Graphical games

Graph-theoretic models for multiplayer games

Geelon So, agso@eng.ucsd.edu
Algorithmic game theory reading group — November 9, 2021

Graphical games

Graphical games capture games with:

- large number of players
- direct influences that are local/sparse
  - payoffs of each player can be determined locally (actions of adjacent players and itself)

Traffic example

Consider all drivers on the road in San Diego.

- The payoff for each driver depends on the action of cars immediately around them.
- Effects like traffic or accidents can propagate beyond the immediate neighborhoods.

Multiplayer game

Consider a **multiplayer game** with n players with:

- **binary action space**: $a_i \in \{0, 1\}$ is the action of the ith player
- **payoff matrix**: $M_i(\mathbf{a}) \in [0, 1]$ is the payoff for player i
  - the joint action is given by $\mathbf{a} = (a_1, \ldots, a_n)$

Model

A **graphical game** with n players consists of:

- a graph G with n vertices
- local payoff matrices $M_1, \ldots, M_n$
  - given a player i, let $N(i)$ be its **neighborhood** (set of adjacent players and itself)
  - a payoff matrix M_i is **local** if $M_i(\mathbf{a}) = M_i(\mathbf{a}')$ for all $\mathbf{a}\big|_{N(i)} = \mathbf{a}'\big|_{N(i)}$

Nash equilibria

- A **mixed strategy** for player i is given by the probability $p_i \in [0, 1]$ of choosing 1.
- A **joint mixed strategy** $\mathbf{p} = (p_1, \ldots, p_n)$ describes a product distribution, where each player chooses their actions independently.
- The **expected payoff** to player i is given by:

$$M_i(\mathbf{p}) = \mathop{\mathbb{E}}_{\mathbf{a} \sim \mathbf{p}} \left[M_i(\mathbf{a}) \right].$$

- A **Nash equilibrium** is a joint mixed strategy $\mathbf{p}$ such that for any $\mathbf{p}'$ differing only in one coordinate i (so that $p_j = p'_j$ for $j \neq i$),

$$M_i(\mathbf{p}) \geq M_i(\mathbf{p}').$$

Traffic example

Question: what would a Nash equilibria mean for the traffic example?

Correlated equilibria

- More generally, the players may have random but correlated actions.
- A distribution P over the joint action space $\{0,1\}^n$ is a **correlated equilibrium** if for each player i and binary action $b \in \{0,1\}$,

$$\mathop{\mathbb{E}}_{\mathbf{a}\sim P}\left[M_i(\mathbf{a}) \,\middle|\, a_i = b\right] \geq \mathop{\mathbb{E}}_{\mathbf{a}\sim P}\left[M_i(\mathbf{a}[i:\neg b]) \,\middle|\, a_i = b\right],$$

  where $\mathbf{a}[i:\neg b]$ sets the ith coordinate of $\mathbf{a}$ to $\neg b$.
  - That is, the player is not incentivized to deviate from its jointly chosen action.
  - Mechanistically, a trusted party draws a joint action $\mathbf{a}$ according to P and privately distributes to each player i only their component a_i.

Traffic example

A traffic light could be informally understood as shared/public randomness.

- This allows player to coordinate their randomness.
- Since no player gains a greater payoff by unilaterally running a light, following traffic rules can be understood as a correlated equilibrium.

Problems of interest

1. Nash and correlated equilibria are models of ‘rational behavior’. We would like to compute or approximate them efficiently leveraging topology.
    - “If your PC cannot find the equilibria, then neither can the market.” Kamal Jain
    - Even *representing* a correlated equilibrium may require bits exponential in player
2. Understand relationship between network structures and strategic outcomes.
3. Develop connections with graphical models
    - “In probibalistic inference the interactions are stochastic, whereas in graphical games they are strategic (best response).” Michael Kearns (Kearns, 2007)

Computing Nash Equilibria

Tree graphical games

Consider a graphical game where the graph is a tree.

- Without loss of generality, assume the tree has a **root**.
  - Remaining vertices are either **internal nodes** or **leaves**.
  - Given a vertex V, any vertex on the path from V down to the root is **downstream**
  - Any vertex on the path from V up to a leaf is **upstream**
- Given a vertex V, define family of **subgames** G_u^V on V and its upstream vertices:
  - if U is the child of V (if it exists), then G_u^V corresponds to the game induced by fixing the action of player U to mixed strategy u
  - this is well-defined because of the locality of the payoff matrices
- We can find equilibria by a belief propagation-style algorithm.

Finding Nash equilibria: TreeNash

1. start upstream (leaves) and inductively work way downstream (root)
2. each vertex can have parents upstream and at most a single child downstream
3. for a vertex V and child U, define map $T_{V,U}(v, u) \in \{0, 1\}$ so that:

$$T_{V,U}(v, u) = 1 \quad \Leftrightarrow \quad \exists \text{ Nash equilibrium on } G_u^V \text{ where } V \text{ plays } v \text{ and } U \text{ plays } u,$$

where $v, u \in [0, 1]$ are mixed strategies.

- each vertex V receives $T_{P_1,V}, \ldots, T_{P_k,V}$ from its parents
- the vertex V computes $T_{V,U}$ by:

$$T_{V,U}(v, u) = \begin{cases} 1 & \exists (p_1, \ldots, p_k, v, u) \text{ is a local Nash equilibrium} \\ 0 & \text{o.w.} \end{cases}$$

where $(p_1, \ldots, p_k, v, u)$ is a local Nash equilibrium if v is optimal mixed strategy for V given that each parent P_i plays p_i and U plays u, and if $T_{P_i,V}(p_i, v) = 1$ for all i.

Approximation algorithm

For each player, discretize the mixed strategy space $[0,1]^n$ by a τ-grid:

$$\{0, \tau, 2\tau, \ldots, 1\}^n.$$

Lemma

Let G be a graph with maximal degree d, and let $(g, M_1, \ldots, M_n)$ be a graphical game. Let $\mathbf{p}$ be a Nash equilibrium and $\mathbf{q}$ be the nearest mixed strategy on the τ-grid in L_1-distance. Then $\mathbf{q}$ is a $d\tau$-approximate Nash equilibrium.

- A joint strategy $\mathbf{q}$ is an ε-approximate Nash equilibrium if for each player i, deviating from strategy q_i improves their payoff by at most ε.

Complexity of ε-approximation algorithm

- To store $T_{V,U}$ we need $(1/\tau)^2$ bits of memory.
- To compute $T_{V,U}$, we need to query each entry of $T_{P_i,V}$, using $(1/\tau)^{2k}$ lookups.
- Running time of algorithm polynomial in $1/\varepsilon$, n, 2^d

Exact algorithm

- It turns out that $T_{V,U} : [0,1]^2 \to \{0,1\}$ can be represented as a finite union of axis-aligned rectangular regions.
  - The number of regions multiplies with increasing depth.
  - Worst-case bound on number of regions for root is $\exp(n)$.
- There is an exponential time algorithm computing all exact Nash equilibria.

Nash equilibria on general graphs

- Need to compute $T_{V,U}$ for each edge (V, U)
  - Even for approximate version, no guarantee to be polynomial time
- The algorithm finds a superset of all (approximate) Nash equilibria
  - Even if $T_{V,U}(p_v, p_u) = 1$ for all edges, $\mathbf{p}$ is not necessarily a Nash equilibria
  - Paring down superset may be computationally expensive

Representing and computing correlated equilibria

The representation problem

Because a correlated equilibrium P is no longer a joint distribution, even writing it down may take exponentially many bits.

- A mixture of Nash equilibria is always a correlated equilibria; there are simple games with exponentially many NE.
- It turns out that if P is a CE, then there is a Markov random field Q on the graph equivalent to P. Such Q's can be efficiently represented.

Expected payoff equivalence

Definition

Two distributions P and Q over joint actions $\mathbf{a} \in \{0,1\}^n$ are ***expected payoff equivalent*** *$P \equiv_{\mathrm{EP}} Q$ if the expected payoffs for each player is the same:*

$$\mathop{\mathbb{E}}_{\mathbf{a}\sim P}\left[M_i(\mathbf{a})\right] = \mathop{\mathbb{E}}_{\mathbf{a}\sim Q}\left[M_i(\mathbf{a})\right].$$

Local network equivalence

Definition

*Given graph G, two distributions P and Q over joint actions are **local neighborhood equivalent** $P \equiv_{\mathrm{LN}} Q$ with respect to the graph G if for each player i, the marginal distributions of P and Q on its neighborhood $N(i)$ are equal.*

- Notice that LN equivalence depends only on the graph structure.

LN is stronger than EP equivalence

Lemma

- *If* $P \equiv_{\mathrm{LN}} Q$*, then* $P \equiv_{\mathrm{EP}} Q$*.*
- *If* $P \not\equiv_{\mathrm{LN}} Q$*, then there exists payoff matrices* $M_1, \ldots, M_n$ *such that* $P \not\equiv_{\mathrm{EP}} Q$*.*

LN preserves CE

Lemma

For any graphical game $(G, M_1, \ldots, M_n)$, *if* P *is a correlated equilibrium and if* $P \equiv_{\mathrm{LN}} Q$, *then* Q *is also a correlated equilibrium.*

All CE are LN equivalent to an MRF on G

Theorem

For all graphical games $(G, M_1, \ldots, M_n)$, and for any correlated equilibrium P, there exists distribution Q such that:

(i) *Q is also a correlated equilibrium*

(ii) *Q is expected payoff equivalent to P*

(iii) *Q can be represented as a Markov random field on G*

Note: an MRF on G has representation complexity linear in the size of $G = (V, E)$.

Markov random fields

Definition

*A **Markov random field** is a pair (G, Ψ) where:*

- *G is an **undirected graph** on n vertices*
- *Ψ is a set of **potential functions** for each local neighborhood $N(i)$,*

$$\psi_i : \{0, 1\}^{N(i)} \to [0, \infty).$$

An MRF defines a probability distribution by:

$$P(\mathbf{a}) \equiv \frac{1}{Z} \left(\prod_{i=1}^{n} \psi_i(\mathbf{a}_{N(i)}) \right),$$

where Z is a normalization and $\mathbf{a}_{N(i)}$ is the projection of $\mathbf{a}$ onto coordinates in $N(i)$.

LN equivalence to MRFs

Lemma

Let P be a joint distribution over $\{0, 1\}^G$. There exists a distribution Q representable as a MRF on G such that $P \equiv_{\text{LN}} Q$ with respect to G.

Proof idea.

Max entropy distribution with the same marginal distributions is in an exponential family factorizing over the graph. □

Computing correlated equilibria on trees

We just need to determine the marginal distributions on each neighborhood $N(i)$. For every player i and assignment $\mathbf{a}_{N(i)}$, define the variable $P^{(i)}(\mathbf{a}_{N(i)})$. Consider the LP:

1. *CE constraints:* for all players i and action $b \in \{0, 1\}$,

$$\sum_{\mathbf{a}_{N(i)}:a_i=b} P^{(i)}(\mathbf{a}_{N(i)}) M_i(\mathbf{a}_{N(i)}) \geq \sum_{\mathbf{a}_{N(i)}:a_i=b} P^{(i)}(\mathbf{a}_{N(i)}) M_i\big(\mathbf{a}_{N(i)}[i : \neg b]\big).$$

2. *Marginal constraints:* for all players i,

$$\forall \mathbf{a}_{N(i)}, \quad P^{(i)}(\mathbf{a}_{N(i)}) \geq 0 \qquad \text{and} \qquad \sum_{\mathbf{a}_{N(i)}} P^{(i)}(\mathbf{a}_{N(i)}) = 1.$$

3. *Intersection consistency constraints:* for all players i and j,

$$P^{(i)}(\mathbf{a}_{N(i)\cap N(j)}) = P^{(j)}(\mathbf{a}_{N(i)\cap N(j)}).$$

Solutions of LP are CE

Theorem

Given a tree graphical game, any solution to the above LP is a correlated equilibrium.

- If the tree has degree d, this LP has $O(n2^d)$ variables and $O(n2^d)$ linear inequalities.

References

Michael Kearns. Graphical games. *Algorithmic game theory*, 3:159–180, 2007.