Optimistic Online Learning in Symmetric Cone Games

Anas Barakat

Joint work with Wayne Lin, John Lazarsfeld, Antonios Varvitsiotis

July 24th - ICCOPT 2025

SINGAPORE UNIVERSITY OF TECHNOLOGY AND DESIGN

Game Theory Formalism

- Finite number of players $\mathcal{N} := \{1, \cdots, N\}$.
- Strategy set of player i: $\mathcal{X}_i$, joint strategy set: $\mathcal{X} := \prod_{i=1}^{N} \mathcal{X}_i$.
- Utility $u_i : \mathcal{X} \to \mathbb{R}$ assumed to be concave w.r.t. its ith variable and differentiable.

ZOOM on the strategy spaces of agents (typically convex).

Examples of Games

- **Finite normal-form games** $\mathcal{X}_i = \Delta(\mathcal{A}_i)$: Golden standard.
- **Convex games with ball strategy sets** $\mathcal{X}_i = B(0, r_i)$.
- **PSD matrix games** $\mathcal{X}_i = \Delta_{\mathbb{S}^n_+}$.
  - Each player controls a PSD matrix variable (e.g. a signal covariance matrix) .
  - Applications in wireless communication networks for the competitive maximization of mutual information in interfering networks [Arslan et al., 2007, Scutari et al., 2008, Mertikopoulos and Moustakas, 2015, Majlesinasab et al., 2019].
- **Quantum games** $\mathcal{X}_i = \Delta_{\mathbb{H}^n_+}$.
  - Strategies of the players are quantum states represented by density matrices.
  - Utility is the expected value of a measurement on the joint state.

Example 1: Distance Metric Learning

- E.g. Learn a Mahalanobis distance given a dataset $\{x_i\}_{1\leq i\leq N}$ where $x_i \in \mathbb{R}^d$.
- [Ying and Li, 2012]

$$\begin{aligned} &\max_{M\in\mathbb{S}^d_+} \quad \min_{(i,j)\in\mathcal{D}} \quad \underbrace{d_M^2(x_i, x_j)}_{(x_i-x_j)^T M(x_i-x_j)} \\ &\text{s.t.} \quad \sum_{(i,j)\in\mathcal{S}} d_M^2(x_i, x_j) \leq 1 \end{aligned} \tag{1}$$

Simplex-Spectraplex Game

$$\min_{x\in\Delta^{m-1}} \max_{Y\in\Delta_{\mathbb{S}^d_+}} f(x, Y) := \langle Y, \mathcal{A}(x)\rangle + \langle b, x\rangle + \langle C, Y\rangle \ , \tag{2}$$

where $\mathcal{A} : \mathbb{R}^m \to \mathbb{S}^d$ is the linear map given by $\mathcal{A}(x) = \sum_{i=1}^m x_i A_i$ for some $A_i \in \mathbb{S}^d$.

- Smoothing [Nesterov, 2007] in $\mathcal{O}(1/\varepsilon)$ iterations, interior point methods ...

Example 2: Fermat-Weber Problem

$$\min_{x \in B(0,R)} \left\{ g(x) := \sum_{i=1}^{p} \|A_i x - b_i\|_2 \right\} .$$

- **Min-max reformulation:** variational characterization of the maximal eigenvalue,

$$g(x) = \sum_{i=1}^{p} \lambda_{\max}(\bar{A}_i x - \bar{b}_i) = \sum_{i=1}^{p} \max_{\bar{u}_i \in \Delta_{\mathcal{K}}} \langle \bar{u}_i, \bar{A}_i x - \bar{b}_i \rangle_{\mathcal{J}} = \max_{y \in \Delta_{\mathcal{K}}^{p}} \langle y, \bar{A}x - \bar{b} \rangle_{\mathcal{J}^p}$$

where $\bar{A}_i := \left(0, A_i^T\right)^T \in \mathbb{R}^{(m+1)\times d}$ and $\bar{b}_i := \left(0, b_i^T\right)^T \in \mathbb{R}^{m+1}$, $\mathcal{J} = \mathbb{L}^{d+1}$.

Second-Order Cone Min-Max Game

$$\min_{\bar{x}=(1/2,\tilde{x}) \in \Delta_{\mathbb{L}_+^{d+1}}} \quad \max_{y \in \prod_{i=1}^{p} \Delta_{\mathbb{L}_+^{d+1}}} \left\{ f(\bar{x}, y) := \langle \tilde{A}\bar{x}, y \rangle - \langle \bar{b}, y \rangle \right\}, \; \tilde{A} := 2R(0 \quad \bar{A}) . \qquad (3)$$

- IP method of [Xue and Ye, 1997], extension of smoothing [Baes, 2006].

Questions

- Can we unify all these games?
- Can we learn Nash equilibria in the 2-player zero-sum setting using a **SINGLE ALGORITHM** for all these games simultaneously?

Two-player zero-sum SCGs

$$\min_{x \in \Delta_{\mathcal{K}_1}} \max_{y \in \Delta_{\mathcal{K}_2}} f(x, y)$$

SHORT ANSWERS

- Can we unify all these games?

SYMMETRIC CONE GAMES

- Can we learn Nash equilibria efficiently in the 2-player zero-sum setting using a single algorithm for all these games simultaneously?

OPTIMISTIC SYMMETRIC CONE MULTIPLICATIVE WEIGHTS

Outline

1. **Symmetric Cones and SC Games**

2. **Optimistic Online Learning in Symmetric Cone Games**

3. **Applications to Min-Max Problems over Symmetric Cones**

Two-player zero-sum SCGs

$$\min_{x \in \Delta_{\mathcal{K}_1}} \max_{y \in \Delta_{\mathcal{K}_2}} f(x, y)$$

v

Examples of Symmetric Cones and corresponding EJAs

EJA $\mathcal{J}$	Inner product $\langle x, y\rangle$	Jordan product $x \circ y$	Cone of squares $\mathcal{K}$
Euclidean space $(\mathbb{R}^n)$	$\sum_{i=1}^n x_i y_i$	$(x_i y_i)_{i=1,\cdots,n}$	nonnegative orthant $(\mathbb{R}^n_+)$
Real sym. matrices $(\mathbb{S}^n)$	$\mathrm{tr}(xy)$	$\frac{1}{2}(xy + yx)$	PSD cone $(\mathbb{S}^n_+)$
Jordan spin algebra $(\mathbb{L}^n)$	$2\sum_{i=1}^n x_i y_i$	$(\bar{x}^\top \bar{y}, x_1\bar{y} + y_1\bar{x})$	second-order cone $(\mathbb{L}^n_+)$

- Characterization of SCs and formalism of Euclidean Jordan Algebras [Faraut and Korányi, 1994]: *Any* symmetric cone is the cone of squares $\{x \circ x : x \in \mathcal{J}\}$ of some EJA $\mathcal{J}$.

Strategy sets: Generalized Simplexes

Generalized Simplex

If $(\mathcal{J}, \circ)$ is an EJA and $\mathcal{K}$ its cone of squares,

$$\Delta_{\mathcal{K}} := \{x \in \mathcal{K} : \mathrm{tr}(x) = 1\}.$$

Symmetric Cone $\mathcal{K}$	Generalized Simplex $\Delta_{\mathcal{K}}$
Nonnegative orthant $\mathbb{R}^n_+$	Simplex $\Delta^{n-1} = \{x \in \mathbb{R}^n_+ : \sum_{i=1}^n x_i = 1\}$
Real PSD symmetric matrices $\mathbb{S}^n_+$	Spectraplex $\Delta_{\mathbb{S}^n_+} = \{X \in \mathbb{S}^n_+ : \mathrm{Tr}(X) = 1\}$
PSD Hermitian matrices $\mathbb{H}^n_+$	Spectraplex $\Delta_{\mathbb{H}^n_+} = \{X \in \mathbb{H}^n_+ : \mathrm{Tr}(X) = 1\}$
Second-order cone $\mathbb{L}^n_+$	Ball $\Delta_{\mathbb{L}^n_+} = \{(\frac{1}{2}, x) \in \mathbb{R}^n : \|x\|_2 \leq \frac{1}{2}\}$

Definition: Symmetric Cone Games

- Finite number of players $\mathcal{N} := \{1, \cdots, N\}$.
- Strategy set of player i: generalized simplex $\Delta_{\mathcal{K}_i}$ where $\mathcal{K}_i$ symmetric cone.
  - Notation: Space of joint strategies $\mathcal{X} := \prod_{i \in \mathcal{N}} \Delta_{\mathcal{K}_i}$.
- Utility $u_i : \mathcal{X} \to \mathbb{R}$ assumed to be concave w.r.t. its ith variable and differentiable.
  - Notation: Payoff vector,

$$\forall x \in \mathcal{X},\ m(x) = (m_i(x))_{i \in \mathcal{N}}\,, \quad m_i(x) = \nabla_{x_i} u_i(x_i, x_{-i}), \forall i \in \mathcal{N}\,.$$

Examples of Symmetric Cone Games

- Finite normal-form games
- Convex games with ball strategy sets
- PSD matrix games
- Quantum games

- Nonnegative orthant
- 2nd order cone
- Cone of PSD matrices
- Cone of complex PSD matrix

Online Learning in Symmetric-Cone Games

Initialization: $x_i^1 \in \Delta_{\mathcal{K}_i}, \forall i \in \mathcal{N}$

For $t = 1, \ldots, T$:

For $i \in \mathcal{N}$ (simultaneously):

Player i observes their (and only theirs) payoff $m_i^t \in \mathcal{J}_i$

Player i receives linear payoff $\langle m_i^t, x_i^t \rangle$

Player i computes their next strategy $x_{t+1}^i \in \Delta_{\mathcal{K}_i}$

Output: $\bar{x}_T^i = \frac{1}{T}\sum_{t=1}^T x_i^t \, , \forall i \in \mathcal{N}$.

*canonical EJA inner product $\langle x, y \rangle = tr(x \circ y)$.

Playerwise Regret

$$r_i(T) := \sup_{x \in \Delta_{\mathcal{K}_i}} \sum_{t=1}^{T} u_i^t(x, x_{-i}^t) - u_i(x_i^t, x_{-i}^t)\,, \quad \forall i \in \mathcal{N}\,.$$

Optimistic Follow The Regularized Leader

- Fix player $i \in \mathcal{N}$.

OFTRL

$$x^{t+1} = \operatorname*{argmax}_{x \in \Delta_{\mathcal{K}}} \left\{ \underbrace{\eta}_{\text{step size}} \left\langle \sum_{k=1}^{t} m^k + \tilde{m}^{t+1}, x \right\rangle - \underbrace{\Phi(x)}_{\text{SC regularizer}} \right\}$$

where $m^k = \nabla_{x_i} u_i(x_i^k, x_{-i}^k)$ and $(\tilde{m}^t)$ is a predictor sequence, typically $\tilde{m}^{t+1} = m^t$.

- Why optimism?

- Which regularizer?

Why Optimistic Online Learning? Related work

- In min-max problems (under min-max theorem):
  - average regret is bounded above by $\varepsilon \implies$ time-average iterate is an ε-saddle point.

Fully adversarial setting	Game setting
Average regret	Average regret in 2pZS and N-player NF
$\mathcal{O}(T^{-1/2})$	$\tilde{\mathcal{O}}(T^{-1})$

- [Daskalakis et al., 2011, Chiang et al., 2012, Rakhlin and Sridharan, 2013, Syrgkanis et al., 2015, Daskalakis et al., 2021] ...
- [Vasconcelos et al., 2023] zero-sum quantum games, optimization perspective.

Regularizer? Symmetric Cone Negative Entropy

Negative Entropy

Given the EJA $\mathcal{J}$ and its cone of squares $\mathcal{K}$, $\Phi_{\text{ent}} : \text{int}(\mathcal{K}) \to \mathbb{R}$:

$$\forall x \in \text{int}(\mathcal{K}), \quad \Phi_{\text{ent}}(x) = \text{tr}(x \circ \ln x) = \sum_{i=1}^{r} \lambda_i \ln \lambda_i \,, \tag{SCNE}$$

- $x = \sum_{i=1}^{r} \lambda_i q_i \in \text{int}(\mathcal{K})$ spectral decomposition of x.
- $\ln : \text{int}(\mathcal{K}) \to \mathcal{J}$ Löwner extension of the scalar log, i.e. $\ln x = \sum_{i=1}^{r} \ln(\lambda_i) q_i$.
- Note that $\lambda_i > 0$ for every $1 \leq i \leq r$ as $x \in \text{int}(\mathcal{K})$.

- Exponential mapping $\exp : \mathcal{J} \to \text{int}(\mathcal{K})$ is defined by $\exp(x) = \sum_{i=1}^{r} \exp(\lambda_i)\, q_i$.

Strong Convexity of Symmetric Cone Negative Entropy

Theorem

Let $(\mathcal{J}, \circ)$ be an EJA and let $\mathcal{K}$ be its cone of squares. Then,

$$\forall x, y \in \operatorname{int}(\Delta_{\mathcal{K}}), \quad D_{\Phi_{\text{ent}}}(x, y) \geq \frac{1}{2}\|x - y\|_{tr,1}^2 . \tag{4}$$

where $D_{\Phi_{\text{ent}}}$ is the Bregman divergence $D_{\Phi_{\text{ent}}}(x, y) = \operatorname{tr}(x \circ \ln x - x \circ \ln y + y - x)$.

- Proof sketch:

$$D_{\Phi}(x, y) \underset{(1+2)}{\geq} D_{\Phi}(T(x), T(y)) = \mathrm{KL}(u(x)||u(y)) \geq \|u(x) - u(y)\|_1^2 = \frac{1}{2}\|x - y\|_{tr,1}^2 ,$$

1. The diagonal mapping is a convex combination of EJA automorphisms
2. Use joint convexity of the relative entropy and properties of automorphisms.

- Alternative proof [Baes, 2006] (Hessian and duality).

Optimistic Symmetric Cone Multiplicative Weights Update

- Regularizer Φ = Symmetric Cone Negative Entropy

OSCMWU Algorithm

$$w^{t+1} = \eta \left(\sum_{k=1}^{t} m^k + \tilde{m}^{t+1} \right), \qquad x^{t+1} = \frac{\exp(w^{t+1})}{\mathrm{tr}(\exp(w^{t+1}))}, \quad \forall t \geq 1,$$

where $(\tilde{m}^t)$ is a predictor sequence, typically $\tilde{m}^{t+1} = m^t$.

- $\tilde{m}^{t+1} = 0$: SCMWU introduced and studied recently in [Canyakmaz et al., 2023]

Regret in Symmetric Cone Games

Theorem [Syrgkanis et al., 2015]

Under smoothness of payoff vectors, if each player $i \in \mathcal{N}$ runs OSCMWU for T rounds on $\Delta_{\mathcal{K}_i}$ with stepsize $\eta = 1/(2\sqrt{N \sum_{i=1}^N L_i^2})$ and set $\| \cdot \| = \| \cdot \|_{tr,1}$. Then

$$\sum_{i=1}^{N} r_i(T) \leq 2 \left(\sum_{i=1}^{N} R_i \right) \cdot \sqrt{N \sum_{i=1}^{N} L_i^2}\,, \tag{5}$$

where $r_i(T)$ is the i-th player's regret and $R_i = \sup_{x \in \Delta_{\mathcal{K}_i}} \Phi(x) - \inf_{x \in \Delta_{\mathcal{K}_i}} \Phi(x)$.

Application to 2-Player Zero-Sum Symmetric Cone Games

Min-Max Problem

$$\min_{x \in \Delta_{\mathcal{K}_1}} \max_{y \in \Delta_{\mathcal{K}_2}} f(x, y),$$

- $f : \mathcal{J}_1 \times \mathcal{J}_2 \to \mathbb{R}$ is convex-concave and differentiable.
- $\mathcal{K}_1, \mathcal{K}_2$ are arbitrary symmetric cones and $\Delta_{\mathcal{K}_1}, \Delta_{\mathcal{K}_2}$ their generalized simplexes.

Theorem (2-player Zero-Sum SCG) - adaptation of folklore result

If both players run OSCMWU with stepsize $\eta = 1/(2\sqrt{2(L_1^2 + L_2^2)})$,

$$T \geq \frac{2(\ln r_1 + \ln r_2)\sqrt{2(L_1^2 + L_2^2)}}{\varepsilon} \implies \left(\bar{x}_T = \frac{1}{T}\sum_{t=1}^{T} x^t, \bar{y}_T = \frac{1}{T}\sum_{t=1}^{T} y^t\right) \quad \varepsilon\text{-SP}.$$

$r_i = \operatorname{rank}(\mathcal{J}_i), i = 1, 2.$

Conclusion

- **SCGs:** Exploit geometric structure and unify several existing problems.
  - Normal form, quantum, PSD games, convex games with ball strategy sets.
- **OSCMWU:** single algorithm applying in a unified way instead of there exist ad-hoc algorithms for special cases of our setting.
  - **Log dependence on the intrinsic dimension of the problem**.
- **Applications** beyond normal-form and quantum games.
  - simplex-spectraplex, second-order cone games.

References I

Arslan, G., Demirkol, M. F., and Song, Y. (2007).
Equilibrium efficiency improvement in MIMO interference systems: A decentralized stream control approach.
IEEE Transactions on Wireless Communications, 6(8):2984–2993.

Baes, M. (2006).
Spectral functions and smoothing techniques on Jordan algebras.
Catholic University of Louvain. Ph. D Thesis.

Canyakmaz, I., Lin, W., Piliouras, G., and Varvitsiotis, A. (2023).
Multiplicative updates for online convex optimization over symmetric cones.
arXiv preprint arXiv:2307.03136.

References II

Chiang, C.-K., Yang, T., Lee, C.-J., Mahdavi, M., Lu, C.-J., Jin, R., and Zhu, S. (2012).
Online optimization with gradual variations.
In *Conference on Learning Theory*, pages 6–1. JMLR Workshop and Conference Proceedings.

Daskalakis, C., Deckelbaum, A., and Kim, A. (2011).
Near-optimal no-regret algorithms for zero-sum games.
In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 235–254. SIAM.

Daskalakis, C., Fishelson, M., and Golowich, N. (2021).
Near-optimal no-regret learning in general games.
Advances in Neural Information Processing Systems, 34:27604–27616.

References III

Faraut, J. and Korányi, A. (1994).
Analysis on symmetric cones.
Oxford university press.

Majlesinasab, N., Yousefian, F., and Feizollahi, M. J. (2019).
A first-order method for monotone stochastic variational inequalities on semidefinite matrix spaces.
In *2019 American Control Conference (ACC)*, pages 169–174. IEEE.

Mertikopoulos, P. and Moustakas, A. L. (2015).
Learning in an uncertain world: MIMO covariance matrix optimization with imperfect feedback.
IEEE Transactions on Signal Processing, 64(1):5–18.

References IV

Nesterov, Y. (2007).
Smoothing technique and its applications in semidefinite optimization.
Mathematical Programming, 110(2):245–259.

Rakhlin, S. and Sridharan, K. (2013).
Optimization, learning, and games with predictable sequences.
Advances in Neural Information Processing Systems, 26.

Scutari, G., Palomar, D. P., and Barbarossa, S. (2008).
Competitive design of multiuser MIMO systems based on game theory: A unified view.
IEEE Journal on Selected Areas in Communications, 26(7):1089–1103.

Syrgkanis, V., Agarwal, A., Luo, H., and Schapire, R. E. (2015).
Fast convergence of regularized learning in games.
Advances in Neural Information Processing Systems, 28.

References V

Vasconcelos, F., Vlatakis-Gkaragkounis, E.-V., Mertikopoulos, P., Piliouras, G., and Jordan, M. I. (2023).
A quadratic speedup in finding Nash equilibria of quantum zero-sum games.
Quantum Techniques in Machine Learning Conference 2023, arXiv:2311.10859.

Xue, G. and Ye, Y. (1997).
An efficient algorithm for minimizing a sum of Euclidean norms with applications.
SIAM Journal on Optimization, 7(4):1017–1036.

Ying, Y. and Li, P. (2012).
Distance metric learning with eigenvalue optimization.
The Journal of Machine Learning Research, 13(1):1–26.

Symmetric Cones and Euclidean Jordan Algebras I

Definition of symmetric cones

- A cone $\mathcal{K}$ in an inner product space is *symmetric* if self-dual and homogeneous.

Let $\mathcal{J}$ be a finite dim. vector space with a bilinear product $\circ : \mathcal{J} \times \mathcal{J} \to \mathcal{J}$.

- $(\mathcal{J}, \circ)$ Jordan algebra if:

$$\forall x, y \in \mathcal{J}, \quad x \circ y = y \circ x \,, \quad \underbrace{x^2}_{x \circ x} \circ (x \circ y) = x \circ (x^2 \circ y) \,.$$

- $(\mathcal{J}, \circ, (\cdot, \cdot))$: *Euclidean* Jordan algebra over $\mathbb{R}$ if equipped with an associative inner product $(\cdot, \cdot)$, i.e. for all $x, y, z \in \mathcal{J}, (x \circ y, z) = (y, x \circ z)$.

Characterization of Symmetric cones

- *Any* symmetric cone is the cone of squares $\{x \circ x : x \in \mathcal{J}\}$ of some EJA $\mathcal{J}$ [Faraut and Korányi, 1994]

A Simplex-Spectraplex Game (2/2) I

- **Iteration complexity of OSCMWU:** $T \geq \frac{4(\ln m + \ln d) \max_i \|A_i\|_{\mathrm{tr},\infty}}{\varepsilon}$ iterations.
- **Prior work:**
  - Smoothing technique [Nesterov, 2007] in $\mathcal{O}(1/\varepsilon)$ iterations (similar).
  - Frank-Wolfe method + smoothing [Ying and Li, 2012] in $\mathcal{O}(1/\varepsilon^2)$ iterations but only $\mathcal{O}(d^2)$ runtime cost.
  - Interior point methods: high precision but prohibitive per-iteration cost for large-scale problems.

A 2nd-Order Cone Min-Max Game: (2/2) I

$$\min_{\bar{x}=(1/2,\tilde{x})\in\Delta_{\mathbb{L}_+^{d+1}}} \max_{y\in\prod_{i=1}^{p}\Delta_{\mathbb{L}_+^{d+1}}} \left\{ f(\bar{x},y) := \langle \tilde{A}\,\bar{x}, y\rangle - \langle \bar{b}, y\rangle \right\}, \; \tilde{A} := 2R(0 \quad \bar{A})\,.$$

- **Iteration complexity of OSCMWU:** $T \geq \frac{4(p+1)L\ln 2}{\varepsilon}$ iterations.
- **Prior work:**
  - Extension of the smoothing technique of Nesterov to EJAs [Baes, 2006].
  - Interior point method of [Xue and Ye, 1997] requires fewer iterations but much higher per-iteration cost.

Regret Bounded by Variation in Utilities [Syrgkanis et al., 2015] I

RVU

For (x^t) generated by OFTRL with $\tilde{m}^{t+1} = m^t$ and a regularizer Φ that is 1-strongly convex w.r.t. a norm $\|\cdot\|$, for all $T \geq 1$,

$$\forall x \in \Delta_{\mathcal{K}}, \quad \sum_{t=1}^{T} f^t(x) - f^t(x^t) \leq \frac{R}{\eta} + \eta \sum_{t=1}^{T} \|m^t - m^{t-1}\|_*^2 - \frac{1}{4\eta} \sum_{t=1}^{T} \|x^t - x^{t-1}\|^2 ,$$

where $R = \sup_{x \in \Delta_{\mathcal{K}}} \Phi(x) - \inf_{x \in \Delta_{\mathcal{K}}} \Phi(x)$, $\|\cdot\|_*$ is the dual norm and $\langle \cdot, \cdot \rangle$ the EJA inner product.

Online Symmetric-Cone Optimization (OSCO)

Initialization: $x^1 \in \Delta_{\mathcal{K}} := \{x \in \mathcal{K} : \mathrm{tr}(x) = 1\}$

For $t = 1, \ldots, T$:

Observe the payoff $m^t \in \mathcal{J}$

Receive linear payoff $\langle m^t, x^t \rangle$

Compute new iterate $x^{t+1} \in \Delta_{\mathcal{K}}$