Differentially Private Equilibrium Finding in Polymatrix Games

Mingyang Liu, Gabriele Farina, Asuman Ozdaglar

Massachusetts Institute of Technology

LIDS

ICLR

Preliminaries

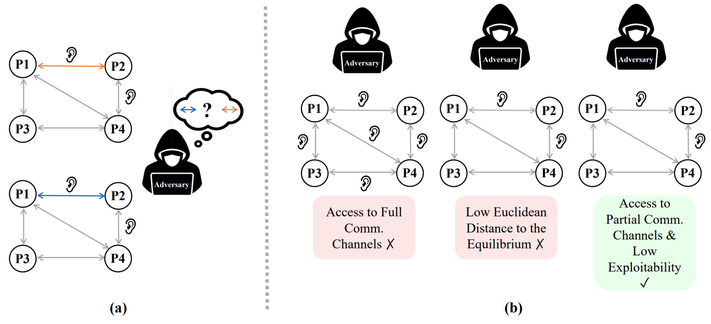

Figure 1 (a): An illustration of adjacent polymatrix games defined in Definition 2.1. The nodes represent the players, and the lines represent the edges of the polymatrix game. The two games differ at the blue/orange edge, and the adversary cannot differentiate these two games with certainty from his observations. 👂 on an edge means that the communication channels between those two players can be accessed by the adversary.

Polymatrix Games. Polymatrix games can be written as a tuple

$$\mathcal{G} := \left([N], E, \{\mathcal{A}_i\}_{i\in[N]}, \{\boldsymbol{U}_{i,j}\}_{(i,j)\in E}\right),$$

where

- $[N]$ is the set of players.
- E is the set of edges, where each $(i,j) \in E$ indicates that players i, j interact with each other.
- $\{\mathcal{A}_i\}_{i\in[N]}$ is the set of players' action set, which means player $i \in [N]$ chooses actions in $\mathcal{A}_i$.
- Let $A := \max_{i\in[N]} |\mathcal{A}_i|$ be the size of the largest action set.
- $\boldsymbol{U}_{i,j} \in [-1,1]^{\mathcal{A}_i \times \mathcal{A}_j}$ is the utility matrix between player $(i,j) \in E$.

Definition 2.1 (Game adjacency)**.** Given two polymatrix games $\mathcal{G} = \left([N], E, \{\mathcal{A}_i\}_{i\in[N]}, \{\boldsymbol{U}_{i,j}\}_{(i,j)\in E}\right)$ and $\mathcal{G}' = \left([N'], E', \{\mathcal{A}'_i\}_{i\in[N']}, \{\boldsymbol{U}'_{i,j}\}_{(i,j)\in E'}\right)$, they are said to be *adjacent*, indicated as $\mathcal{G} \sim \mathcal{G}'$, if

1. $N = N'$, $E = E'$ and $\mathcal{A}_i = \mathcal{A}'_i$ for any $i \in [N]$; and
2. except for an edge $(i,j) \in E$, $\boldsymbol{U}_{i',j'} = \boldsymbol{U}'_{i',j'}$ and $\boldsymbol{U}_{j',i'} = \boldsymbol{U}'_{j',i'}$ for any $(i',j') \in E \setminus \{(i,j),(j,i)\}$.

Definition 2.2 ((ϵ,δ)-Differential Privacy)**.** For an $\epsilon \geq 0$ and $\delta \geq 0$, an iterative distributed algorithm for finding equilibria is (ϵ,δ)-differentially private, if and only if for any two adjacent polymatrix game $\mathcal{G}, \mathcal{G}'$, any timestep $t > 0$ and any set of observations $\mathcal{S} \subseteq \mathcal{O}^t$,

$$\mathcal{P}_{\mathcal{G}}\left(\{\boldsymbol{\theta}\colon \mathcal{R}_{\mathcal{G}}(\boldsymbol{\theta}) \in \mathcal{S}\}\right) \leq e^{\epsilon}\mathcal{P}_{\mathcal{G}'}\left(\{\boldsymbol{\theta}\colon \mathcal{R}_{\mathcal{G}'}(\boldsymbol{\theta}) \in \mathcal{S}\}\right) + \delta. \tag{1}$$

$\mu_{\mathcal{G}}^{(t)}$ is the distribution over the adversary's observation at timestep t

Definition 2.3 ((α,ϵ)-Rényi Differential Privacy)**.** For $\alpha > 1$ and $\epsilon \geq 0$, an iterative distributed algorithm for finding equilibria is (α,ϵ)-Rényi differentially private, if and only if for any two adjacent polymatrix game $\mathcal{G}, \mathcal{G}'$ and timestep $t > 0$,

$$D_\alpha\left(\mu_{\mathcal{G}}^{(t)}, \mu_{\mathcal{G}'}^{(t)}\right) := \frac{1}{\alpha-1}\log \mathbb{E}_{\boldsymbol{o}\sim\mu_{\mathcal{G}'}^{(t)}}\left[\left(\frac{\mu_{\mathcal{G}}^{(t)}(\boldsymbol{o})}{\mu_{\mathcal{G}'}^{(t)}(\boldsymbol{o})}\right)^{\alpha}\right] \leq \epsilon, \tag{2}$$

where ϵ is also called the *privacy budget*.

Lemma 2.4. *If an algorithm satisfies (α,ϵ)-Rényi DP for $\alpha > 1$, then for any $\delta \in (0,1)$, the algorithm also satisfies $\left(\epsilon + \frac{\log(1/\delta)}{\alpha-1}, \delta\right)$-DP. Moreover, (∞,ϵ)-Rényi DP is equivalent to $(\epsilon,0)$-DP.*

Impossibility Results

- Suppose the algorithm outputs $\pi = (\pi_1, \pi_2, \ldots, \pi_N)$ as the equilibrium
- The distribution over the adversary's observation is $\mu_{\mathcal{G}}$

$$\text{Accuracy: } \frac{1}{N}\sum_{i=1}^{N}\mathbb{E}\left[\max_{\widehat{\pi}_i\in\Delta^{\mathcal{A}_i}}\left\langle \pi_i - \widehat{\pi}_i, \boldsymbol{g}_i^{\boldsymbol{\pi}}\right\rangle\right] \leq \zeta$$

$$\text{Privacy: } D_\alpha\left(\mu_{\mathcal{G}}, \mu_{\mathcal{G}'}\right) \leq \epsilon \quad \forall \mathcal{G} \sim \mathcal{G}'.$$

Lemma 3.1. *For any $N \geq 12$, there exists two zero-sum adjacent polymatrix games with N players so that for any algorithm guaranteeing* (3) *and* (4)*, we have*

$$\zeta \geq \min\left\{\frac{3\exp(-2\epsilon)}{112}, \frac{1}{112}\right\}. \tag{5}$$

- Suppose the distribution over the adversary's local observation on player i's information channel is $\mu_{\mathcal{G},i}$

$$\text{Accuracy: } \frac{1}{N}\sum_{i=1}^{N}\mathbb{E}\left[\left\|\pi_i - \mathrm{Proj}_{\Delta^{\mathcal{A}_i},*}(\pi_i)\right\|^2\right] \leq \zeta$$

$$\text{Privacy: } \frac{1}{N}\sum_{i=1}^{N} D_\alpha\left(\mu_{\mathcal{G},i}, \mu_{\mathcal{G}',i}\right) \leq \epsilon \quad \forall \mathcal{G} \sim \mathcal{G}'.$$

Lemma 3.2. *For any $N \geq 8$, there exists two zero-sum adjacent polymatrix games with N players so that for any algorithm guaranteeing* (6) *and* (7)*, then*

$$\zeta \geq \min\left\{\frac{3}{8}\exp(-4\epsilon), \frac{1}{16}\right\}. \tag{8}$$

Algorithm

Algorithm 1 Differentially Private CCE Computation in Polymatrix Games

Input: Player index i
Initialize $\pi_i^{(0)}$ as uniform distribution over $\mathcal{A}_i$
Let $\overline{\mathcal{N}} := N \cdot (\sum_{i=1}^{N} \frac{1}{|\mathcal{N}(i)|})^{-1}$ be the harmonic mean of players' degrees
Initialize $\tau_i \leftarrow \frac{(\overline{\mathcal{N}})^{5/9}}{|\mathcal{N}(i)|\log N}$
for $t = 0, 1, \ldots, T$ **do**
  Sample $\boldsymbol{n}_i^{(t)} \sim \mathcal{N}\left(\boldsymbol{0}, \sigma^2 \boldsymbol{I}^{\mathcal{A}_i}\right)$
  Broadcast $\pi_i^{(t)} + \boldsymbol{n}_i^{(t)}$ to neighbors $j \in \mathcal{N}(i)$
  for $j \in \mathcal{N}(i)$ **do**
    Receive $\pi_j^{(t)} + \boldsymbol{n}_j^{(t)}$ and compute $\overline{\pi}_j^{(t)} \leftarrow \mathrm{Proj}_{\Delta^{\mathcal{A}_j}}\left(\pi_j^{(t)} + \boldsymbol{n}_j^{(t)}\right)$
  end for
  Gradient $\overline{\boldsymbol{g}}_i^{(t)} \leftarrow -\frac{1}{|\mathcal{N}(i)|}\sum_{j\in\mathcal{N}(i)} \boldsymbol{U}_{i,j}\overline{\pi}_j^{(t)}$

$$\pi_i^{(t+1)} \leftarrow \operatorname*{argmin}_{\pi_i\in\Delta^{\mathcal{A}_i}} \left\langle \pi_i, \overline{\boldsymbol{g}}_i^{(t)}\right\rangle + \tau_i\|\pi_i\|^2 + \frac{1}{\eta}\left\|\pi_i - \overline{\pi}_i^{(t)}\right\|^2. \tag{9}$$

{We remark that (9) is equivalent to $\pi_i^{(t+1)} \leftarrow \mathrm{Proj}_{\Delta^{\mathcal{A}_i}}\left(\frac{\overline{\pi}_i^{(t)} - \eta\overline{\boldsymbol{g}}_i^{(t)}}{1+\eta\tau_i}\right)$.}
end for

Intuition

- Use adaptive regularizers whose magnitude is **inversely proportional** to each player's degree.
- When a player has low degree, its strategy becomes more sensitive to changes in the utility matrix between neighboring games.
- Therefore, a stronger regularizer is needed to stabilize the strategy and make it harder for the adversary to distinguish between the games.

Convergence Guarantee

Dense Graph: $\overline{\mathcal{N}} \geq N^p$

$$\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\mathbb{E}\left[\left\langle \boldsymbol{g}_i^{\boldsymbol{\pi}^{(t+1)}}, \pi_i^{(t+1)} - \pi_i\right\rangle\right] \leq \mathcal{O}\left(\frac{(\log N)^{2/3}}{N^{4p/27}}\right),$$

$$\frac{1}{N}\sum_{i=1}^{N} D_\alpha\left(\mu_{\mathcal{G},i}^{(T)}, \mu_{\mathcal{G}',i}^{(T)}\right) \leq \mathcal{O}\left(\frac{\alpha(\log N)^{2/3}}{N^{4p/27}}\right).$$

Sparse Graph: Every player's degree is constant

$$\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\mathbb{E}\left[\left\langle \boldsymbol{g}_i^{\boldsymbol{\pi}^{(t+1)}}, \pi_i^{(t+1)} - \pi_i\right\rangle\right] \leq \mathcal{O}\left(\frac{1}{(\log N)^{1/3}}\right)$$

$$\frac{1}{N}\sum_{i=1}^{N} D_\alpha\left(\mu_{\mathcal{G},i}^{(T)}, \mu_{\mathcal{G}',i}^{(T)}\right) \leq \mathcal{O}\left(\frac{\alpha}{(\log N)^{1/3}}\right).$$

Experiments

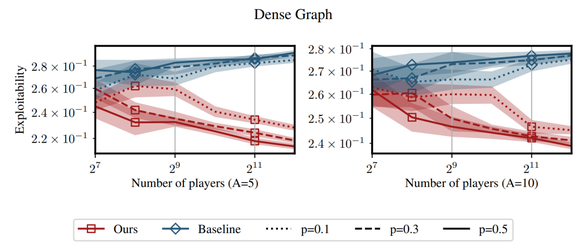

Figure 3: We evaluate the exploitability of our algorithm and a baseline on dense graphs. Each node (player) connects to another node independently with probability p, and duplicate edges are then removed. All players have action sets of size A. Both the baseline and our method are run under the same differential privacy budget.