Learning in finite games and convex potential games

Walid Krichene

May 12, 2014

Outline

1. Learning in finite games
   - Equilibria
   - Regret
   - Blackwell approachability

2. Learning in convex potential games
   - The routing game and Nash equilibria
   - Regret
   - Mirror descent on potential function
   - Convergence of a dense subsequence
   - Strong convergence

Introduction

- Finite player game, played repeatedly.
- Player k has mixed strategy $p^k(t)$, which obeys an update rule.

Introduction

- Finite player game, played repeatedly.
- Player k has mixed strategy $p^k(t)$, which obeys an update rule.
- Questions:
  - Does $(p(t))_{t \in \mathbb{N}}$ converge to an equilibrium set?

Introduction

- Finite player game, played repeatedly.
- Player k has mixed strategy $p^k(t)$, which obeys an update rule.
- Questions:
  - Does $(p(t))_{t\in\mathbb{N}}$ converge to an equilibrium set?
  - Does $(\bar{p}(t))_{t\in\mathbb{N}}$ converge to an equilibrium set?

$$\bar{p}(t) = \frac{1}{t}\sum_{\tau=1}^{t} p(\tau)$$

Introduction

- Finite player game, played repeatedly.
- Player k has mixed strategy $p^k(t)$, which obeys an update rule.
- Questions:
  - Does $(p(t))_{t\in\mathbb{N}}$ converge to an equilibrium set?
  - Does $(\bar{p}(t))_{t\in\mathbb{N}}$ converge to an equilibrium set?

$$\bar{p}(t) = \frac{1}{t}\sum_{\tau=1}^{t} p(\tau)$$

Which equilibrium sets? Nash / Correlated (Auman) / Hannan

$$\mathcal{N} \subseteq \mathcal{C} \subseteq \mathcal{H}$$

Results

Game class	Dynamics	Convergence
finite	external-regret minimizing	$\bar{p}(t) \to \mathcal{H}$
finite	internal-regret minimizing	$\bar{p}(t) \to \mathcal{C}$
2 player zero-sum	external-regret minimizing	$\bar{p}(t) \to \mathcal{N}$
2 players, 2 actions	fictitious play	$\bar{p}(t) \to \mathcal{N}$

The setting

- K players
- player k: finite action set S_k

The setting

- K players
- player k: finite action set S_k
- joint action $i \in S_1 \times \cdots \times S_K$
- loss $\ell_k(i) \in [0, 1]$

The setting

- K players
- player k: finite action set S_k
- joint action $i \in S_1 \times \cdots \times S_K$
- loss $\ell_k(i) \in [0, 1]$
- players randomize: mixed strategy profile $p \in \Delta^{S_1 \times \cdots \times S_K}$
- If players randomize independently,

$$p(i) = p^{(1)}_{i_1} \dots p^{(K)}_{i_K}$$

- joint random action $I \in S_1 \times \cdots \times S_K$, $I \sim p$
- expected loss $\mathbb{E}_p[\ell_k] = \sum_i p(i)\ell_k(i)$
- write $\ell_k(p) = \mathbb{E}_p[\ell_k]$
  (extension of ℓ from $S_1 \times \cdots \times S_K$ to $\Delta^{S_1 \times \cdots \times S_K}$)

Outline

1. Learning in finite games
   - Equilibria
   - Regret
   - Blackwell approachability
2. Learning in convex potential games
   - The routing game and Nash equilibria
   - Regret
   - Mirror descent on potential function
   - Convergence of a dense subsequence
   - Strong convergence

Nash equilibrium

Nash equilibrium

p is a Nash eq. if it is a product distribution and for all k and all p'

$$\ell_k(p^{(k)}, p^{(-k)}) \leq \ell_k(p', p^{(-k)})$$

- No one has an incentive to unilaterally deviate.

Correlated equilibria

Player actions may be correlated.

Correlated equilibria (Aumann)

- Probability distribution p over $S_1 \times \cdots \times S_K$, not necessarily a product distribution

Correlated equilibria

Player actions may be correlated.

Correlated equilibria (Aumann)

- Probability distribution p over $S_1 \times \cdots \times S_K$, not necessarily a product distribution
- Some coordinator gives advice $I \sim P$
- If player unilaterally deviates from j, worse-off in expectation.

Correlated equilibria

Player actions may be correlated.

Correlated equilibria (Aumann)

- Probability distribution p over $S_1 \times \cdots \times S_K$, not necessarily a product distribution
- Some coordinator gives advice $I \sim P$
- If player unilaterally deviates from j, worse-off in expectation.

"Switching j to j' is a bad idea". For all j, j'

$$\sum_{i:i^k=j} p(i)[\ell_k(j, i^{(-k)}) - \ell^{(k)}(j', i^{(-k)})] \leq 0$$

Correlated equilibria

Player actions may be correlated.

Correlated equilibria (Aumann)

- Probability distribution p over $S_1 \times \cdots \times S_K$, not necessarily a product distribution
- Some coordinator gives advice $I \sim P$
- If player unilaterally deviates from j, worse-off in expectation.

"Switching j to j' is a bad idea". For all j, j'

$$\sum_{i : i^k = j} p(i)[\ell_k(j, i^{(-k)}) - \ell^{(k)}(j', i^{(-k)})] \leq 0$$

Structure of correlated equilibria

- p vector in $\Delta^{S_1 \times \cdots \times S_K}$
- Linear inequalities in $p \Rightarrow \mathcal{C}$ is a closed convex polyhedron
- Convex combination of Nash eq $\Rightarrow$ correlated eq. (but not $\Leftarrow$)

Hannan set

Hannan set

- $p \in \mathcal{H}$ if for all j

$$\ell_k(p) \leq \ell_k(\delta^j \times p^{(-k)})$$

Hannan set

Hannan set

- $p \in \mathcal{H}$ if for all j

$$\ell_k(p) \leq \ell_k(\delta^j \times p^{(-k)})$$

- Same as the Nash condition but p not necessarily product distribution
- $p \in \mathcal{H}$ and p is a product $\Leftrightarrow p \in \mathcal{N}$
- $\mathcal{H}$ is a convex polyhedron
- $\mathcal{C} \subseteq \mathcal{H}$

Outline

1. Learning in finite games
   - Equilibria
   - Regret
   - Blackwell approachability
2. Learning in convex potential games
   - The routing game and Nash equilibria
   - Regret
   - Mirror descent on potential function
   - Convergence of a dense subsequence
   - Strong convergence

Repeated play

- k maintains $p^{(k)}(t)$, draws action $I^{(k)}(t) \sim p_t^{(k)}$
- observes all players' actions $I(t) = (I^{(1)}(t), \dots, I^{(K)}(t))$

Repeated play

- k maintains $p^{(k)}(t)$, draws action $I^{(k)}(t) \sim p_t^{(k)}$
- observes all players' actions $I(t) = (I^{(1)}(t), \dots, I^{(K)}(t))$
- Uncoupled play: player only knows his loss function.

Regret

Regret

Compare expected loss to loss had we played differently.
Let $\psi : S_k \to S_k$.
This defines distribution $\phi(p)$

$$\phi(p)(i) = p(\psi(i^k), i^{-k})$$

Regret:

$$R_\phi(t) = \sum_{\tau=1}^{t} \ell_k(p(\tau)) - \sum_{\tau=1}^{t} \ell_k(\phi(p(\tau)))$$

Regret

Regret

Compare expected loss to loss had we played differently.
Let $\psi : S_k \to S_k$.
This defines distribution $\phi(p)$

$$\phi(p)(i) = p(\psi(i^k), i^{-k})$$

Regret:

$$R_\phi(t) = \sum_{\tau=1}^{t} \ell_k(p(\tau)) - \sum_{\tau=1}^{t} \ell_k(\phi(p(\tau)))$$

- Idea: fix a class of functions Φ. This defines a regret vector $R(t) = (R_\phi(t))_{\phi \in \Phi}$.
- We want $\lim_{t\to\infty} d(\frac{R(t)}{t}, \mathbb{R}_-^{|\Phi|}) = 0$

Regret

By linearity

$$\frac{1}{t} R_\phi(t) = \frac{1}{t} \sum_{\tau=1}^{t} \sum_i p(\tau)(i) \left(\ell_k(i) - \ell_k(\psi(i^k), i^{(-k)}) \right)$$

$$= \sum_i \bar{p}(t)(i) \left(\ell_k(i) - \ell_k(\psi(i^k), i^{(-k)}) \right)$$

$$= \ell_k(\bar{p}(t)) - \ell_k(\phi(\bar{p}(t)))$$

Regret

External regret

For all $j \in S_k$, let ψ_j map any action to j.

$$\phi_j : p \mapsto \delta^j \times p^{(-k)}$$

$$R_{ext}(t) \in \mathbb{R}^{|S_k|}$$

Regret

External regret

For all $j \in S_k$, let ψ_j map any action to j.

$$\phi_j : p \mapsto \delta^j \times p^{(-k)}$$

$$R_{ext}(t) \in \mathbb{R}^{|S_k|}$$

$$\frac{1}{t} R_{\phi_j}(t) = \ell_k(\bar{p}(t)) - \ell_k(\delta^j \times \bar{p}^{(-k)}(t))$$

Regret

External regret

For all $j \in S_k$, let ψ_j map any action to j.

$$\phi_j : p \mapsto \delta^j \times p^{(-k)}$$

$$R_{ext}(t) \in \mathbb{R}^{|S_k|}$$

$$\frac{1}{t} R_{\phi_j}(t) = \ell_k(\bar{p}(t)) - \ell_k(\delta^j \times \bar{p}^{(-k)}(t))$$

Convergence to Hannan set

$$\frac{1}{t} R^{ext}(t) \to \mathbb{R}_-^{|S_k|} \Leftrightarrow \bar{p}(t) \to \mathcal{H}$$

Regret

Internal regret

For all $j, j' \in S_k \times S_k$, let $\psi_{j \to j'}$ the function play j' instead of j.

$$R_{int}(t) \in \mathbb{R}^{|S_k|(|S_k|-1)}$$

Regret

Internal regret

For all $j, j' \in S_k \times S_k$, let $\psi_{j \to j'}$ the function play j' instead of j.

$$R_{int}(t) \in \mathbb{R}^{|S_k|(|S_k|-1)}$$

$$\frac{1}{t} R_{\phi_{j \to j'}}(t) = \sum_i \bar{p}(t)(i) \left(\ell_k(i) - \ell_k(\psi_{j \to j'}(i^k), i^{(-k)} \right)$$

$$= \sum_{i : i^k \neq j} \bar{p}(t)(i).0 + \sum_{i : i^k = j} \bar{p}(t)(i) \left(\ell_k(j, i^{(-k)}) - \ell_k(j', i^{(-k)}) \right)$$

Regret

Internal regret

For all $j, j' \in S_k \times S_k$, let $\psi_{j \to j'}$ the function play j' instead of j.

$$R_{int}(t) \in \mathbb{R}^{|S_k|(|S_k|-1)}$$

$$\frac{1}{t} R_{\phi_{j \to j'}}(t) = \sum_i \bar{p}(t)(i) \left(\ell_k(i) - \ell_k(\psi_{j \to j'}(i^k), i^{(-k)} \right)$$
$$= \sum_{i : i^k \neq j} \bar{p}(t)(i).0 + \sum_{i : i^k = j} \bar{p}(t)(i) \left(\ell_k(j, i^{(-k)}) - \ell_k(j', i^{(-k)}) \right)$$

Convergence to Correlated equilibria

$$\frac{1}{t} R_{int}(t) \to \mathbb{R}_-^{|S_k|(|S_k|-1)} \Leftrightarrow \bar{p}(t) \to \mathcal{C}$$

Summary

Nash equilibria $\mathcal{N}$

p product distribution. For all j

$$\ell_k(p^{(k)}, p^{(-k)}) \leq \ell_k(\delta^j, p^{(-k)})$$

Correlated equilibria $\mathcal{C}$

p not necessarily a product distribution. For all j, j'

$$\sum_{i:i^k=j} p(i)\ell_k(j, i^{(-k)}) \leq \sum_{i:i^k=j} p(i)\ell_k(j', i^{(-k)})]$$

Hannan set $\mathcal{H}$

p not necessarily a product distribution. For all j

$$\ell_k(p^{(k)}, p^{(-k)}) \leq \ell_k(\delta^j, p^{(-k)})$$

Fictitious play

On iteration t, play best response to cumulative losses

$$L^{(k)}(t)(i^{(k)}) = \sum_{\tau=1}^{t} \ell(\tau)(i^{(k)}, I^{(-k)})$$

Fictitious play

On iteration t, play best response to cumulative losses

$$L^{(k)}(t)(i^{(k)}) = \sum_{\tau=1}^{t} \ell(\tau)(i^{(k)}, I^{(-k)})$$

Fictitious play

$$i^{(k)}(t+1) \in \arg\min_{i \in S^k} L^{(k)}(t)(i)$$

equivalent to

$$\min_{p^{(k)} \in \Delta^{S_k}} \left\langle p^{(k)}, L^{(k)}(t) \right\rangle$$

Fictitious play

On iteration t, play best response to cumulative losses

$$L^{(k)}(t)(i^{(k)}) = \sum_{\tau=1}^{t} \ell(\tau)(i^{(k)}, I^{(-k)})$$

Fictitious play

$$i^{(k)}(t+1) \in \arg\min_{i \in S^{k}} L^{(k)}(t)(i)$$

equivalent to

$$\min_{p^{(k)} \in \Delta^{S_k}} \left\langle p^{(k)}, L^{(k)}(t) \right\rangle$$

- Only converges in very special cases (e.g. 2 players, 2 actions each)
- Not Hannan-consistent in general ($\frac{1}{t} R_{ext}(t) \not\to 0$)

Regret-minimization Vs. fictitious play

Fictitious play

$$p^{(k)}(t+1) \in \arg\min_{p^{(k)} \in \Delta^{\mathcal{S}_k}} \left\langle p^{(k)}, L^{(k)}(t) \right\rangle$$

Regret-minimization

Some algorithms can be written as

$$p^{(k)}(t+1) \in \arg\min_{p^{(k)} \in \Delta^{\mathcal{S}_k}} \left\langle p^{(k)}, L^{(k)}(t) \right\rangle + \frac{1}{\gamma} R(p^{(k)})$$

Algorithms with sublinear regret

How can we find algorithms such that

$$\frac{R(t)}{t} \to 0$$

- Blackwell approachability (reduces the problem)

Outline

Blackwell approachability

- Vector $r(i,j)$
- Extended bilinearly $r(p,q) = \sum_i \sum_j p(i)q(j)r(i,j)$
- Assume $\|r(i,j)\|_2 \leq 1$
- A closed convex set $S \subseteq B(0,1)$ is approachable if $\exists$ a randomized strategy such that for all sequences $J(t)$,

$$\lim_{t\to\infty} d\left(\frac{1}{t}\sum_{\tau\leq t} r(I(\tau), J(\tau)), S\right) = 0$$

Blackwell approachability: Half spaces

- $H_{a,c} = \{r : \langle a, r \rangle \leq c\}$

Blackwell approachability: Half spaces

- $H_{a,c} = \{r : \langle a, r \rangle \leq c\}$
- Define new game: loss of (i, j) is $\langle a, r(i, j) \rangle := (M_a)_{i,j}$. Defines a matrix M_a
- This a scalar game

Blackwell approachability: Half spaces

- $H_{a,c} = \{r : \langle a, r \rangle \leq c\}$
- Define new game: loss of (i, j) is $\langle a, r(i,j) \rangle := (M_a)_{i,j}$. Defines a matrix M_a
- This a scalar game

Half spaces

$H_{a,c}$ is approachable iff $c \geq V = \min_{p \in \Delta} \max_{q \in \Delta} \langle p, M_a q \rangle$

- in particular, if $H_{a,c}$ is approachable, then $\exists p \in \Delta : \forall q, \; \langle p, M_a q \rangle \leq c$
- Call it $p^*(H_{a,c})$.
- $\forall j, \; r(p^* H_{a,c}, j) \in H_{a,c}$

Blackwell approachability theorem

Theorem: Blackwell approachability

S is approachable iff every half space containing S is approachable.

Blackwell approachability theorem

Theorem: Blackwell approachability

S is approachable iff every half space containing S is approachable.

- $\Rightarrow$ if S is approachable, any super set is approachable.

Blackwell approachability theorem

Theorem: Blackwell approachability

S is approachable iff every half space containing S is approachable.

- $\Rightarrow$ if S is approachable, any super set is approachable.
- $\Leftarrow$ Assume every half space containing S is approachable.
  - average regret $A_t = \frac{1}{t} \sum_{\tau \leq t} r(i_\tau, j_\tau)$
  - at t, if $A_{t-1} \notin S$
    project $\pi_S(A_{t-1})$
    define H_{t-1}: through $\pi_S(A_{t-1})$, $\perp (A_{t-1} - \pi_S(A_{t-1}))$.
  - By assumption H_{t-1} is approachable. Play $p^*(H_{t-1})$
  - Next step: $r(p^*(H_{t-1}), j_t) \in H_{t-1}$

Blackwell approachability theorem

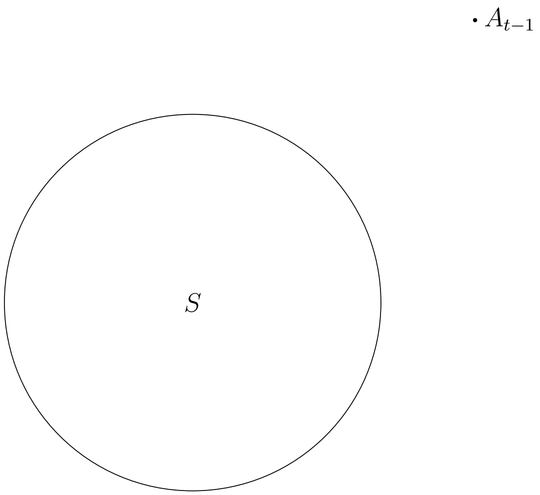

Figure : Instantaneous regret $r(p_t, J_t)$ is forced in H_{t-1}

Blackwell approachability theorem

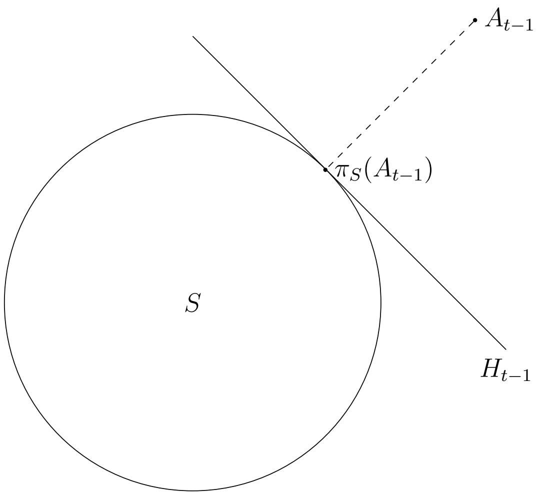

Figure : Instantaneous regret $r(p_t, J_t)$ is forced in H_{t-1}

Blackwell approachability theorem

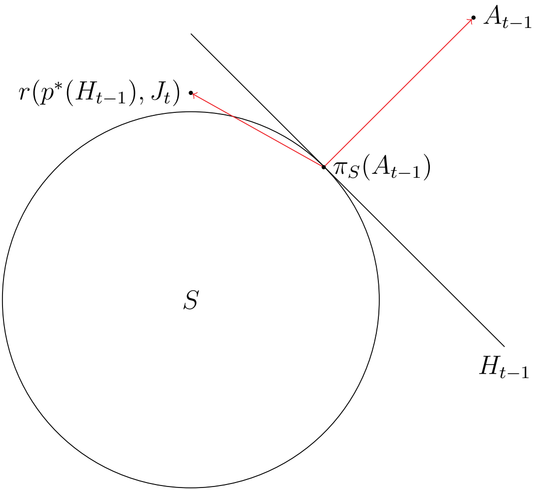

Figure : Instantaneous regret $r(p_t, J_t)$ is forced in H_{t-1}

Blackwell approachability theorem

$$
\begin{aligned}
d(A_t, S)^2 &\leq \|A_t - \pi_S(A_{t-1})\|_2^2 \\
&= \|\frac{t-1}{t}A_{t-1} + \frac{1}{t}r(p_t, j_t) - (\frac{t-1}{t} + \frac{1}{t})\pi_S(A_{t-1})\|_2^2 \\
&= \left(\frac{t-1}{t}\right)^2 d(A_{t-1}, S)^2 + \frac{1}{t^2}\|r(p_t, j_t) - \pi_S(A_{t-1})\|_2^2 \\
&\quad + 2\frac{t-1}{t^2}\langle A_{t-1} - \pi_S(A_{t-1}), r(p_t, j_t) - \pi_S(A_{t-1})\rangle \\
&\leq \left(\frac{t-1}{t}\right)^2 d(A_{t-1}, S)^2 + \frac{1}{t^2}2^2
\end{aligned}
$$

Blackwell approachability theorem

$$
\begin{aligned}
d(A_t, S)^2 &\leq \|A_t - \pi_S(A_{t-1})\|_2^2 \\
&= \|\frac{t-1}{t}A_{t-1} + \frac{1}{t}r(p_t, j_t) - (\frac{t-1}{t} + \frac{1}{t})\pi_S(A_{t-1})\|_2^2 \\
&= \left(\frac{t-1}{t}\right)^2 d(A_{t-1}, S)^2 + \frac{1}{t^2}\|r(p_t, j_t) - \pi_S(A_{t-1})\|_2^2 \\
&\quad + 2\frac{t-1}{t^2}\langle A_{t-1} - \pi_S(A_{t-1}), r(p_t, j_t) - \pi_S(A_{t-1})\rangle \\
&\leq \left(\frac{t-1}{t}\right)^2 d(A_{t-1}, S)^2 + \frac{1}{t^2}2^2
\end{aligned}
$$

multiply by t^2, then by induction

$$t^2 d(A_t, S)^2 \leq 4t$$

thus

$$d(A_t, S) \leq \frac{2}{\sqrt{t}}$$

Blackwell approachability theorem

- In fact, can show almost sure convergence to S.

Potential based approachability

- S approachable. Want to define a class of algorithms which approach S.
- Potential $\phi \geq 0$ convex, $\phi(x) = 0 \Leftrightarrow x \in S$.
- e.g. $\phi(s) = d(x, S)^2$

Potential based approachability

- S approachable. Want to define a class of algorithms which approach S.
- Potential $\phi \geq 0$ convex, $\phi(x) = 0 \Leftrightarrow x \in S$.
- e.g. $\phi(s) = d(x, S)^2$

Algorithm

Same idea: if $A_{t-1} \notin S$, use the potential to choose a strategy p_t such that

$$\langle A_{t-1} - \pi_S(A_{t-1}), r(p_t, j) - \pi_S(A_{t-1}) \rangle \leq 0 \; \forall j$$

where π_S is the Bregman projection on S

$$\begin{aligned} \pi_S(a) &= \arg\min_{r \in S} \phi(r) - \phi(a) - \langle \nabla \phi(a), r - a \rangle \\ &= \arg\min_{r \in S} - \langle \nabla \phi(a), r \rangle \end{aligned}$$

Back to regret

- Choose $S = \mathbb{R}^d_-$.
- To apply approachability need to guarantee

$$\langle d, r(p_t, j)\rangle \leq 0 \ \forall j$$

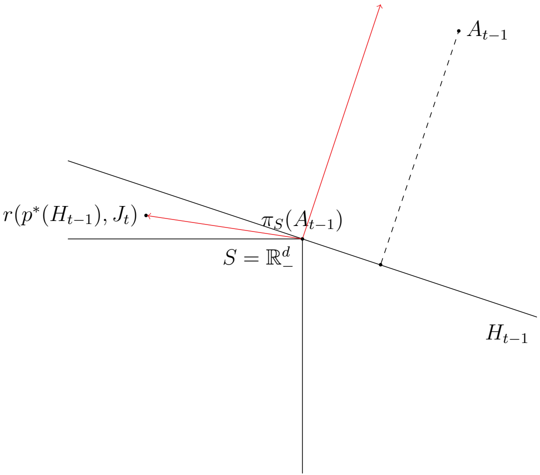

Back to regret

$$\langle d, r(p_t, j)\rangle \leq 0 \ \forall j$$

- Recall

$$r_i(p, j) = \langle p, \ell(j)\rangle - \ell_i(j)$$

- take $p = \frac{d}{\|d\|_1}$

$$\langle d, r(p, j)\rangle = \sum_i \left(\langle d/\|d\|_1, \ell(j)\rangle\, d_i - \ell_i(j) d_i\right) = 0$$

Convergence of $p(t)$?

We have $\bar{p}(t) \to S$. Can we show $p(t) \to S$?

Convergence of $p(t)$?

We have $\bar{p}(t) \to S$. Can we show $p(t) \to S$?
Show that $p(t)$ converges.

- If $p(t) \to p_\infty$, then $\bar{p}(t) \to p_\infty$, so $\{p_\infty\} \in S$.
- One possible technique:
  - consider a continuous-time version of the dynamics

$$\dot{p}(t) = F(p(t))$$

  - show that $\mathbf{p}$ converges
  - show that discrete trajectories approach continuous trajectories

Extensions

- Continuous time
  Internal regret, external regret: Hart and Mas-Colell (2003)

Extensions

- Continuous time
  Internal regret, external regret: Hart and Mas-Colell (2003)
- Continuum of players
  Potential games in continuous time, dynamics with a positive correlation condition. Sandholm (2001)
  Congestion games with an approximate replicator update

Extensions

- Continuous time
  Internal regret, external regret: Hart and Mas-Colell (2003)
- Continuum of players
  Potential games in continuous time, dynamics with a positive correlation condition. Sandholm (2001)
  Congestion games with an approximate replicator update
- Continuous action spaces
  Convex compact sets, minimizing exp-concave loss functions. Hazan et. al (2006)
  Convex compact sets, learning correlated equilibria. Stoltz and Lugosi (2006)
  Compact locally convex, Lipschitz loss functions.

Outline

Routing game

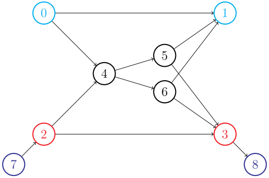

Figure : Example network

- Graph (V, E)
- source-sink pairs, (s_k, t_k): paths $\mathcal{P}_k$

Routing game

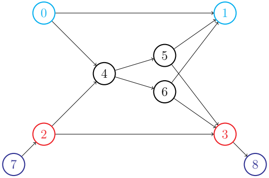

Figure : Example network

- Graph (V, E)
- source-sink pairs, (s_k, t_k): paths $\mathcal{P}_k$
- Players choose a distribution over paths $\mu^k \in \Delta^{\mathcal{P}_k}$,
- μ determines edge flows $\phi = M\mu$ (linear function)

Routing game

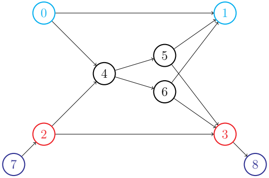

Figure : Example network

- Graph (V, E)
- source-sink pairs, (s_k, t_k): paths $\mathcal{P}_k$
- Players choose a distribution over paths $\mu^k \in \Delta^{\mathcal{P}_k}$,
- μ determines edge flows $\phi = M\mu$ (linear function)
- Congestion on edge e: $c_e : \phi_e \mapsto c_e(\phi_e)$, increasing
- want to minimize personal latency $\ell_p^k(\mu) = \sum_{e \in p} c_e(\phi_e)$

Player = infinitesimal amount of flow.
μ = combined decision of all players.

More precisely

- Measurable set of players $(\mathcal{X}_k, \mathcal{S}_k, m_k)$, atomless
- $m_k(\mathcal{X}_k)$ finite

More precisely

- Measurable set of players $(\mathcal{X}_k, \mathcal{S}_k, m_k)$, atomless
- $m_k(\mathcal{X}_k)$ finite
- Every player $x \in \mathcal{X}_k$ chooses distribution $\pi^k(x)$
- This induces population distribution $\mu^k(x) = \int_{\mathcal{X}_k} \pi^k(x) dm(x)$

Nash equilibria

Nash equilibrium

μ is a Nash equilibrium if for all k, for all $p \in \mathcal{P}_k$ with positive mass, $\ell_p^k(\mu)$ is minimal on $\mathcal{P}_k$

$$\ell_p^k(\mu) \leq \ell_{p'}^k(\mu) \ \forall p' \in \mathcal{P}_k$$

- How to compute Nash equilibria?

Nash equilibria

Nash equilibrium

μ is a Nash equilibrium if for all k, for all $p \in \mathcal{P}_k$ with positive mass, $\ell_p^k(\mu)$ is minimal on $\mathcal{P}_k$

$$\ell_p^k(\mu) \leq \ell_{p'}^k(\mu) \; \forall p' \in \mathcal{P}_k$$

- How to compute Nash equilibria? Convex formulation

Potential function

μ is a Nash equilibrium iff it minimizes a potential function

$$\begin{aligned} &\text{minimize}_{\phi, \mu \in \Delta^{\mathcal{P}_1} \times \cdots \times \Delta^{\mathcal{P}_K}} && \sum_e \int_0^{\phi_e} c_e(u) du \\ &\text{subject to} && \forall e, \sum_k \sum_{p \ni e} \mu_p^k = \phi_e \end{aligned}$$

Nash equilibria

Convex potential function

$$V(\mu) = \sum_e \int_0^{(M\mu)_e} c_e(u)du$$

- V is convex: composition of $\phi \mapsto \sum_e \int_0^{\phi_e} c_e(u)du$ (strongly convex) with $\mu \mapsto M\mu$
- $\nabla_{\mu^k} V(\mu) = \ell^k(\mu)$
- M non injective in general, V weakly convex, minimizer not unique.

Motivation for a learning model

- How do players find a Nash equilibrium?

Motivation for a learning model

- How do players find a Nash equilibrium?
  Ideally: distributed, and has minimal information requirements.
  - loss on the player's path
  - loss on population paths
  - all edge losses
  - all congestion functions and edge flows

Motivation for a learning model

- How do players find a Nash equilibrium?
  Ideally: distributed, and has minimal information requirements.
  - loss on the player's path
  - loss on population paths
  - all edge losses
  - all congestion functions and edge flows
- Player dynamics: given $\mu^{k(t)}$, $\ell^k(\mu^{(t)})$, choose $\mu^{k(t+1)}$

Outline

The Hedge algorithm

Hedge algorithm

- Update the distribution according to observed loss

$$\mu_p^{k(t+1)} \propto \mu_p^{k(t)} e^{-\eta \ell_p^{k(t)}}$$

Regret Bound

- Assume losses are in $[0, \rho]$.
- Expected loss is $\langle \mu^{k(t)}, \ell^k(\mu^{(t)}) \rangle$

$$R^{k(T)} = \sum_{t=1}^{T} \left\langle \mu^{k(t)}, \ell^k(\mu^{(t)}) \right\rangle - \min_p \sum_{t=1}^{T} \ell_p^{k(t)}$$

Regret Bound

- Assume losses are in $[0, \rho]$.
- Expected loss is $\langle \mu^{k(t)}, \ell^k(\mu^{(t)}) \rangle$

$$R^{k(T)} = \sum_{t=1}^{T} \left\langle \mu^{k(t)}, \ell^k(\mu^{(t)}) \right\rangle - \min_p \sum_{t=1}^{T} \ell_p^{k(t)}$$

Regret bound

$$\frac{R^{k(T)}}{T} \leq \frac{\rho \ln \mu_{\min}^{k(0)}}{T\eta} + \rho\eta$$

Motivation for regret: in general finite games, equilibria are characterized in terms of regret (more details later).

Convergence to approximate Nash equilibria

Regret bound

$$\frac{R^{k(T)}}{T} \leq \frac{\rho \ln \mu_{\min}^{k(0)}}{T\eta} + \rho\eta$$

Convergence to approximate Nash equilibria

If an update satisfies the regret bound, then for all $\epsilon > 0$, for η small enough, $\mu^{(t)}$ converges

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mu^{(t)} \in \mathcal{N}_\epsilon$$

$\mathcal{N}_\epsilon = \{\mu : V(\mu) < V_\mathcal{N} + \epsilon\}$: ϵ-approximate Nash equilibrium.

Convergence to approximate Nash equilibria

Regret bound

$$\frac{R^{k(T)}}{T} \leq \frac{\rho \ln \mu_{\min}^{k(0)}}{T\eta} + \rho\eta$$

Convergence to approximate Nash equilibria

If an update satisfies the regret bound, then for all $\epsilon > 0$, for η small enough, $\mu^{(t)}$ converges

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mu^{(t)} \in \mathcal{N}_\epsilon$$

$\mathcal{N}_\epsilon = \{\mu : V(\mu) < V_\mathcal{N} + \epsilon\}$: ϵ-approximate Nash equilibrium.

Proof: show

$$V(\bar{\mu}^{(T)}) - V(\mu^*) \leq \sum_k \frac{R^{k(T)}}{T}$$

Outline

Mirror Descent

Mirror Descent

Consider the convex problem

$$\text{minimize}_{\mu \in \Delta} V(\mu)$$

Algorithm 1 Mirror Descent Method

1: **for** $t \in \mathbb{N}$ **do**

2: $\mu^{(t+1)} = \arg\min_{\mu \in \Delta} \left\langle \nabla V(\mu^{(t)}), \mu \right\rangle + \frac{1}{\eta} D_\psi(\mu, \mu^{(t)})$

3: **end for**

where D_ψ is a Bregman divergence

$$D_\psi(\mu, \nu) = \psi(\mu) - \psi(\nu) - \langle \nabla \psi(\nu), \mu - \nu \rangle$$

Mirror descent on potential function

Mirror descent on V

- Take V potential function defined earlier. $\nabla_{\mu^k} V(\mu) = \ell^k(\mu)$

Mirror descent on potential function

Mirror descent on V

- Take V potential function defined earlier. $\nabla_{\mu^k} V(\mu) = \ell^k(\mu)$
- Take $\psi(\mu) = \sum_k \sum_p \mu_p^k \ln \mu_p^k$
- Then $D_\psi(\mu, \nu) = \sum_k D_{KL}(\mu^k, \nu^k)$

Mirror descent on potential function

Mirror descent on V

- Take V potential function defined earlier. $\nabla_{\mu^k} V(\mu) = \ell^k(\mu)$
- Take $\psi(\mu) = \sum_k \sum_p \mu_p^k \ln \mu_p^k$
- Then $D_\psi(\mu, \nu) = \sum_k D_{KL}(\mu^k, \nu^k)$
- Update:

$$\mu^{(t+1)} = \arg \min_{\mu \in \Delta^1 \times \cdots \times \Delta^K} \sum_k \left(\left\langle \nabla_{\mu^k} V(\mu), \mu^k \right\rangle + \frac{1}{\eta} D_{KL}(\mu^k, \mu^{k(t)}) \right)$$

Mirror descent on potential function

Mirror descent on V

- Take V potential function defined earlier. $\nabla_{\mu^k} V(\mu) = \ell^k(\mu)$
- Take $\psi(\mu) = \sum_k \sum_p \mu_p^k \ln \mu_p^k$
- Then $D_\psi(\mu, \nu) = \sum_k D_{KL}(\mu^k, \nu^k)$
- Update:

$$\mu^{(t+1)} = \arg \min_{\mu \in \Delta^1 \times \cdots \times \Delta^K} \sum_k \left(\left\langle \nabla_{\mu^k} V(\mu), \mu^k \right\rangle + \frac{1}{\eta} D_{KL}(\mu^k, \mu^{k(t)}) \right)$$

- Objective separates

$$\mu^{k(t+1)} = \arg \min_{\mu \in \Delta^k} \left\langle \nabla_{\mu^k} V(\mu), \mu^k \right\rangle + \frac{1}{\eta} D_\psi(\mu^k, \mu^{k(t)})$$

Mirror descent on potential function

Mirror descent on V

- Take V potential function defined earlier. $\nabla_{\mu^k} V(\mu) = \ell^k(\mu)$
- Take $\psi(\mu) = \sum_k \sum_p \mu_p^k \ln \mu_p^k$
- Then $D_\psi(\mu, \nu) = \sum_k D_{KL}(\mu^k, \nu^k)$
- Update:

$$\mu^{(t+1)} = \arg \min_{\mu \in \Delta^1 \times \cdots \times \Delta^K} \sum_k \left(\left\langle \nabla_{\mu^k} V(\mu), \mu^k \right\rangle + \frac{1}{\eta} D_{KL}(\mu^k, \mu^{k(t)}) \right)$$

- Objective separates

$$\mu^{k(t+1)} = \arg \min_{\mu \in \Delta^k} \left\langle \nabla_{\mu^k} V(\mu), \mu^k \right\rangle + \frac{1}{\eta} D_\psi(\mu^k, \mu^{k(t)})$$

- Solution: Hedge algorithm with learning rate η

$$\mu_p^{k(t+1)} \propto \mu_p^{k(t)} e^{-\eta \ell_p^{k(t)}}$$

Mirror Descent on potential function

Convergence of Mirror Descent method

$$V\left(\frac{1}{T}\sum_{t\leq T}\mu^{(t)}\right) - V(\mu^*) \leq \frac{1}{\eta T}\sum_{k} D_{KL}(\mu^{*k}, \mu^{k(0)}) + \frac{\eta\rho^2}{2}$$

No need to use a regret bound.

Mirror Descent on potential function

Convergence of Mirror Descent method

$$V\left(\frac{1}{T}\sum_{t\leq T}\mu^{(t)}\right) - V(\mu^*) \leq \frac{1}{\eta T}\sum_k D_{KL}(\mu^{*k}, \mu^{k(0)}) + \frac{\eta\rho^2}{2}$$

No need to use a regret bound.

General result

Convergence of $\bar{\mu}^{(T)} = \frac{1}{T}\sum_{t\leq T}\mu^{(t)}$ for

- Any potential game (that is, $\exists V$ convex such that $\nabla_{\mu^k} V(\mu) = \ell^k(\mu)$)
- Any Mirror Descent method (choose your favorite ψ, strongly convex w.r.t. $\|\cdot\|_1$)

to N_ϵ.

Mirror Descent with decreasing rates

MD with time-varying rates

$$\mu^{k(t+1)} = \arg\min_{\mu \in \Delta^k} \left\langle \nabla_{\mu^k} V(\mu), \mu^k \right\rangle + \frac{1}{\eta_t} D_\psi(\mu^k, \mu^{k(t)})$$

Mirror Descent with decreasing rates

MD with time-varying rates

$$\mu^{k(t+1)} = \arg\min_{\mu \in \Delta^k} \left\langle \nabla_{\mu^k} V(\mu), \mu^k \right\rangle + \frac{1}{\eta_t} D_\psi(\mu^k, \mu^{k(t)})$$

Then, can show

$$V\left(\frac{\sum_{t \le T} \eta_t \mu^{(t)}}{\sum_{t \le T} \eta_t} \right) - V(\mu^*) \le \frac{1}{\sum_{t \le T} \eta_t} D^{(0)} + \frac{\rho^2}{2} \frac{\sum_{t \le T} \eta_t^2}{\sum_{t \le T} \eta_t}$$

Mirror Descent with decreasing rates

MD with time-varying rates

$$\mu^{k(t+1)} = \arg\min_{\mu \in \Delta^k} \left\langle \nabla_{\mu^k} V(\mu), \mu^k \right\rangle + \frac{1}{\eta_t} D_\psi(\mu^k, \mu^{k(t)})$$

Then, can show

$$V\left(\frac{\sum_{t \le T} \eta_t \mu^{(t)}}{\sum_{t \le T} \eta_t}\right) - V(\mu^*) \le \frac{1}{\sum_{t \le T} \eta_t} D^{(0)} + \frac{\rho^2}{2} \frac{\sum_{t \le T} \eta_t^2}{\sum_{t \le T} \eta_t}$$

if $\eta_t = \frac{1}{\sqrt{t}}$, convergence in $\frac{\ln T}{\sqrt{T}}$

Discounted hedge algorithm

Have a similar bound for discounted regret

Discounted regret

$$R^{k(T)} = \sum_{t \leq T} \left\langle \mu^{k(t)}, \eta_t \ell^k(\mu^{(t)}) \right\rangle - \min_p \sum_{t=1}^{T} \eta_t \ell_p^{k(t)}$$

Interpretation: discount losses over time.

Discounted regret bound

$$\frac{R^{k(T)}}{\sum_{t \leq T} \eta_t} \leq \frac{\rho \log \mu_{\min}^{k(0)}}{\sum_{t \leq T} \eta_t} + \frac{\rho}{8} \frac{\sum_{t \leq T} \eta_t^2}{\sum_{t \leq T} \eta_t}$$

Outline

1. Learning in finite games
   - Equilibria
   - Regret
   - Blackwell approachability
2. Learning in convex potential games
   - The routing game and Nash equilibria
   - Regret
   - Mirror descent on potential function
   - Convergence of a dense subsequence
   - Strong convergence

Convergence of a dense subsequence

$$V\left(\frac{\sum_{t\leq T}\eta_t\mu^{(t)}}{\sum_{t\leq T}\eta_t}\right) - V(\mu^*) \leq \frac{1}{\sum_{t\leq T}\eta_t}D^{(0)} + \frac{\rho^2}{2}\frac{\sum_{t\leq T}\eta_t^2}{\sum_{t\leq T}\eta_t}$$

$$\bar{\mu}^{(t)} \to \mathcal{N}$$

Convergence of a dense subsequence

$$V\left(\frac{\sum_{t\le T}\eta_t\mu^{(t)}}{\sum_{t\le T}\eta_t}\right) - V(\mu^*) \le \frac{1}{\sum_{t\le T}\eta_t}D^{(0)} + \frac{\rho^2}{2}\frac{\sum_{t\le T}\eta_t^2}{\sum_{t\le T}\eta_t}$$

$$\bar{\mu}^{(t)} \to \mathcal{N}$$

Theorem

Under MD with appropriate η_t, a dense subsequence of $(\mu^{(t)})_t$ converges to $\mathcal{N}$

- Subsequence $(\mu^{(t)})_{t\in\mathcal{T}}$ converges
- $\lim_{T\to\infty} \frac{\sum_{t\in\mathcal{T}:t\le T}\eta_t}{\sum_{t\le T}\eta_t} = 1$

Convergence of a dense subsequence

$$V\left(\frac{\sum_{t\leq T}\eta_t\mu^{(t)}}{\sum_{t\leq T}\eta_t}\right) - V(\mu^*) \leq \frac{1}{\sum_{t\leq T}\eta_t}D^{(0)} + \frac{\rho^2}{2}\frac{\sum_{t\leq T}\eta_t^2}{\sum_{t\leq T}\eta_t}$$

$$\bar{\mu}^{(t)} \to \mathcal{N}$$

Theorem

Under MD with appropriate η_t, a dense subsequence of $(\mu^{(t)})_t$ converges to $\mathcal{N}$

- Subsequence $(\mu^{(t)})_{t\in\mathcal{T}}$ converges
- $\lim_{T\to\infty} \frac{\sum_{t\in\mathcal{T}:t\leq T}\eta_t}{\sum_{t\leq T}\eta_t} = 1$

Proof.

- Absolute Cesàro convergence implies convergence of a dense subsequence.

Simulations

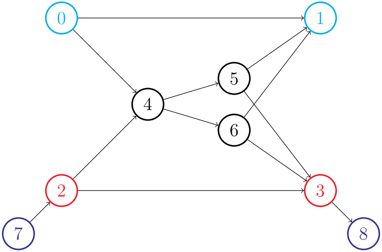

Figure : Example network

Simulations

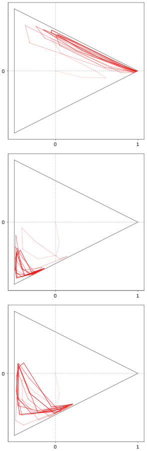

(a) Trajectories $(\mu^{k(t)})_t$.

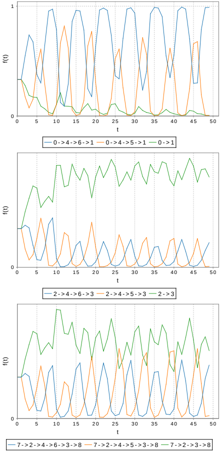

(b) Path flows $\mu_p^{k(t)}, p \in \mathcal{P}_k$

Figure : Constant learning rate $\eta = 0.7$. The trajectories do not converge.

Simulations

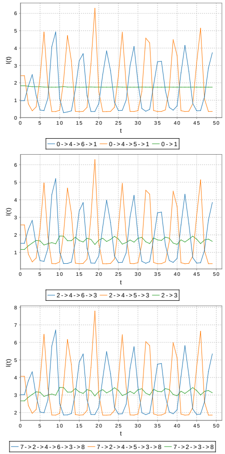

(a) Path losses $\ell_p^k(\mu^{(t)})$

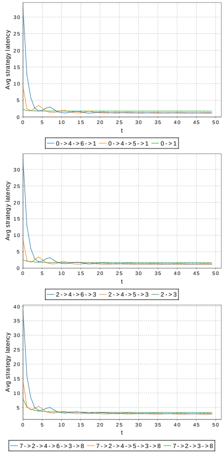

(b) Path losses for the means, $\ell_p^k(\frac{1}{t}\sum_{\tau \le t} \mu^{(\tau)})$

Figure : Path latencies

Simulations

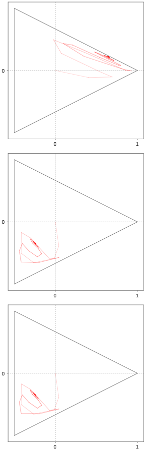

(a) Trajectories $(\mu^{k(t)})_t$.

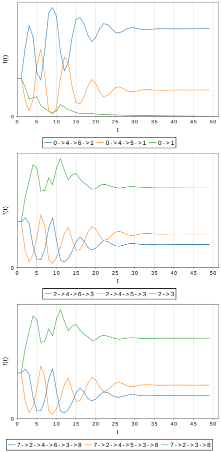

(b) Path flows $\mu_p^{k(t)}, p \in \mathcal{P}_k$

Figure : harmonic sequence of learning rates $\eta_t = \frac{1}{1+t/10}$

Simulations

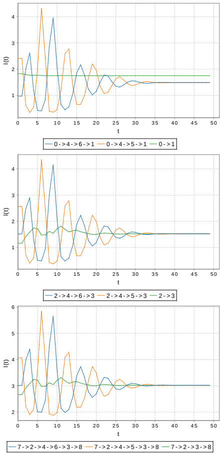

Figure : Path losses $\ell_p^k(\mu^{(t)})$

Outline

Sufficient conditions for convergence of $(\mu^{(t)})_t$

- Have $\bar{\mu}^{(t)} \to \mathcal{N}$.

Sufficient conditions for convergence of $(\mu^{(t)})_t$

- Have $\bar{\mu}^{(t)} \to \mathcal{N}$.

Sufficient condition 1

If $(\mu^{(t)})$ converges, then $\mu^{(t)} \to \mu^* \in \mathcal{N}$.

Sufficient conditions for convergence of $(\mu^{(t)})_t$

- Have $\bar{\mu}^{(t)} \to \mathcal{N}$.

Sufficient condition 1

If $(\mu^{(t)})$ converges, then $\mu^{(t)} \to \mu^* \in \mathcal{N}$.

Sufficient condition 2

If $V(\mu^{(t)})$ converges (weaker, $\mu^{(t)}$ need not converge), then

- $V(\mu^{(t)}) \to V_*$
- $\mu^{(t)} \to \mathcal{N}$ (V is continuous, $\mu \in \Delta$ compact)

Replicator dynamics

Imagine an underlying continuous time. Updates happen at $\eta_1, \eta_1 + \eta_2, \ldots$

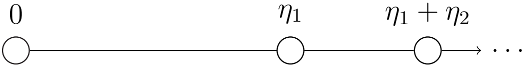

Figure : Underlying continuous time

Replicator dynamics

Imagine an underlying continuous time. Updates happen at $\eta_1, \eta_1 + \eta_2, \dots$

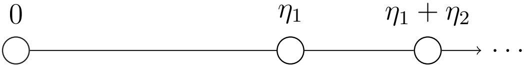

Figure : Underlying continuous time

In the update equation $\mu_p(t+1) \propto \mu_p(t) e^{-\eta_t \ell_p(t)}$, take $\eta_t \to 0$
We obtain the autonomous ODE:

Replicator equation

$$\forall p \in \mathcal{P}_k, \frac{d\mu_p^k}{dt} = \mu_p^k(\langle \ell^k(\mu), \mu^k \rangle - \ell_p^k(\mu))/\rho \quad (1)$$

Also in evolutionary game theory.

Replicator dynamics

Replicator equation

$$\forall p \in \mathcal{P}_k, \frac{d\mu_p^k}{dt} = \mu_p^k(\langle \ell^k(\mu), \mu^k \rangle - \ell_p^k(\mu))/\rho$$

Replicator dynamics

Replicator equation

$$\forall p \in \mathcal{P}_k, \frac{d\mu_p^k}{dt} = \mu_p^k(\langle \ell^k(\mu), \mu^k \rangle - \ell_p^k(\mu))/\rho$$

Theorem

Every solution of the ODE (1) converges to the set of its stationary points.

Replicator dynamics

Replicator equation

$$\forall p \in \mathcal{P}_k, \frac{d\mu_p^k}{dt} = \mu_p^k(\langle \ell^k(\mu), \mu^k \rangle - \ell_p^k(\mu))/\rho$$

Theorem

Every solution of the ODE (1) converges to the set of its stationary points.

Proof: V is a Lyapunov function.

REP method

Discretization of the continuous-time replicator dynamics

$$\mu_p^{k(t+1)} - \mu_p^{k(t)} = \eta_t \mu_p^{k(t)} \left(\frac{\langle \ell^k(\mu^{(t)}), \mu^{k(t)} \rangle - \ell_p^k(\mu^{(t)})}{\rho} \right) + \eta_t U_p^{k(t+1)}$$

$(U^{(t)})_{t \geq 1}$ perturbations that satisfy for all $T > 0$,

$$\lim_{\tau_1 \to \infty} \max_{\tau_2 : \sum_{t=\tau_1}^{\tau_2} \eta_t < T} \left\| \sum_{t=\tau_1}^{\tau_2} \eta_t U^{(t+1)} \right\| = 0$$

Convergence to Nash equilibria

Theorem

Consider a potential game with convex potential V.
Under any MD algorithm which is approximate REP, $\mu^{(t)} \to \mathcal{N}$.

Convergence to Nash equilibria

Theorem

Consider a potential game with convex potential V.
Under any MD algorithm which is approximate REP, $\mu^{(t)} \to \mathcal{N}$.

Proof uses two facts

- Affine interpolation of $\mu^{(t)}$ is an asymptotic pseudo trajectory for the ODE.
- V is a Lyapunov function for Nash equilibria.

REP update

In particular for $U = 0$, we obtain

$$\mu_p^{k(t+1)} - \mu_p^{k(t)} = \eta_t \mu_p^{k(t)} \left(\frac{\left\langle \ell^k(\mu^{(t)}), \mu^{k(t)} \right\rangle - \ell_p^k(\mu^{(t)})}{\rho} \right)$$

REP update

In particular for $U = 0$, we obtain

$$\mu_p^{k(t+1)} - \mu_p^{k(t)} = \eta_t \mu_p^{k(t)} \left(\frac{\left\langle \ell^k(\mu^{(t)}), \mu^{k(t)} \right\rangle - \ell_p^k(\mu^{(t)})}{\rho} \right)$$

- Also MD algorithm:

$$\mu^{(t+1)} \in \arg\min_{\mu \in \Delta} \left\langle \frac{\ell(\mu^{(t)})}{\rho}, \mu \right\rangle + \frac{1}{\eta_t} D(\mu \| \mu^{(t)})$$

- $D(\mu||\nu) = \frac{1}{2} \sum_p \nu_p \left(\frac{\mu_p}{\nu_p} - 1 \right)^2$ (KKT conditions)

That's all, folks!

A convex formulation of Nash equilibria

μ is a Nash equilibrium iff it minimizes a potential function

$$\begin{aligned} &\text{minimize}_{\mu \geq 0, \phi \geq 0} && \sum_e \int_0^{\phi_e} \ell_e(u) du \\ &\text{subject to} && \forall e, \sum_{p \ni e} \mu_p = \phi_e \quad \sum_p \mu_p = 1 \end{aligned}$$

A convex formulation of Nash equilibria

μ is a Nash equilibrium iff it minimizes a potential function

$$\begin{aligned} \text{minimize}_{\mu \geq 0, \phi \geq 0} \quad & \sum_e \int_0^{\phi_e} \ell_e(u) du \\ \text{subject to} \quad & \forall e, \ \sum_{p \ni e} \mu_p = \phi_e \quad \sum_p \mu_p = 1 \end{aligned}$$

partial Lagrangian

$$\sum_e \int_0^{\phi_e} \ell_e(u) du + \sum_e v_e \left(\sum_{p \ni e} \mu_p - \phi_e \right) - w \left(\sum_p \mu_p - 1 \right) - \sum_e \lambda_e \phi_e$$

A convex formulation of Nash equilibria

μ is a Nash equilibrium iff it minimizes a potential function

$$\text{minimize}_{\mu \geq 0, \phi \geq 0} \quad \sum_e \int_0^{\phi_e} \ell_e(u) du$$

$$\text{subject to} \qquad \forall e, \sum_{p \ni e} \mu_p = \phi_e \qquad \sum_p \mu_p = 1$$

partial Lagrangian

$$\sum_e \int_0^{\phi_e} \ell_e(u) du + \sum_e v_e \left(\sum_{p \ni e} \mu_p - \phi_e \right) - w \left(\sum_p \mu_p - 1 \right) - \sum_e \lambda_e \phi_e$$

A convex formulation of Nash equilibria

μ is a Nash equilibrium iff it minimizes a potential function

$$\text{minimize}_{\mu \geq 0, \phi \geq 0} \quad \sum_e \int_0^{\phi_e} \ell_e(u) du$$

$$\text{subject to} \quad \forall e, \sum_{p \ni e} \mu_p = \phi_e \quad \sum_p \mu_p = 1$$

partial Lagrangian

$$\sum_e \int_0^{\phi_e} \ell_e(u) du + \sum_e v_e \left(\sum_{p \ni e} \mu_p - \phi_e \right) - w \left(\sum_p \mu_p - 1 \right) - \sum_e \lambda_e \phi_e$$

A convex formulation of Nash equilibria

μ is a Nash equilibrium iff it minimizes a potential function

$$\text{minimize}_{\mu \geq 0, \phi \geq 0} \quad \sum_e \int_0^{\phi_e} \ell_e(u)du$$
$$\text{subject to} \quad \forall e, \sum_{p \ni e} \mu_p = \phi_e \quad \sum_p \mu_p = 1$$

partial Lagrangian

$$\sum_e \int_0^{\phi_e} \ell_e(u)du + \sum_e v_e \left(\sum_{p \ni e} \mu_p - \phi_e \right) - w \left(\sum_p \mu_p - 1 \right) - \sum_e \lambda_e \phi_e$$

A convex formulation of Nash equilibria

μ is a Nash equilibrium iff it minimizes a potential function

$$\begin{aligned} \text{minimize}_{\mu \geq 0, \phi \geq 0} \quad & \sum_e \int_0^{\phi_e} \ell_e(u) du \\ \text{subject to} \quad & \forall e, \sum_{p \ni e} \mu_p = \phi_e \quad \sum_p \mu_p = 1 \end{aligned}$$

partial Lagrangian

$$\sum_e \int_0^{\phi_e} \ell_e(u) du + \sum_e v_e \left(\sum_{p \ni e} \mu_p - \phi_e \right) - w \left(\sum_p \mu_p - 1 \right) - \sum_e \lambda_e \phi_e$$

Optimality conditions

- stationarity

$$\begin{aligned} \forall e, \quad & \ell_e(\phi_e) - v_e - \lambda_e = 0 \\ \forall p, \quad & \sum_{e \in p} v_e - w = 0 \end{aligned}$$

- complementary slackness: $\lambda_e \phi_e = 0$ for all e

A convex formulation of Nash equilibria

- stationarity

$$\forall e, \qquad \ell_e(\phi_e) - v_e - \lambda_e = 0$$
$$\forall p, \qquad \sum_{e \in p} v_e - w = 0$$

- complementary slackness: $\lambda_e \phi_e = 0$ for all e

for all paths p, the cost on the path is

$$\ell_p(\mu) = \sum_{e \in p} \ell(\phi_e)$$

A convex formulation of Nash equilibria

- stationarity

$$\forall e, \qquad \ell_e(\phi_e) - v_e - \lambda_e = 0$$

$$\forall p, \qquad \sum_{e \in p} v_e - w = 0$$

- complementary slackness: $\lambda_e \phi_e = 0$ for all e

for all paths p, the cost on the path is

$$\ell_p(\mu) = \sum_{e \in p} \ell(\phi_e) = \sum_{e \in p} (v_e + \lambda_e)$$

A convex formulation of Nash equilibria

- stationarity

$$\forall e, \qquad \ell_e(\phi_e) - v_e - \lambda_e = 0$$
$$\forall p, \qquad \sum_{e \in p} v_e - w = 0$$

- complementary slackness: $\lambda_e \phi_e = 0$ for all e

for all paths p, the cost on the path is

$$\ell_p(\mu) = \sum_{e \in p} \ell(\phi_e) = \sum_{e \in p} (v_e + \lambda_e) = w + \sum_{e \in p} \lambda_e$$

A convex formulation of Nash equilibria

- stationarity

$$\forall e, \qquad \ell_e(\phi_e) - v_e - \lambda_e = 0$$

$$\forall p, \qquad \sum_{e \in p} v_e - w = 0$$

- complementary slackness: $\lambda_e \phi_e = 0$ for all e

for all paths p, the cost on the path is

$$\ell_p(\mu) = \sum_{e \in p} \ell(\phi_e) = \sum_{e \in p} (v_e + \lambda_e) = w + \sum_{e \in p} \lambda_e$$

- if $\mu_p > 0$, $\ell_p(\mu) = w$
- if $\mu_p = 0$, $\ell_p(\mu) \geq w$

A convex formulation of Nash equilibria

- stationarity

$$\forall e, \qquad \ell_e(\phi_e) - v_e - \lambda_e = 0$$

$$\forall p, \qquad \sum_{e \in p} v_e - w = 0$$

- complementary slackness: $\lambda_e \phi_e = 0$ for all e

for all paths p, the cost on the path is

$$\ell_p(\mu) = \sum_{e \in p} \ell(\phi_e) = \sum_{e \in p} (v_e + \lambda_e) = w + \sum_{e \in p} \lambda_e$$

- if $\mu_p > 0$, $\ell_p(\mu) = w$
- if $\mu_p = 0$, $\ell_p(\mu) \geq w$

This is the set of Nash equilibria