Towards characterizing the first-order query complexity of learning (approximate) Nash equilibria in zero-sum matrix games

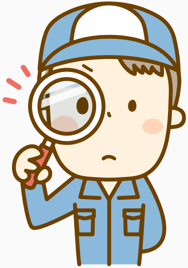

Wouter M. Koolen

CWI Centrum Wiskunde & Informatica

UNIVERSITY OF TWENTE.

MDS Seminar
University of Twente
Monday 6th November, 2023

Warm Thanks

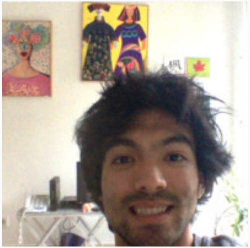

Hédi Hadiji

Sarah Sachs

Tim van Erven

Where we are heading today

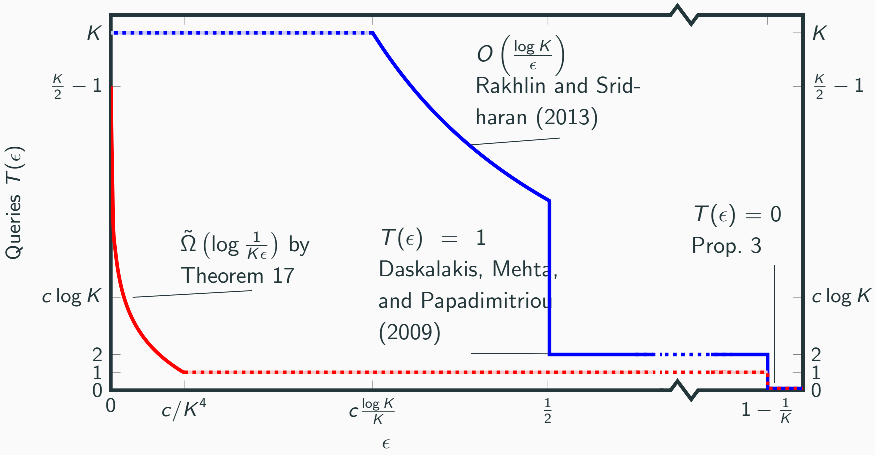

Outline

1. Motivation
2. Upper Bounds
3. New Lower Bounds
4. State of the Art
5. Conclusion

Games!

Lots of interest, old and new, in solving **convex-concave** min-max problems

$$\min_{p \in \mathcal{P}} \max_{q \in \mathcal{Q}} f(p, q)$$

Games!

Lots of interest, old and new, in solving **convex-concave** min-max problems

$$\min_{p \in \mathcal{P}} \max_{q \in \mathcal{Q}} f(p, q)$$

- Economics
- Optimisation
- Machine learning (GANs)
- Online learning and Bandits (Track-and-Stop)
- ...

What is a solution?

Given $\epsilon \geq 0$, we aim to find an approximate saddle point

$$(p_\star, q_\star) \in \mathcal{P} \times \mathcal{Q},$$

satisfying

$$\max_{q \in \mathcal{Q}} f(p_\star, q) - \min_{p \in \mathcal{P}} f(p, q_\star) \ \leq \ 2\epsilon$$

How are we going to find that solution

We consider the first-order query model.

We start with an unknown f from a known class $\mathcal{F}$.

Interaction protocol

In rounds $1, 2, \ldots, T$

- Learner issues query (p, q)
- Learner receives **feedback** $(\nabla_p f(p, q), \nabla_q f(p, q))$

The learner outputs an ϵ-optimal saddle point $(p_\star, q_\star)$.

How are we going to find that solution

We consider the first-order query model.

We start with an unknown f from a known class $\mathcal{F}$.

Interaction protocol

In rounds $1, 2, \ldots, T$

- Learner issues query (p, q)
- Learner receives **feedback** $(\nabla_p f(p, q), \nabla_q f(p, q))$

The learner outputs an ϵ-optimal saddle point $(p_\star, q_\star)$.

Query complexity

How many first-order queries are necessary and sufficient for a sequential learner to output an approximate saddle point for any $f \in \mathcal{F}$?

Making the question more tractable

Consider **special case** of zero-sum matrix games (bilinear functions over probability simplex):

$$\mathcal{P} = \mathcal{Q} = \triangle_K, \qquad \mathcal{F} = \left\{ f(p,q) = p^\intercal M q \;\middle|\; M \in [\pm 1]^{K \times K} \right\}$$

$$(\nabla_p f(p,q), \nabla_q f(p,q)) = (Mq, M^\intercal p)$$

Making the question more tractable

Consider **special case** of zero-sum matrix games (bilinear functions over probability simplex):

$$\mathcal{P} = \mathcal{Q} = \triangle_K, \qquad \mathcal{F} = \left\{ f(p,q) = p^\intercal M q \;\middle|\; M \in [\pm 1]^{K \times K} \right\}$$

$$(\nabla_p f(p,q), \nabla_q f(p,q)) = (Mq, M^\intercal p)$$

Algorithms since Brown (1951), up to Rakhlin and Sridharan (2013).

Lower bounds remain elusive.

$\Rightarrow$ Optimal query complexity **unknown**.

What is known: Upper Bounds

1951: First iterative methods by Brown (1951) and Robinson (1951).

1999: Freund and Schapire (1999) discovered the relation to Regret Bounds: Can compute an ϵ-Nash-equilibrium with T iterations, where

$$T = O\left(\frac{\log K}{\epsilon^2}\right)$$

2011: Daskalakis, Deckelbaum, and Kim (2011) can compute an ϵ-Nash-equilibrium with T iterations, where

$$T = O\left(\frac{f(K)}{\epsilon}\right)$$

2013: Rakhlin and Sridharan (2013) can compute an ϵ-Nash-equilibrium with T iterations, where

$$T = O\left(\frac{\log K}{\epsilon}\right)$$

What is known: Lower Bounds

2018: Ouyang and Xu (2021) Showed a lower bound on the query complexity for Saddle-Point Problems with curvature and rotationally invariant constraint sets.

2020: Ibrahim et al. (2020) adapted Nesterov's lower bound technique for games. This requires a two-step linear span assumption.

Our Results

Lower Bounds:

- $K/2 - 1$ queries needed for learning **exact** $\epsilon = 0$ equilibrium
- $\Omega\left(\frac{\log \frac{1}{\epsilon K}}{\log K}\right)$ queries required when $\epsilon \leq \frac{1}{K^4}$.

Upper Bounds:

- If entries in known countable set, say $M \in \mathbb{Q}^{K \times K}$, one query suffices.

Outline

One Query Suffices

Suppose $M_{ij} \in \{0, \ldots, n-1\}$ for some $n \geq 1$.

One Query Suffices

Suppose $M_{ij} \in \{0, \ldots, n-1\}$ for some $n \geq 1$.

Consider query (p, q) with p arbitrary and $q_j \propto n^{-j}$. Then the i^{th} entry of the feedback (to the p player) is

$$\nabla_p f(p, q)_i \; = \; \sum_{j=1}^{K} M_{ij} q_j \; \propto \; \sum_{j=1}^{K} M_{ij} n^{-j}$$

One Query Suffices

Suppose $M_{ij} \in \{0, \ldots, n-1\}$ for some $n \geq 1$.

Consider query (p, q) with p arbitrary and $q_j \propto n^{-j}$. Then the i^{th} entry of the feedback (to the p player) is

$$\nabla_p f(p, q)_i \; = \; \sum_{j=1}^{K} M_{ij} q_j \; \propto \; \sum_{j=1}^{K} M_{ij} n^{-j}$$

This is the i^{th} row of M written in base n.

One Query Suffices

Suppose $M_{ij} \in \{0, \ldots, n-1\}$ for some $n \geq 1$.

Consider query (p, q) with p arbitrary and $q_j \propto n^{-j}$. Then the i^{th} entry of the feedback (to the p player) is

$$\nabla_p f(p, q)_i \;=\; \sum_{j=1}^{K} M_{ij} q_j \;\propto\; \sum_{j=1}^{K} M_{ij} n^{-j}$$

This is the i^{th} row of M written in base n.

We recover the entire matrix in **one** query.

One Query Suffices

Suppose $M_{ij} \in \{0, \dots, n-1\}$ for some $n \geq 1$.

Consider query (p, q) with p arbitrary and $q_j \propto n^{-j}$. Then the i^{th} entry of the feedback (to the p player) is

$$\nabla_p f(p, q)_i = \sum_{j=1}^{K} M_{ij} q_j \propto \sum_{j=1}^{K} M_{ij} n^{-j}$$

This is the i^{th} row of M written in base n.

We recover the entire matrix in **one** query.

Theorem

One query suffices if entries M_{ij} in a known **countable** *set.*

Upshot

Ok, so we found some lame coding trick exploiting infinite-precision.

Upshot

Ok, so we found some lame coding trick exploiting infinite-precision.

Why is this **cool**?

Upshot

Ok, so we found some lame coding trick exploiting infinite-precision.

Why is this **cool**?

Existing lower bound techniques construct matrices with entries from a **finite alphabet**, and hence must be powerless.

- Nesterov-style lower bound (Ibrahim et al., 2020)
- Rademacher entries (Orabona and Pál, 2018)
- Reduction to hard combinatorial / submodular instance (Babichenko, 2016; Hart and Nisan, 2018)

Outline

Ingredients and main Ideas

The fully mixed case is **equalised** $Mq_\star = M^\intercal p_\star \propto \mathbf{1}$.

Ingredients and main Ideas

The fully mixed case is **equalised** $Mq_\star = M^\intercal p_\star \propto \mathbf{1}$.

We consider as our class of matrices perturbations of (scaled) identity I_K, for which saddle points are near uniform and fully mixed. So we do not hit the non-negativity constraints.

Ingredients and main Ideas

The fully mixed case is **equalised** $Mq_\star = M^\intercal p_\star \propto \mathbf{1}$.

We consider as our class of matrices perturbations of (scaled) identity I_K, for which saddle points are near uniform and fully mixed. So we do not hit the non-negativity constraints.

We show that if the span of the feedback includes $\mathbf{1}$, then the learner knows an exact equilibrium.

Ingredients and main Ideas

The fully mixed case is **equalised** $M q_\star = M^\intercal p_\star \propto \mathbf{1}$.

We consider as our class of matrices perturbations of (scaled) identity I_K, for which saddle points are near uniform and fully mixed. So we do not hit the non-negativity constraints.

We show that if the span of the feedback includes $\mathbf{1}$, then the learner knows an exact equilibrium.

To prove a lower bound, our job is to sequentially respond to the queries in a consistent way, while keeping $\mathbf{1}$ outside/**far from** the span of the feedback for as long as possible.

Technique

We work with matrices

$$B_0 \;=\; \mathcal{B}_{\|\cdot\|_{1,\infty}}\left(\frac{I_K}{2}, \frac{1}{16K^2}\right) \;=\; \left\{M \in [\pm 1]^{K \times K} \text{ s.t. } \left|M_{ij} - \frac{\delta_{i=j}}{2}\right| \le \frac{1}{16K^2}\right\}.$$

Any $M \in B_0$ has

- All equilibria of M are fully mixed
- Non-zero value $\min_p \max_q p^\intercal M q > 0$.

Main Idea

Consider t rounds with queries

$$(p_s, q_s)_{s \leq t}$$

and feedback

$$(\ell_s^{(p)}, \ell_s^{(q)})_{s \leq t}$$

Consistent matrices are

$$\mathcal{E}_t \;=\; \left\{ M \in B_0 \middle| M^\intercal p_s = \ell_s^{(q)} \text{ and } M q_s = \ell_s^{(p)} \text{ for all } s \leq t \right\}$$

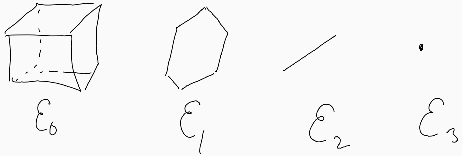

Key Insight

Lemma

Let (p, q) be a **common Nash equilibrium** *for all $M \in \mathcal{E}_t \neq \emptyset$. Then $p \in \mathrm{Span}(p_{1:t})$ and $q \in \mathrm{Span}(q_{1:t})$.*

Key Insight

Lemma

Let (p, q) be a **common Nash equilibrium** *for all $M \in \mathcal{E}_t \neq \emptyset$. Then $p \in \mathrm{Span}(p_{1:t})$ and $q \in \mathrm{Span}(q_{1:t})$.*

Proof.

Fix $M \in \mathcal{E}_t$. Let $\bar{p} = p - \mathrm{Proj}_{\mathrm{Span}(p_{1:t})}(p)$, and let $u_q \neq \mathbf{0}$ be orthogonal to $\mathrm{Span}(q_{1:t})$. Then $M' = M + \alpha \bar{p} u_q^\intercal \in \mathcal{E}_t$. But then

$$\mathbf{0} \;=\; (M - M')^\intercal p = \alpha(\bar{p}^\intercal p) u_q \;=\; \alpha \|\bar{p}\|^2 u_q$$

So $\bar{p} = \mathbf{0}$ and hence $p \in \mathrm{Span}(p_{1:t})$. □

Key Insight

Lemma

Let (p, q) *be a* **common Nash equilibrium** *for all* $M \in \mathcal{E}_t \neq \emptyset$*. Then* $p \in \mathrm{Span}(p_{1:t})$ *and* $q \in \mathrm{Span}(q_{1:t})$*.*

Proof.

Fix $M \in \mathcal{E}_t$. Let $\bar{p} = p - \mathrm{Proj}_{\mathrm{Span}(p_{1:t})}(p)$, and let $u_q \neq \mathbf{0}$ be orthogonal to $\mathrm{Span}(q_{1:t})$. Then $M' = M + \alpha \bar{p} u_q^\intercal \in \mathcal{E}_t$. But then

$$\mathbf{0} \;=\; (M - M')^\intercal p = \alpha(\bar{p}^\intercal p) u_q \;=\; \alpha \|\bar{p}\|^2 u_q$$

So $\bar{p} = \mathbf{0}$ and hence $p \in \mathrm{Span}(p_{1:t})$. □

Under same assumption, $\mathbf{1} \in \mathrm{Span}(\ell_{1:t}^{(p)}) \cap \mathrm{Span}(\ell_{1:t}^{(q)})$.

Keeping $\mathbf{1}$ from the span of the feedback

Theorem

For $T \le K/2 - 1$ rounds we can maintain $M_t \in \mathcal{E}_t$ s.t. $\mathbf{1} \notin \mathrm{Span}(\ell^{(q)}_{1:T})$.

Keeping 1 from the span of the feedback

Theorem

For $T \leq K/2 - 1$ rounds we can maintain $M_t \in \mathcal{E}_t$ s.t. $\mathbf{1} \notin \mathrm{Span}(\ell^{(q)}_{1:T})$.

By induction on t.

For the base case, we pick $M_0 = I_K/2 \in \mathcal{E}_0$.

Upon query p_{t+1} with fresh part $\bar{p}_{t+1} = p_{t+1} - \mathrm{Proj}_{\mathrm{Span}(p_{1:t})}(p_{t+1})$, set

$$M_{t+1} \; = \; M_t + \frac{\bar{p}_{t+1}}{\left\|\bar{p}_{t+1}\right\|^2} u_t^\intercal$$

where we pick u_t orthogonal to $\mathbf{1}$, as well as to

- $\mathrm{Span}(q_{1:t})$ (consistent with past feedback $\ell^{(p)}_t$)
- $\mathrm{Span}(\ell^{(q)}_{1:t})$ (proof artifact)
- $M_t^\intercal p_{t+1}$ (the threat)

The new feedback is $\ell^{(q)}_{t+1} = M_{t+1}^\intercal p_{t+1} = M_t^\intercal p_{t+1} + u_t$. If $\mathbf{1} = \sum_{s=1}^t \alpha_s \ell^{(q)}_s + \alpha_{t+1} \ell^{(q)}_{t+1}$ then $0 = \mathbf{1}^\intercal u_t = \alpha_{t+1} \|u_t\|$, so $\alpha_{t+1} = 0$.

Result

We can keep going until all dimensions are exhausted and we cannot pick u_t orthogonal to $\mathrm{Span}(q_{1:t}, \ell_{1:t}^{(q)}, \mathbf{1}, M_{t+1}^{\intercal} p_t)$ of $2t+2$ vectors. We obtain

Theorem

The query complexity of learning the exact $\epsilon = 0$ nash equilibrium in the first-order query model is $T \geq K/2 - 1$.

Quantitative case

Our result for approximate $\epsilon > 0$ equilibria is based on keeping $\left\| \mathbf{1} - \mathrm{Proj}_{\mathrm{Span}(\ell^{(q)}_{1:t})}(\mathbf{1}) \right\|$ big (instead of non-zero).

Quantitative case

Our result for approximate $\epsilon > 0$ equilibria is based on keeping $\left\| \mathbf{1} - \mathrm{Proj}_{\mathrm{Span}(\ell_{1:t}^{(q)})}(\mathbf{1}) \right\|$ big (instead of non-zero).

Lemma

We can ensure

$$\left\| \mathbf{1} - \mathrm{Proj}_{\mathrm{Span}(\ell_{1:T}^{(q)})}(\mathbf{1}) \right\|^2 \geq K \left(\frac{1}{8KT^2} \right)^{T+1}$$

Quantitative case

Our result for approximate $\epsilon > 0$ equilibria is based on keeping $\left\| \mathbf{1} - \mathrm{Proj}_{\mathrm{Span}(\ell_{1:t}^{(q)})}(\mathbf{1}) \right\|$ big (instead of non-zero).

Lemma

We can ensure

$$\left\| \mathbf{1} - \mathrm{Proj}_{\mathrm{Span}(\ell_{1:T}^{(q)})}(\mathbf{1}) \right\|^2 \geq K \left(\frac{1}{8KT^2} \right)^{T+1}$$

Lemma

Tuning all ingredients gives query complexity lower bound

$$T \geq \frac{\log \frac{1}{K^4 \epsilon}}{\log K} \wedge K - 3$$

Quantitative case

Our result for approximate $\epsilon > 0$ equilibria is based on keeping $\left\| \mathbf{1} - \mathrm{Proj}_{\mathrm{Span}(\ell_{1:t}^{(q)})}(\mathbf{1}) \right\|$ big (instead of non-zero).

Lemma

We can ensure

$$\left\| \mathbf{1} - \mathrm{Proj}_{\mathrm{Span}(\ell_{1:T}^{(q)})}(\mathbf{1}) \right\|^2 \;\geq\; K \left(\frac{1}{8KT^2} \right)^{T+1}$$

Lemma

Tuning all ingredients gives query complexity lower bound

$$T \;\geq\; \frac{\log \frac{1}{K^4 \epsilon}}{\log K} \wedge K - 3$$

Constant for $\epsilon = \frac{1}{K^c}$. **Insightful** e.g. when $\epsilon = K^{-K}$.

Outline

1. Motivation
2. Upper Bounds
3. New Lower Bounds
4. State of the Art
5. Conclusion

What we have discussed

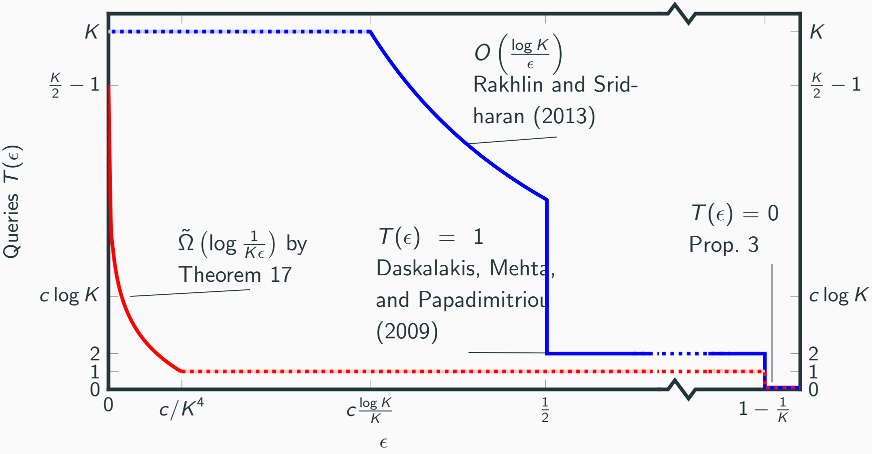

Outline

1. Motivation
2. Upper Bounds
3. New Lower Bounds
4. State of the Art
5. Conclusion

Conclusion

First non-trivial lower bounds.

Lots of open space.

Need even sharper techniques.

Thanks!

References i

- Babichenko, Y. (2016). **"Query complexity of approximate Nash equilibria"**. In: **Journal of the ACM (JACM)** 63.4.
- Brown, G. W. (1951). **"Iterative solution of games by fictitious play"**. In: **Act. Anal. Prod Allocation** 13.1.
- Daskalakis, C., A. Deckelbaum, and A. Kim (2011). **"Near-optimal no-regret algorithms for zero-sum games"**. In: **Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms**. SIAM.
- Daskalakis, C., A. Mehta, and C. Papadimitriou (2009). **"A note on approximate Nash equilibria"**. In: **Theoretical Computer Science** 410.17.
- Freund, Y. and R. E. Schapire (1999). **"Adaptive game playing using multiplicative weights"**. In: **Games and Economic Behavior** 29.1-2.
- Hart, S. and N. Nisan (2018). **"The query complexity of correlated equilibria"**. In: **Games and Economic Behavior** 108.

- Ibrahim, A., W. Azizian, G. Gidel, and I. Mitliagkas (2020). **"Linear Lower Bounds and Conditioning of Differentiable Games".** In: **Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event**. Vol. 119. Proceedings of Machine Learning Research.
- Orabona, F. and D. Pál (2018). **"Scale-free online learning".** In: **Theoretical Computer Science** 716. Special Issue on ALT 2015.
- Ouyang, Y. and Y. Xu (2021). **"Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems".** In: **Mathematical Programming** 185.1.
- Rakhlin, S. and K. Sridharan (2013). **"Optimization, learning, and games with predictable sequences".** In: **Advances in Neural Information Processing Systems** 26.
- Robinson, J. (1951). **"An iterative method of solving a game".** In: **Annals of mathematics**.