L8 Other notions of equilibria

CS 280 Algorithmic Game Theory

Ioannis Panageas

Relaxing Nash equilibrium

- NASH is computationally hard.

Question: Are there other equilibrium notions that are computationally tractable?

Relaxing Nash equilibrium

- NASH is computationally hard.

Question: Are there other equilibrium notions that are computationally tractable?

Answer: Correlated equilibria, i.e., relaxing the product distribution assumption.

Example (Correlated eq.)

	Chicken-out	Dare
Chicken-out	**0, 0**	**-2, 1**
Dare	**1, -2**	**-10, -10**

Suppose agents are recommended $(C, D), (D, C), (C, C)$ with probability $\frac{1}{3}$ each.

Example (Correlated eq.)

	Chicken-out	Dare
Chicken-out	**0, 0**	**-2, 1**
Dare	**1, -2**	**-10, -10**

Suppose agents are recommended $(C, D), (D, C), (C, C)$ with probability $\frac{1}{3}$ each.

- If agent row is recommended to choose C, then column is recommended to play C or D with equal probability. Expected payoff of row is $\frac{1}{2} \cdot 0 + \frac{1}{2}(-2) = -1$ which is greater than switching to D (expected payoff is -4.5).

Example (Correlated eq.)

	Chicken-out	Dare
Chicken-out	**0, 0**	**-2, 1**
Dare	**1, -2**	**-10, -10**

Suppose agents are recommended $(C, D), (D, C), (C, C)$ with probability $\frac{1}{3}$ each.

- If agent row is recommended to choose C, then column is recommended to play C or D with equal probability. Expected payoff of row is $\frac{1}{2} \cdot 0 + \frac{1}{2}(-2) = -1$ which is greater than switching to D (expected payoff is -4.5).
- If agent row is recommended to choose D, then column is recommended to play C. Expected payoff of row is 1 which is greater than switching to C (expected payoff is 0).

Example (Correlated eq.)

	Chicken-out	Dare
Chicken-out	**0, 0**	**-2, 1**
Dare	**1, -2**	**-10, -10**

Suppose agents are recommended $(C,D),(D,C),(C,C)$ with probability $\frac{1}{3}$ each.

- If agent row is recommended to choose C, then column is recommended to play C or D with equal probability. Expected payoff of row is $\frac{1}{2}\cdot 0 + \frac{1}{2}(-2) = -1$ w ed payoff is -4.5).
- If agent row is is recommended to play C. Exp han switching to C (expected pa

Similarly for column player! (C,D), (D,C) **and** (C,C) **with** $1/3$ **each is a correlated eq.**

Definitions

Definition (Recall). *A game is specified by*

- *set of n players* $[n] = \{1, ..., n\}$
- *For each player i a set of strategies/actions S_i.*
- *set of strategy profiles* $S = S_1 \times ... \times S_n$.
- *Each agent i has a utility* $u_i : S \to [-1, 1]$ *denoting the payoff of i.*

Definitions

Definition (Recall). *A game is specified by*

- *set of n players $[n] = \{1, ..., n\}$*
- *For each player i a set of strategies/actions S_i.*
- *set of strategy profiles $S = S_1 \times ... \times S_n$.*
- *Each agent i has a utility $u_i : S \to [-1, 1]$ denoting the payoff of i.*

Definition (Correlated Equilibrium). *Correlated equilibrium is a distribution χ over S such that for all agents i and strategies b, b' of i*

$$\mathbb{E}_{s\sim\chi}[u_i(b, s_{-i})|s_i = b] \geq \mathbb{E}_{s\sim\chi}[u_i(b', s_{-i})|s_i = b].$$

Definitions

Definition (Recall). *A game is specified by*

- *set of* n *players* $[n] = \{1, ..., n\}$
- *For each player* i *a set of strategies/actions* S_i.
- *set of strategy profiles* $S = S_1 \times ... \times S_n$.
- *Each agent* i *has a utility* $u_i : S \to [-1, 1]$ *denoting the payoff of* i.

Definition (Correlated Equilibrium). *Correlated equilibrium is a distribution* χ *over* S *such that for all agents* i *and strategies* b, b' *of* i

$$\mathbb{E}_{s \sim \chi}[u_i(b, s_{-i}) | s_i = b] \geq \mathbb{E}_{s \sim \chi}[u_i(b', s_{-i}) | s_i = b].$$

Similarly for all agents i *and swapping functions* $f : S_i \to S_i$,

$$\mathbb{E}_{s \sim \chi}[u_i(s_i, s_{-i})] \geq \mathbb{E}_{s \sim \chi}[u_i(f(s_i), s_{-i})].$$

Correlated equilibrium and Nash

Remarks:

- Knowing an agent her recommended action, she can infer something about other players' moves. Yet she is better off playing the recommended action.
- Suppose χ is a product distribution. Then correlated eq. corresponds to Nash eq.

Correlated equilibrium and Nash

Remarks:

- Knowing an agent her recommended action, she can infer something about other players' moves. Yet she is better off playing the recommended action.

Set of Nash equilibria $\subseteq$ Set of correlated equilibria.

Example (Coarse Correlated eq.)

	✊	✋	✌
✊	0, 0	-1, 1	1,-1
✋	1,-1	0, 0	-1, 1
✌	-1, 1	1, -1	0, 0

Suppose the actions $(R, P), (R, S), (P, R), (P, S), (S, R), (S, P)$ are chosen with probability $\frac{1}{6}$ each.

Example (Coarse Correlated eq.)

	✊	✋	✌
✊	0, 0	-1, 1	1,-1
✋	1,-1	0, 0	-1, 1
✌	-1, 1	1, -1	0, 0

Suppose the actions $(R, P), (R, S), (P, R), (P, S), (S, R), (S, P)$ are chosen with probability $\frac{1}{6}$ each.

- If agent row plays R, agent column responds with either P or S with equal probability. If column deviates (say starts responding with paper higher possibility) she will incur more loss when row plays S (recall row plays R as well S with equal probability).

Example (Coarse Correlated eq.)

			
	0, 0	-1, 1	1,-1
	1,-1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

Suppose the actions $(R, P), (R, S), (P, R), (P, S), (S, R), (S, P)$ are chosen with probability $\frac{1}{6}$ each.

- If agent column is instructed to play P then she knows that other player is playing either R or S and column has average payoff 0. She can change then to R and improve payoff to $1/2$ compared to zero if she plays recommended action. In this case, column could exploit knowledge of action recommendation to improve her payoff.

Definitions

Definition (Coarse Correlated Equilibrium). *Coarse correlated equilibrium is a distribution χ over S such that for all agents i and strategies b of i*

$$\mathbb{E}_{s\sim\chi}[u_i(s)] \geq \mathbb{E}_{s\sim\chi}[u_i(b, s_{-i})].$$

Remark: The difference between coarse correlated and correlated is that we can choose a ``smart” swap function, namely $f$ ``knows” the distribution χ.

Definitions

Definition (**Coarse Correlated Equilibrium**). *Coarse correlated equilibrium is a distribution* χ *over* S *such that for all agents* i *and strategies* b *of* i

$$\mathbb{E}_{s\sim\chi}[u_i(s)] \geq \mathbb{E}_{s\sim\chi}[u_i(b, s_{-i})].$$

Set of correlated equilibria $\subseteq$ **Set of coarse correlated equilibria.**

Full picture of Inclusions

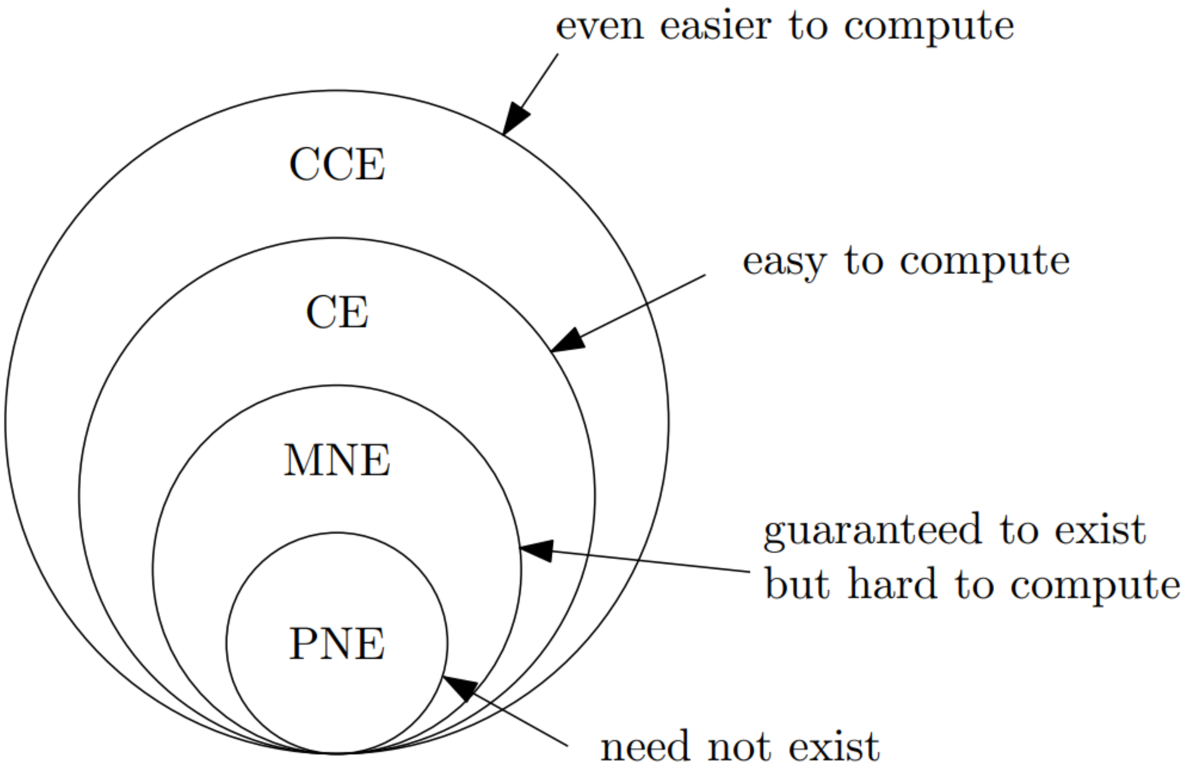

Online learning in Games

Definition. *At each time step* $t = 1 \ldots T$*.*

- *Each player i chooses* $x_i^{(t)} \in \Delta_i$ *(simplex).*
- *Each player experiences payoff* $u_i(x^{(t)})$ *and observes all players strategies* $x_j^{(t)}$.

Online learning in Games

Definition. *At each time step* $t = 1 \ldots T$.

- *Each player i chooses* $x_i^{(t)} \in \Delta_i$ *(simplex).*
- *Each player experiences payoff* $u_i(x^{(t)})$ *and observes all players strategies* $x_j^{(t)}$.

Player's i goal is to minimize the (time average) Regret, that is:

$$\frac{1}{T}\left[\max_{a \in S_i} \sum_{t=1}^{T} u_i(a, x_{-i}^{(t)}) - \sum_{t=1}^{T} u_i(x^{(t)})\right].$$

If Regret → 0 as T → ∞, the algorithm is called no-regret.

A no-regret Algorithm

Definition (Online Gradient Descent). *Let* $\ell_t : \mathcal{X} \to \mathbb{R}$ *be family of convex functions, differentiable and L-Lipschitz in some compact convex set* $\mathcal{X}$ *of diameter D. Online GD is defined:*

```
Initialize at some $x_0$.
For t:=1 to T do
    1. $y_t = x_t - \alpha_t \nabla \ell_t(x_t)$.
    2. $x_{t+1} = \Pi_{\mathcal{X}}(y_t)$.
Regret: $\frac{1}{T}\left(\sum_{t=1}^{T} \ell_t(x_t) - \min_x \sum_{t=1}^{T} \ell_t(x)\right)$.
```

A no-regret Algorithm

Definition (Online Gradient Descent). *Let* $\ell_t : \mathcal{X} \to \mathbb{R}$ *be family of convex functions, differentiable and L-Lipschitz in some compact convex set* $\mathcal{X}$ *of diameter D. Online GD is defined:*

```
Initialize at some $x_0$.
For t:=1 to T do
    1. $y_t = x_t - \alpha_t \nabla \ell_t(x_t)$.
    2. $x_{t+1} = \Pi_{\mathcal{X}}(y_t)$.
Regret: $\frac{1}{T}\left(\sum_{t=1}^{T} \ell_t(x_t) - \min_x \sum_{t=1}^{T} \ell_t(x)\right)$.
```

step-size

$$\ell_t = -u_i(x^{(t)})$$

Analysis of Online GD for L-Lipschitz

Theorem (Online Gradient Descent). *Let $\ell_t : \mathcal{X} \to \mathbb{R}$ be family of convex functions, differentiable and L-Lipschitz in some compact convex set $\mathcal{X}$ of diameter D. It holds*

$$\left(\frac{1}{T}\sum_{t=1}^{T}\ell_t(x_t) - \min_{x}\sum_{t=1}^{T}\ell_t(x)\right) \leq \frac{3}{2}\frac{LD}{\sqrt{T}},$$

with appropriately choosing $\alpha = \frac{D}{L\sqrt{t}}$.

Remarks:

- If we want error ϵ, we need $T = \Theta\left(\frac{L^2D^2}{\epsilon^2}\right)$ iterations.
- I could have written Multiplicative Weights Update. This is another no-regret

algorithm! Same regret guarantees, i.e., $O\left(\frac{1}{\sqrt{T}}\right)$.

Analysis of Online GD for L-Lipschitz

Proof. Let x^* be the argmin of $\sum \ell_t(x)$.

$$\begin{aligned}\ell_t(x_t) - \ell_t(x^*) &\leq \nabla \ell_t(x_t)^\top (x_t - x^*) \text{ convexity,} \\ &= \frac{1}{\alpha_t}(x_t - y_t)^\top (x_t - x^*) \text{ definition of GD,}\end{aligned}$$

Analysis of Online GD for L-Lipschitz

Proof. Let x^* be the argmin of $\sum \ell_t(x)$.

$$\begin{aligned}
\ell_t(x_t) - \ell_t(x^*) &\leq \nabla \ell_t(x_t)^\top (x_t - x^*) \text{ convexity,} \\
&= \frac{1}{\alpha_t}(x_t - y_t)^\top (x_t - x^*) \text{ definition of GD,} \\
&= \frac{1}{2\alpha_t}\left(\|x_t - x^*\|_2^2 + \|x_t - y_t\|_2^2 - \|y_t - x^*\|_2^2\right) \text{ law of Cosines,} \\
&= \frac{1}{2\alpha_t}\left(\|x_t - x^*\|_2^2 - \|y_t - x^*\|_2^2\right) + \frac{\alpha_t}{2}\|\nabla \ell_t(x_t)\|_2^2 \text{ Def. of } y_t,
\end{aligned}$$

Analysis of Online GD for L-Lipschitz

Proof. Let x^* be the argmin of $\sum \ell_t(x)$.

$$\begin{aligned}
\ell_t(x_t) - \ell_t(x^*) &\leq \nabla \ell_t(x_t)^\top (x_t - x^*) \text{ convexity,} \\
&= \frac{1}{\alpha_t}(x_t - y_t)^\top (x_t - x^*) \text{ definition of GD,} \\
&= \frac{1}{2\alpha_t}\left(\|x_t - x^*\|_2^2 + \|x_t - y_t\|_2^2 - \|y_t - x^*\|_2^2\right) \text{ law of Cosines,} \\
&= \frac{1}{2\alpha_t}\left(\|x_t - x^*\|_2^2 - \|y_t - x^*\|_2^2\right) + \frac{\alpha_t}{2}\|\nabla \ell_t(x_t)\|_2^2 \text{ Def. of } y_t, \\
&\leq \frac{1}{2\alpha_t}\left(\|x_t - x^*\|_2^2 - \|y_t - x^*\|_2^2\right) + \frac{\alpha_t L^2}{2} \text{ Lipschitz,} \\
&\leq \frac{1}{2\alpha_t}\left(\|x_t - x^*\|_2^2 - \|x_{t+1} - x^*\|_2^2\right) + \frac{\alpha_t L^2}{2} \text{ projection.}
\end{aligned}$$

Analysis of Online GD for L-Lipschitz

Proof cont. Since

$$\ell_t(x_t) - \ell_t(x^*) \leq \frac{1}{2\alpha_t}\left(\|x_t - x^*\|_2^2 - \|x_{t+1} - x^*\|_2^2\right) + \frac{\alpha_t L^2}{2},$$

taking the telescopic sum we have

$$\begin{aligned}\sum_{t=1}^{T}\left(\ell_t(x_t) - \ell_t(x^*)\right) &\leq \sum_{t=1}^{T}\|x_t - x^*\|_2^2\left(\frac{1}{2\alpha_t} - \frac{1}{2\alpha_{t-1}}\right) + \frac{L^2}{2}\sum_{t=1}^{T}\alpha_t. \\ &\leq \frac{D^2}{2}\sum_{t=1}^{T}\left(\frac{1}{\alpha_t} - \frac{1}{\alpha_{t-1}}\right) + \frac{L^2}{2}\sum_{t=1}^{T}\alpha_t.\end{aligned}$$

Analysis of Online GD for L-Lipschitz

Proof cont. Since

$$\ell_t(x_t) - \ell_t(x^*) \leq \frac{1}{2\alpha_t}\left(\|x_t - x^*\|_2^2 - \|x_{t+1} - x^*\|_2^2\right) + \frac{\alpha_t L^2}{2},$$

taking the telescopic sum we have

$$\begin{aligned}\sum_{t=1}^{T}(\ell_t(x_t) - \ell_t(x^*)) &\leq \sum_{t=1}^{T}\|x_t - x^*\|_2^2\left(\frac{1}{2\alpha_t} - \frac{1}{2\alpha_{t-1}}\right) + \frac{L^2}{2}\sum_{t=1}^{T}\alpha_t.\\ &\leq \frac{D^2}{2}\sum_{t=1}^{T}\left(\frac{1}{\alpha_t} - \frac{1}{\alpha_{t-1}}\right) + \frac{L^2}{2}\sum_{t=1}^{T}\alpha_t.\\ &\leq \frac{D^2}{2\alpha_T} + \frac{L^2}{2}\sum_{t=1}^{T}\alpha_t \leq \frac{LD}{2}\sqrt{T} + 2\sqrt{T}\frac{LD}{2}.\end{aligned}$$

where we used the fact $\sum \frac{1}{\sqrt{t}} \leq 2\sqrt{T}$ and $\alpha_t = \frac{D}{\sqrt{t}L}$.

Computing coarse correlated equilibria

Suppose that each agent i uses no-regret dynamics (online GD), with $l_t = -u_i(x^{(t)})$ where $x^{(t)}$ is the mixed strategy profile at iterate t.

Computing coarse correlated equilibria

Suppose that each agent i uses no-regret dynamics (online GD), with $l_t = -u_i(x^{(t)})$ where $x^{(t)}$ is the mixed strategy profile at iterate t.

- Let σ^t be the product distribution on S induced by $x^{(t)}$.
- Let σ be the uniform distribution over $\{\sigma^1, \dots, \sigma^T\}$.

Computing coarse correlated equilibria

Suppose that each agent i uses no-regret dynamics (online GD), with $l_t = -u_i(x^{(t)})$ where $x^{(t)}$ is the mixed strategy profile at iterate t.

- Let σ^t be the product distribution on S induced by $x^{(t)}$.
- Let σ be the uniform distribution over $\{\sigma^1, \dots, \sigma^T\}$.

We conclude that for each agent i

$$\mathbb{E}_{s\sim\sigma}[u_i(s)] = \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{s\sim\sigma^t}[u_i(s)]$$

$$\min_{b\in S_i}\mathbb{E}_{s\sim\sigma}[u_i(b, s_{-i})] = \min_{b\in S_i}\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{s\sim\sigma^t}[u_i(b, s_{-i})]$$

Computing coarse correlated equilibria

Suppos

where

- Let
- Let

We co

Lipschitz

$$\min_{b \in S_i} \mathbb{E}_{s\sim\sigma}[u_i(b, s_{-i})] - \mathbb{E}_{s\sim\sigma}[u_i(s)] \le \frac{3}{2}\frac{\sqrt{n}}{\sqrt{T}}$$

$$\mathbb{E}_{s\sim\sigma}[u_i(s)] = \frac{1}{T}\sum_{t=1} \mathbb{E}_{s\sim\sigma^t}[u_i(s)]$$

$$\min_{b \in S_i} \mathbb{E}_{s\sim\sigma}[u_i(b, s_{-i})] = \min_{b \in S_i} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{s\sim\sigma^t}[u_i(b, s_{-i})]$$

Computing coarse correlated equilibria

Suppos

where

- Let
- Let

We co

$$\min_{b \in S_i} \mathbb{E}_{s\sim\sigma}[u_i(b,s_{-i})] - \mathbb{E}_{s\sim\sigma}[u_i(s)] \leq \frac{3}{2}\frac{\overset{\text{Lipschitz}}{\boxed{\sqrt{n}}}}{\sqrt{T}}$$

Choosing $T = \frac{9n}{4\epsilon^2}$ we conclude σ is ϵ-approximate CCE!

$$\mathbb{E}_{s\sim\sigma}[u_i(s)] = \frac{1}{T}\sum_{t=1} \mathbb{E}_{s\sim\sigma^t}[u_i(s)]$$

$$\min_{b \in S_i} \mathbb{E}_{s\sim\sigma}[u_i(b,s_{-i})] = \min_{b \in S_i} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{s\sim\sigma^t}[u_i(b,s_{-i})]$$

Computing coarse correlated equilibria

Suppos

where

- Let
- Let

We co

Lipschitz

$$\min_{b \in S_i} \mathbb{E}_{s\sim\sigma}[u_i(b, s_{-i})] - \mathbb{E}_{s\sim\sigma}[u_i(s)] \leq \frac{3}{2}\frac{\sqrt{n}}{\sqrt{T}}$$

Choosing $T = \frac{9n}{4\epsilon^2}$ we conclude σ is ϵ-approximate CCE!

If we use MWUA, it gives $O\left(\frac{\ln n}{\epsilon^2}\right)$.

$$\mathbb{E}_{s\sim\sigma}[u_i(s)] = \frac{1}{T}\sum_{t=1} \mathbb{E}_{s\sim\sigma^t}[u_i(s)]$$

$$\min_{b \in S_i} \mathbb{E}_{s\sim\sigma}[u_i(b, s_{-i})] = \min_{b \in S_i} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{s\sim\sigma^t}[u_i(b, s_{-i})]$$