max Optimization and Team Zero-Sum Games

Ioannis Panageas
(UC Irvine & Archimedes AI)

Based on joint work with
I. Anagnostides, T. Sandholm, J. Yan

Playing Rock-Paper-Scissors

0, 0	-1, 1	1,-1
1,-1	0, 0	-1, 1
-1, 1	1, -1	0, 0

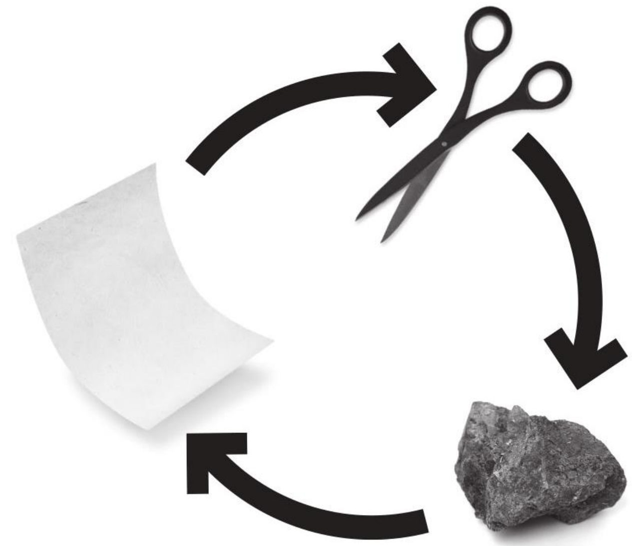

Two-player zero-sum. *Player* $\boldsymbol{y}$ *gets payoff* $\boldsymbol{x}^\top A\boldsymbol{y}$ *and* $\boldsymbol{x}$ *gets* $-\boldsymbol{x}^\top A\boldsymbol{y}$. *A Nash equilibrium* $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ *satisfies the Variational Inequalities*

$$\langle \mathbf{x}^*, A\mathbf{y}^* \rangle \leq \langle \mathbf{x}', A\mathbf{y}^* \rangle \text{ and } \langle \mathbf{x}^*, A\mathbf{y}^* \rangle \geq \langle \mathbf{x}^*, A\mathbf{y}' \rangle.$$

Playing Rock-Paper-Scissors

0, 0	-1, 1	1,-1
1,-1	0, 0	-1, 1
-1, 1	1, -1	0, 0

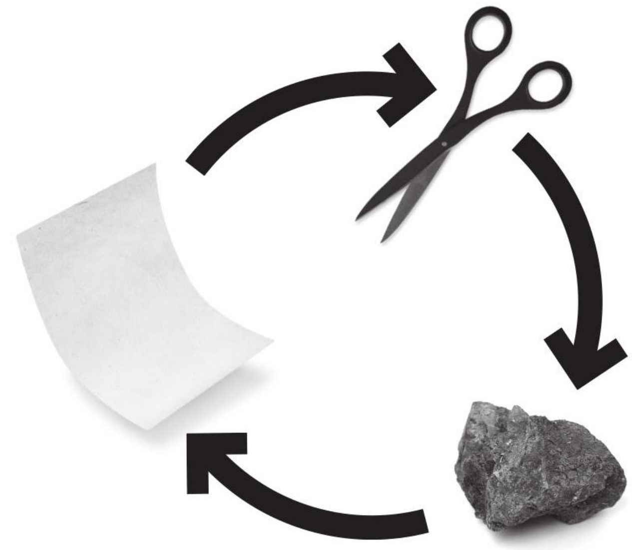

Two-player zero-sum. *Player $\boldsymbol{y}$ gets payoff $\boldsymbol{x}^\top A\boldsymbol{y}$ and $\boldsymbol{x}$ gets $-\boldsymbol{x}^\top A\boldsymbol{y}$. A Nash equilibrium is a solution to*

$$\min_{\boldsymbol{x}\in\Delta_n}\max_{\boldsymbol{y}\in\Delta_m} \boldsymbol{x}^\top A\boldsymbol{y}.$$

From Z.S games to min-max

Min-max optimization. *Let* $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ *be a G-Lipschitz, L-smooth function. Player* $\boldsymbol{y}$ *gets payoff* $f(\boldsymbol{x}, \boldsymbol{y})$ *and* $\boldsymbol{x}$ *gets* $-f(\boldsymbol{x}, \boldsymbol{y})$.

$$\min_{\boldsymbol{x} \in \mathcal{X}} \max_{\boldsymbol{y} \in \mathcal{Y}} f(\boldsymbol{x}, \boldsymbol{y}).$$

From Z.S games to min-max

Min-max optimization. *Let* $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ *be a G-Lipschitz, L-smooth function. Player* $\boldsymbol{y}$ *gets payoff* $f(\boldsymbol{x}, \boldsymbol{y})$ *and* $\boldsymbol{x}$ *gets* $-f(\boldsymbol{x}, \boldsymbol{y})$.

$$\min_{\boldsymbol{x} \in \mathcal{X}} \max_{\boldsymbol{y} \in \mathcal{Y}} f(\boldsymbol{x}, \boldsymbol{y}).$$

- When $f(\boldsymbol{x}, \boldsymbol{y}) = \langle \boldsymbol{x}, A\boldsymbol{y} \rangle$, NE exists, can be computed exactly in poly-time.
- When $f(\boldsymbol{x}, \boldsymbol{y})$ **convex-concave**, NE exists, can be computed in poly(1/ε).

Remark: NE $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ satisfies

$$f(\mathbf{x}^*, \mathbf{y}^*) \le f(\mathbf{x}', \mathbf{y}^*) \text{ and } f(\mathbf{x}^*, \mathbf{y}^*) \ge f(\mathbf{x}^*, \mathbf{y}').$$

- NE not guaranteed to exist for other cases.

Solution concepts

- **First-order NE aka fixed points of GDA:**

$$\langle \mathbf{x}' - \mathbf{x}^*, \nabla_{\mathbf{x}} f(\mathbf{x}^*, \mathbf{y}^*) \rangle \geq -\epsilon \quad \text{and} \ \langle \mathbf{y}' - \mathbf{y}^*, \nabla_{\mathbf{y}} f(\mathbf{x}^*, \mathbf{y}^*) \rangle \leq \epsilon$$

The ***Variational inequalities (VI)*** **always** have a solution. (Hartman-Stampacchia).

Remarks:

- GDA cycles, even for bilinear functions.
- When f **convex, non-concave** **or** **non-convex, concave** **ε-FONE** in poly(1/ε)
- When f **non-convex, non-concave** we do not know…

Solution concepts

- **First-order NE aka fixed points of GDA:**

$$\langle \mathbf{x}' - \mathbf{x}^*, \nabla_{\mathbf{x}} f(\mathbf{x}^*, \mathbf{y}^*) \rangle \geq -\epsilon \quad \text{and } \langle \mathbf{y}' - \mathbf{y}^*, \nabla_{\mathbf{y}} f(\mathbf{x}^*, \mathbf{y}^*) \rangle \leq \epsilon$$

The ***Variational inequalities (VI)*** **always** have a solution. (Hartman-Stampacchia).

Remarks:

- GDA cycles, even for bilinear functions.
- When f **convex, non-concave** **or** **non-convex, concave** **ε-FONE** in poly(1/ε)
- When f **non-convex, non-concave** we do not know…

- **Local NE:** **Relaxed** notion of NE, inequalities hold in a **neighborhood.**

Not guaranteed to exist!

- **Local Stackelberg:** x **in a** **neighborhood** **of** x^*

$$f(\mathbf{x}^*, \mathbf{y}') \leq f(\mathbf{x}^*, \mathbf{y}^*) \leq \max_{\mathbf{y}': \|\mathbf{y}' - \mathbf{y}^*\|_2} f(\mathbf{x}, \mathbf{y}').$$

Team ZS Games

Main focus: Team ZS Games

Team **A** VS Team **B**

- N, M players in teams **A**, **B**.
- Strategy sets $P_1, \ldots, P_N$ and $Q_1, \ldots, Q_M$.
- $U(\boldsymbol{x}, \boldsymbol{y})$ utility of each player from **B**, cost of each player from **A**.

NE in team zs games. *Nash equilibrium are FONE of the min-max*

$$\min_{\boldsymbol{x} \in \Delta(P_1) \times \ldots \times \Delta(P_N)} \max_{\boldsymbol{y} \in \Delta(Q_1) \times \ldots \times \Delta(Q_M)} U(\boldsymbol{x}, \boldsymbol{y}).$$

Other examples: Adversarial training, GANs, robust optimization. Use GDA and hope to stabilize...

An example – Generalized MP

	HH	HT/TH	TT
HH	$1, -1$	$\frac{1}{2}, -\frac{1}{2}$	$-1, 1$
HT/TH	$-\frac{1}{2}, \frac{1}{2}$	$0, 0$	$-\frac{1}{2}, \frac{1}{2}$
TT	$-1, 1$	$\frac{1}{2}, -\frac{1}{2}$	$1, -1$

- ***2** vs **2** players. Each player has two actions $\{\boldsymbol{H}, \boldsymbol{T}\}$.*
- *All vanilla methods you might have heard **cycle**. You can get coarse correlated eq.*
- *NE is the uniform.*

What is known so far

- For ***coupled domains***, ε-FONE is **PPAD**-complete [Daskalakis, Skoulakis, Zampetakis 21]

  Limitation: The construction of [DSZ21] works for ***non-convex linear***.

- For ***Adversarial*** team games ($|\boldsymbol{B}| = 1$) with N players, ε-NE is **CLS**-complete [Anagnostides et al 23].
- In ***Polymatrix*** two team zs with **independent** adversaries, ε-NE is **CLS**-complete. [Hollender, Maystre, Nagarajan 25]
- When $\boldsymbol{f}$ **convex, non-concave or non-convex, concave ε-FONE** in poly(1/ε).

What is known so far

- For ***coupled domains***, ε-FONE is **PPAD**-complete [Daskalakis, Skoulakis, Zampetakis 21]

  Limitation: The construction of [DSZ21] works for ***non-convex linear***.
- For ***Adversarial*** team games ($|\boldsymbol{B}| = 1$) with N players, ε-NE is **CLS**-complete [Anagnostides et al 23].
- In ***Polymatrix*** two team zs with **independent** adversaries, ε-NE is **CLS**-complete. [Hollender, Maystre, Nagarajan 25]
- When $\boldsymbol{f}$ **convex, non-concave or non-convex, concave ε-FONE** in poly(1/ε).

Main questions:

- What is the complexity of min-max for **product** domains.

  Maybe [DSZ21] hardness is because of the constraints [Bernasconi et al 25]?
- What is the complexity of **2 vs 1** in adversarial team games (constant no of players)?
- What is the complexity of **2 vs 2** or maybe **3 vs 3**?

Our main results (min-max)

Def. Symmetric min-max. *f is called anti-symmetric if*

$$f(\boldsymbol{x}, \boldsymbol{y}) = -f(\boldsymbol{y}, \boldsymbol{x}).$$

If f is anti-symmetric, min-max problem is called symmetric.

Theorem (PPAD-completeness for symmetric)**.** *Computing a symmetric $\frac{1}{n^c}$-approximate first-order Nash equilibrium in symmetric n-dimensional min-max optimization is PPAD-complete for any constant $c > 0$.*

Remark. *Theorem holds even for quadratic functions.*

Our main results (min-max)

Def. Symmetric min-max. *f is called anti-symmetric if*

$$f(\boldsymbol{x}, \boldsymbol{y}) = -f(\boldsymbol{y}, \boldsymbol{x}).$$

If f is anti-symmetric, min-max problem is called symmetric.

Theorem (PPAD-completeness for symmetric)**.** *Computing a symmetric $\frac{1}{n^c}$-approximate first-order Nash equilibrium in symmetric n-dimensional min-max optimization is PPAD-complete for any constant $c > 0$.*

Remark. *Theorem holds even for quadratic functions.*

Theorem (FNP-hardness for nonsymmetric)**.** *Computing a nonsymmetric approximate first-order Nash equilibrium in symmetric n-dimensional min-max optimization is FNP-hard.*

Remark. *Our results do not imply hardness for min-max. It is an indication that the problem is hard though.*

Our main results (team games)

Theorem (CLS-completeness for 2 vs 1). *Computing an ϵ-Nash equilibrium in 3-player (that is, 2 vs. 1) adversarial team games is CLS-complete.*

Remark. *Theorem holds even when one restricts to polymatrix, 3-player adversarial team games.*

Theorem (PPAD-completeness for 3 vs 3). *Computing a symmetric $\frac{1}{n^c}$-Nash equilibrium in symmetric, 6-player (3 vs. 3) team zero-sum polymatrix games is PPAD-complete for some constant $c > 0$.*

A simple reduction from 2-player

$1/n^c$-symmetic NE of $(R, R^\top)$ is PPAD-complete [Chen, Deng, Teng 09]

A simple reduction from 2-player

$1/n^c$-symmetic NE of $(R, R^\top)$ is PPAD-complete [Chen, Deng, Teng 09]

$A := \frac{R + R^\top}{2}$ (symmetric matrix)

$A = A^\top$

$C := \frac{R - R^\top}{2}$ (skew symmetric matrix)

$C = -C^\top$

$f(\mathbf{x}, \mathbf{y}) := \langle \mathbf{y}, A\mathbf{y} \rangle - \langle \mathbf{x}, A\mathbf{x} \rangle + \langle \mathbf{y}, C\mathbf{x} \rangle$

$\mathcal{X} \times \mathcal{Y} = \Delta_n \times \Delta_n$

A simple reduction from 2-player

$1/n^c$-symmetic NE of $(R, R^\top)$ is PPAD-complete [Chen, Deng, Teng 09]

$$A := \frac{R + R^\top}{2} \text{ (symmetric matrix)}$$
$$A = A^\top$$

$$C := \frac{R - R^\top}{2} \text{ (skew symmetric matrix)}$$
$$C = -C^\top$$

$$f(\mathbf{x}, \mathbf{y}) := \langle \mathbf{y}, A\mathbf{y} \rangle - \langle \mathbf{x}, A\mathbf{x} \rangle + \langle \mathbf{y}, C\mathbf{x} \rangle$$
$$\mathcal{X} \times \mathcal{Y} = \Delta_n \times \Delta_n$$

Claim. *Any symmetric* (x^*, x^*) ϵ*-FONE of* f *is an* ϵ*-NE of* $(R, R^\top)$.

A simple reduction from 2-player

$1/n^c$-symmetic NE of $(R, R^\top)$ is PPAD-complete [Chen, Deng, Teng 09]

$A := \frac{R + R^\top}{2}$ (symmetric matrix)

$A = A^\top$

$C := \frac{R - R^\top}{2}$ (skew symmetric matrix)

$C = -C^\top$

$f(\mathbf{x}, \mathbf{y}) := \langle \mathbf{y}, A\mathbf{y} \rangle - \langle \mathbf{x}, A\mathbf{x} \rangle + \langle \mathbf{y}, C\mathbf{x} \rangle$

$\mathcal{X} \times \mathcal{Y} = \Delta_n \times \Delta_n$

Claim. *Any symmetric* (x^*, x^*) *ϵ-FONE of f is an ϵ-NE of* $(R, R^\top)$.

ϵ-symmetic NE in 2-player symmetric games

P-**time** →

ϵ-symmetic FONE in symmetric min-max

Remarks

- The idea was to reduce from the VI problem (which is *PPAD-hard*)

$$\langle \boldsymbol{x} - \boldsymbol{x}^*, \boldsymbol{F}(\boldsymbol{x}^*) \rangle \leq \boldsymbol{\epsilon}$$

- Can reprove [DSZ21] if one considers constraints of the form idea

$$-\delta \leq \; x_i - y_i \leq \delta \text{ for appropriate } \delta.$$

- In the same spirit, [Bernasconi et al 25] shows for *box* constraints.

A reduction from max-clique

Graph $G([n], E)$, we construct $A_{i,j} = \begin{cases} 1/2 & \text{if } i = j, \\ 1 & \text{if } \{i, j\} \in E, \\ 0 & \text{otherwise.} \end{cases}$

Claim A. *Optimal NE in (A, A) is uniform with support on a max-clique.*

➡ Value is $\frac{k}{2k^2} + \frac{k(k-1)}{k^2} = 1 - \frac{1}{2k}$.

Best NE has value $1 - \frac{1}{2k}$ iff G has a max clique of size k.

A reduction from max-clique

Graph $G([n], E)$, we construct $A_{i,j} = \begin{cases} 1/2 & \text{if } i = j, \\ 1 & \text{if } \{i, j\} \in E, \\ 0 & \text{otherwise.} \end{cases}$

Claim A. *Optimal NE in (A, A) is uniform with support on a max-clique.*

Value is $\frac{k}{2k^2} + \frac{k(k-1)}{k^2} = 1 - \frac{1}{2k}$.

Best NE has value $1 - \frac{1}{2k}$ iff G has a max clique of size k.

Claim B. *Any ϵ-NE that does not have value ϵ-close to optimal, has value at most $1 - \frac{1}{2k} - \frac{1}{n^2k^4} + O(\epsilon)$ for some $\epsilon = 1/poly(n)$.*

Every ϵ-NE has value either at least $1 - \frac{1}{2k} - O(\epsilon)$ OR

at most $1 - \frac{1}{2k} - \frac{1}{n^2k^4} + O(\epsilon)$.

A reduction from max-clique

Symmetric identical payoff game $(\boldsymbol{B}, \boldsymbol{B})$

- $\boldsymbol{V} = \boldsymbol{1} - \frac{\boldsymbol{1}}{\boldsymbol{2k}}$
- $\boldsymbol{r} = \boldsymbol{1} - \frac{\boldsymbol{1}}{\boldsymbol{2k}} - \frac{\boldsymbol{1}}{\boldsymbol{2n^2k^4}} + \boldsymbol{O}(\boldsymbol{\varepsilon})$

$$B = \begin{bmatrix} A_{1,1} & \cdots & A_{1,n} & r \\ \vdots & \ddots & \vdots & \vdots \\ A_{n,1} & \cdots & A_{n,n} & r \\ r & \cdots & r & V \end{bmatrix}$$

Claim. *It is FNP-hard to find an* ϵ*-NE that does not have most of the mass on* $B_{n+1,n+1}$.

⇨ Finding two approximate NE with distance $> 1/\text{poly}(n)$ is FNP-hard

A reduction from max-clique

Theorem. *Finding a non-symmetric approximate FONE is FNP-hard.*

Separable min-max:

$$f(\mathbf{x}, \mathbf{y}) := \langle \mathbf{y}, B\mathbf{y} \rangle - \langle \mathbf{x}, B\mathbf{x} \rangle$$

$$\mathcal{X} \times \mathcal{Y} = \Delta_n \times \Delta_n$$

Non-symmetric FONE of f ⟺ Two NE in (B, B)

A reduction from max-clique

Theorem. *Finding a non-symmetric approximate FONE is FNP-hard.*

Separable min-max:

$$f(\mathbf{x}, \mathbf{y}) := \langle \mathbf{y}, B\mathbf{y} \rangle - \langle \mathbf{x}, B\mathbf{x} \rangle$$

$$\mathcal{X} \times \mathcal{Y} = \Delta_n \times \Delta_n$$

Non-symmetric FONE of f ⟷ Two NE in (B, B)

ϵ-non-symmetric FONE

Max-clique

P-time →

- Remark: Proof in the same spirit as in [McLennan and Tourky 10']

2 vs 1 is CLS-complete

ϵ-KKT for quadratic is CLS-hard [Fearnley et al 23]

↓ P-time

symmetric ϵ-NE for (A, A) is CLS-complete [Ghosh and Hollender 24]

↓ P-time

ϵ-NE for 2 vs 1 adversarial team games

2 vs 1 is CLS-complete

ϵ-KKT for quadratic is CLS-hard [Fearnley et al 23]

↓ P-time

symmetric ϵ-NE for (A, A) is CLS-complete [Ghosh and Hollender 24]

↓ P-time

ϵ-NE for 2 vs 1 adversarial team games

Main idea. *Consider a two player symmetric identical payoff (A, A) and add a third player that forces symmetry.*

$$u(\mathbf{x}, \mathbf{y}, \mathbf{z}) := \langle \mathbf{x}, A\mathbf{y} \rangle + \frac{1}{\epsilon} \sum_{i=1}^{n} (z_i(x_i - y_i) + z_{n+i}(y_i - x_i)) + z_{2n+1}.$$

$\mathbf{x}, \mathbf{y}$ maximizers, $\mathbf{z}$ minimizer.

Take away messages and future directions

- We provide strong indication that min-max is hard.
- Complexity of adversarial team games is resolved.
- The complexity of min-max is *still open*.

- *Positive* results for low degree polynomials and well-behaved domains?
- Prove unconditional lower bounds.