FOAM Seminar, April 19, 2024

Characterizing the First-order Query Complexity of Learning (Approximate) Nash Equilibria in Zero-sum Matrix Games

Tim van Erven

UNIVERSITY OF AMSTERDAM

Joint work with:

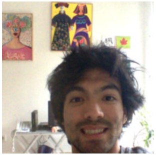

Hédi Hadiji

Sarah Sachs

Wouter Koolen

Thanks to Wouter for making many of the slides!

Outline

Background and Related Work

Identifying a Discrete Matrix is Too Easy

Continuous Matrices are Hard for Exact Nash Equilibria

Extension to Approximate Nash Equilibria

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

What is a solution?

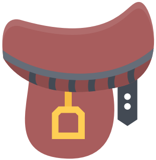

Given $\epsilon \geq 0$, we aim to find an **approximate saddle point / Nash equilibrium**

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

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_t, q_t)
- Learner receives **feedback** $(\nabla_{p_t} f(p_t, q_t), \nabla_{q_t} f(p_t, q_t))$

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_t, q_t)
- Learner receives **feedback** $(\nabla_{p_t} f(p_t, q_t), \nabla_{q_t} f(p_t, q_t))$

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

Query complexity

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

The most classical instance

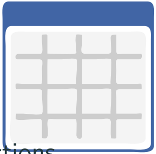

Consider **special case** of **zero-sum matrix games** (bilinear functions over probability simplex):

$$\min_{p \in \triangle_K} \max_{q \in \triangle_K} p^\intercal M q \qquad (M \in [-1, +1]^{K \times K})$$

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

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

The most classical instance

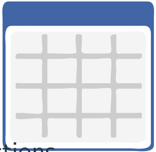

Consider **special case** of **zero-sum matrix games** (bilinear functions over probability simplex):

$$\min_{p \in \triangle_K} \max_{q \in \triangle_K} p^\intercal M q \qquad (M \in [-1,+1]^{K \times K})$$

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

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

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

Lower bounds remain elusive.

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

Where we are heading today

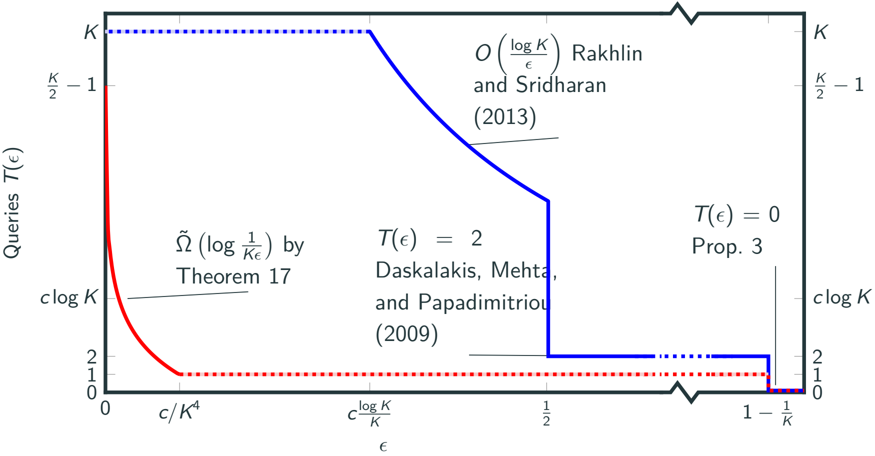

What is known: Upper Bounds

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

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

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

Assumptions on f and domains that exclude our setting:

2018: Ouyang and Xu (2021) show a lower bound on the query complexity for saddle-point problems **with curvature and rotationally invariant constraint sets**.

What is known: Lower Bounds

Assumptions on f and domains that exclude our setting:

2018: Ouyang and Xu (2021) show a lower bound on the query complexity for saddle-point problems **with curvature and rotationally invariant constraint sets**.

Harder query models:

2015: Fearnley et al. (2015) show lower bound when queries (i, j) return **single matrix entry** M_{ij}.

- Technique: construct hard binary matrix $M \in \{0, 1\}^{K \times K}$

2016: Hazan and Koren (2016) show lower bound when queries (p, q) return **best responses** $i^* \in \arg\min_i (Mq)_i$, $j^* \in \arg\max_j (M^\intercal p)_j$.

- Technique: Reduction from submodular optimization over the hypercube by encoding it as a binary matrix $M \in \{0, 1\}^{K \times K}$

Nothing for our setting!

Outline

Discrete entries are too easy!

Theorem (Identifying a Discrete Matrix)

One query suffices *to fully identify M if the entries M_{ij} come from a known countable alphabet.*

- E.g. $M_{ij} \in \{-1, +1\}$
- Implies query complexity is $T(\epsilon) \leq 1$ if we restrict to discrete M!

Discrete entries are too easy!

Theorem (Identifying a Discrete Matrix)

One query suffices *to fully identify M if the entries M_{ij} come from a known countable alphabet.*

- E.g. $M_{ij} \in \{-1, +1\}$
- Implies query complexity is $T(\epsilon) \leq 1$ if we restrict to discrete M!

Rules out all existing lower bound techniques. For instance:

- Hard binary matrix (Fearnley et al., 2015)
- Encoding submodular optimization as binary matrix (Hazan and Koren, 2016)
- Randomly generating a matrix with binary entries (Orabona and Pál, 2018)

Proof Idea: One Query Suffices

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

Proof Idea: 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

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

Proof Idea: 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**.

Proof Idea: 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**.

Proof Idea: 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**.

But query is **very artificial** and fails under numerical imprecision. Should we restrict the query model to only allow more realistic queries?

Proof Idea: 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**.

But query is **very artificial** and fails under numerical imprecision. Should we restrict the query model to only allow more realistic queries? **No!**

Outline

Continuous Matrices are Hard

Theorem (Identifying a Continuous Matrix)

If the entries in M can take any values in $[-1, +1]$, then the number of queries required to fully identify M is ***exactly*** *K.*

- As hard as querying each row/column in turn
- Compare to: 1 query if M is discrete
- Proof approach: carefully count the number of linear constraints imposed by the queries.

Continuous Matrices are Hard

Theorem (Identifying a Continuous Matrix)

If the entries in M can take any values in $[-1, +1]$, then the number of queries required to fully identify M is ***exactly*** *K.*

- As hard as querying each row/column in turn
- Compare to: 1 query if M is discrete
- Proof approach: carefully count the number of linear constraints imposed by the queries.

Theorem (Query Complexity for Exact Equilibria)

The number of queries required to compute an exact Nash equilibrium is at least $T(0) \geq \frac{K}{2} - 1$.

- Essentially as hard as identifying the full matrix!

Proof Ingredients and Main Ideas

Idea: construct adversary **answering** queries by the learner so as to **delay revealing the equilibrium** for as long as possible.

1. Based on the feedback given so far, a subset of consistent matrices remains: every round adds $\leq 2K$ equality constraints.

Proof Ingredients and Main Ideas

Idea: construct adversary **answering** queries by the learner so as to **delay revealing the equilibrium** for as long as possible.

1. Based on the feedback given so far, a subset of consistent matrices remains: every round adds $\leq 2K$ equality constraints.
2. Restrict a priori to nice subset B_0 of matrices M for which the Nash equilibrium $(p^\star, q^\star)$ are **fully mixed**, i.e. have full support. Then they are **equalizer strategies**:

$$Mq_\star = M^\intercal p_\star \propto \mathbf{1}.$$

Proof Ingredients and Main Ideas

Idea: construct adversary **answering** queries by the learner so as to **delay revealing the equilibrium** for as long as possible.

1. Based on the feedback given so far, a subset of consistent matrices remains: every round adds $\leq 2K$ equality constraints.
2. Restrict a priori to nice subset B_0 of matrices M for which the Nash equilibrium $(p^\star, q^\star)$ are **fully mixed**, i.e. have full support. Then they are **equalizer strategies**:

$$Mq_\star = M^\intercal p_\star \propto \mathbf{1}.$$

3. Lemma: the learner **knows** an exact equilibrium **only if** the span of the feedback includes $\mathbf{1}$.

Proof Ingredients and Main Ideas

Idea: construct adversary **answering** queries by the learner so as to **delay revealing the equilibrium** for as long as possible.

1. Based on the feedback given so far, a subset of consistent matrices remains: every round adds $\leq 2K$ equality constraints.
2. Restrict a priori to nice subset B_0 of matrices M for which the Nash equilibrium $(p^\star, q^\star)$ are **fully mixed**, i.e. have full support. Then they are **equalizer strategies**:

$$Mq_\star = M^\intercal p_\star \propto \mathbf{1}.$$

3. Lemma: the learner **knows** an exact equilibrium **only if** the span of the feedback includes $\mathbf{1}$.
4. Our adversary keeps $\mathbf{1}$ out of the span of the feedback for $\frac{K}{2} - 1$ rounds. $\Leftarrow$ "dimension-as-a-resource"

1. Consistent Matrices

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

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

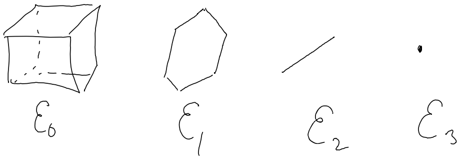

2. Subset B_0 of Nice Matrices

Before we start, we commit that M will be in

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

Any $M \in B_0$ satisfies:

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

3. Known Equilibrium Lemma

Lemma

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

3. Known Equilibrium Lemma

Lemma

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

The learner **knows** an exact equilibrium
only if the span of the feedback includes $\mathbf{1}$:

Corollary

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

3. Known Equilibrium Lemma

Lemma

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

The learner **knows** an exact equilibrium
only if the span of the feedback includes $\mathbf{1}$:

Corollary

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

Proof.

$(p^\star, q^\star)$ fully mixed because $\mathcal{E}_t \subset B_0$. Hence exists $v > 0$ such that

- $\mathbf{1} = vM^\intercal p_\star \in M^\intercal \mathrm{Span}(p_{1:t}) = \mathrm{Span}(\ell^{(q)}_{1:t})$
- $\mathbf{1} = vMq_\star \in M\,\mathrm{Span}(q_{1:t}) = \mathrm{Span}(\ell^{(p)}_{1:t})$

□

4. Keeping $\mathbf{1}$ 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. $\mathbf{1} \notin \mathrm{Span}(\ell_{1:T}^{(q)})$.

4. Keeping $\mathbf{1}$ 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 non-zero 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 $\text{Span}(q_{1:t}, \ell_{1:t}^{(q)}, \mathbf{1}, M_{t+1}^\intercal p_t)$ of $2t+2$ vectors. We obtain

Theorem (Query Complexity for Exact Equilibria)

The number of queries required to compute an exact ($\epsilon = 0$) Nash equilibrium is at least $T(0) \geq \frac{K}{2} - 1$.

Outline

Background and Related Work

Identifying a Discrete Matrix is Too Easy

Continuous Matrices are Hard for Exact Nash Equilibria

Extension to Approximate Nash Equilibria

Approximate Nash Equilibria

The same approach extends from $\epsilon = 0$ to small $\epsilon > 0$:

Theorem (Approximate Nash Equilibria)

The number of queries required to compute a Nash equilibrium for any $\epsilon \leq 1/(e\,2^{10}K^4)$ is at least

$$T(\epsilon) \geq \Big(\frac{-\log(2^{10}K^4\epsilon)}{\log(2^{11/2}K^{5/2}) + \log(-\log(2^{10}K^4\epsilon))} - 1\Big) \wedge \Big(\frac{K}{2} - 1\Big)$$
$$= \tilde{\Omega}\Big(\log\frac{1}{K\epsilon}\Big)$$

Approximate Nash Equilibria

The same approach extends from $\epsilon = 0$ to small $\epsilon > 0$:

Theorem (Approximate Nash Equilibria)

The number of queries required to compute a Nash equilibrium for any $\epsilon \leq 1/(e\,2^{10}K^4)$ *is at least*

$$T(\epsilon) \geq \Big(\frac{-\log(2^{10}K^4\epsilon)}{\log(2^{11/2}K^{5/2}) + \log(-\log(2^{10}K^4\epsilon))} - 1\Big) \wedge \Big(\frac{K}{2} - 1\Big)$$
$$= \tilde{\Omega}\Big(\log\frac{1}{K\epsilon}\Big)$$

Proof approach: Need to keep

$$\mathrm{dist}\Big(\mathbf{1}, \mathrm{Span}(\ell_{1:t}^{(q)})\Big)$$

large enough, instead of only non-zero.

Summary

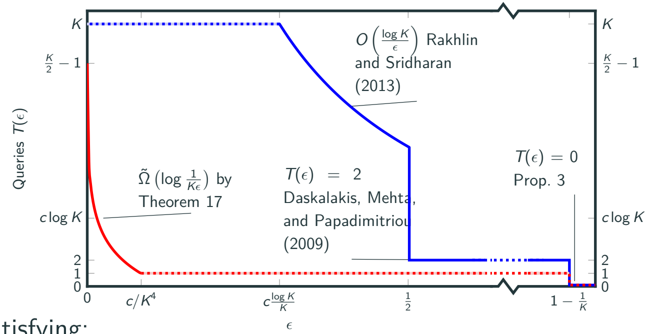

Very satisfying:

- Prior lower bound techniques cannot work, because **discrete matrices are too easy**: 1 query suffices to identify M
- Identifying continuous M is hard: requires K queries
- Computing **exact Nash equilibrium is hard**: $T(0) \geq \frac{K}{2} - 1$

Far from solved:

- For tiny ϵ, we have a **first non-trivial lower bound** on the query complexity $T(\epsilon)$, but it is far from the upper bounds

References i

- H. Hadiji, S. Sachs, T. van Erven and W. M. Koolen. **Towards Characterizing the First-order Query Complexity of Learning (Approximate) Nash Equilibria in Zero-sum Matrix Games.** NeurIPS, 2023.

Brown, George W (1951). **"Iterative solution of games by fictitious play".** In: *Act. Anal. Prod Allocation* 13.1, p. 374.

Daskalakis, Constantinos, Alan Deckelbaum, and Anthony 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, pp. 235–254.

Daskalakis, Constantinos, Aranyak Mehta, and Christos Papadimitriou (2009). **"A note on approximate Nash equilibria".** In: *Theoretical Computer Science* 410.17, pp. 1581–1588.

Fearnley, John, Martin Gairing, Paul W Goldberg, and Rahul Savani (2015). **"Learning Equilibria of Games via Payoff Queries".** In: *Journal of Machine Learning Research* 16, pp. 1305–1344.

References ii

Freund, Yoav and Robert E Schapire (1999). **"Adaptive game playing using multiplicative weights".** In: *Games and Economic Behavior* 29.1-2, pp. 79–103.

Hazan, Elad and Tomer Koren (2016). **"The Computational Power of Optimization in Online Learning".** In: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '16, pp. 128–141.

Orabona, Francesco and Dávid Pál (2018). **"Scale-free online learning".** In: *Theoretical Computer Science* 716. Special Issue on ALT 2015, pp. 50–69.

Ouyang, Yuyuan and Yangyang Xu (2021). **"Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems".** In: *Mathematical Programming* 185.1, pp. 1–35.

Rakhlin, Sasha and Karthik Sridharan (2013). **"Optimization, learning, and games with predictable sequences".** In: *Advances in Neural Information Processing Systems* 26.

References iii

Robinson, Julia (1951). **"An iterative method of solving a game".** In: *Annals of mathematics*, pp. 296–301.