Recurrence vs Convergence

A Geometric Approach to Learning in Games

Davide Legacci, Panayotis Mertikopoulos, Bary Pradelski

October 13, 2023 - SLMath, MMD Seminar

2023-10-13

Recurrence vs Convergence

Recurrence vs Convergence
A Geometric Approach to Learning in Games
Davide Legacci, Panayotis Mertikopoulos, Bary Pradelski
October 13, 2023 - SLMath, MMD Seminar

Mission

What

What is the common property of games that are hard to learn?

How

Mission
What
What is the common property of games that are hard to learn?
How
Mirror descent overview
Convergence and cycles
Combinatorial decomposition for finite normal form games
Application: Two-players first-price sealed-bid auction
Decomposition for general games

- we have algorithms that on some games exhibit good convergence properties
- for example, mirror descent *sometimes* converges to NE (more on this later)
- but we often do not know why! Cf Martin's work here
- Goal: identify properties of games that can explain (non) convergence
- classify games that are intrinsically **hard to learn**
- propose learning algo. that minimizes cycling
- **How?** Two decomposition techniques; one review, one original

Mirror descent overview

Convergence and cycles

Combinatorial decomposition for finite normal form games

Decomposition for general games

Mirror descent overview
Convergence and cycles
Combinatorial decomposition for finite normal form games
Decomposition for general games

Finite normal form games

Goal

Introduce payoff and simultaneous gradient of finite normal form game

Why?

Used in learning algorithm and in geometrical decomposition

Finite normal form games
Goal
Introduce payoff and simultaneous gradient of finite normal form game
Why?
Used in learning algorithm and in geometrical decomposition

Finite normal form game $\Gamma = (\mathcal{N}, \mathcal{A}, u)$

- $\mathcal{N} = \{1, 2, \ldots, N\}$ set of **players**, index i
- Set of **pure strategies** $\mathcal{A}_i = \{1, 2, \ldots, A_i\}$ for each player
- $\mathcal{A} = \prod_{i \in \mathcal{N}} \mathcal{A}_i$ set of **pure strategy profiles**
- Payoff

$$u : \mathcal{A} \to \mathbb{R}^N, \quad a \mapsto (u_1, \ldots, u_N)(a)$$

- $u_i(a)$ = payoff of player $i \in \mathcal{N}$ at pure strategy profile $a \in \mathcal{A}$

Finite normal form game $\Gamma = (\mathcal{N}, \mathcal{A}, u)$
- $\mathcal{N} = \{1, 2, \ldots, N\}$ set of **players**, index i
- Set of **pure strategies** $\mathcal{A}_i = \{1, 2, \ldots, A_i\}$ for each player
- $\mathcal{A} = \prod_{i \in \mathcal{N}} \mathcal{A}_i$ set of **pure strategy profiles**
- Payoff
$u : \mathcal{A} \to \mathbb{R}^N, \quad a \mapsto (u_1, \ldots, u_N)(a)$
- $u_i(a)$ = payoff of player $i \in \mathcal{N}$ at pure strategy profile $a \in \mathcal{A}$

- think of $\mathcal{A}$ as the space of states of the game; an element is a tuple that contains one strategy for each player
- given a strategy profile - a state - each player gets some payoff…
- …and putting these together we get the global payoff
- this is the object we're interested in decomposing

Finite normal form game $\Gamma = (\mathcal{N}, \mathcal{A}, u)$

- $\mathcal{N} = \{1, 2, \ldots, N\}$ set of **players**, index i
- Set of **pure strategies** $\mathcal{A}_i = \{1, 2, \ldots, A_i\}$ for each player
- $\mathcal{A} = \prod_{i \in \mathcal{N}} \mathcal{A}_i$ set of **pure strategy profiles**

- Payoff

$$u : \mathcal{A} \to \mathbb{R}^N, \quad a \mapsto (u_1, \ldots, u_N)(a)$$

- $u_i(a)$ = payoff of player $i \in \mathcal{N}$ at pure strategy profile $a \in \mathcal{A}$

Finite normal form game $\Gamma = (\mathcal{N}, \mathcal{A}, u)$
- $\mathcal{N} = \{1, 2, \ldots, N\}$ set of players, index i
- Set of pure strategies $\mathcal{A}_i = \{1, 2, \ldots, A_i\}$ for each player
- $\mathcal{A} = \prod_{i \in \mathcal{N}} \mathcal{A}_i$ set of pure strategy profiles
- Payoff
$u : \mathcal{A} \to \mathbb{R}^N, \quad a \mapsto (u_1, \ldots, u_N)(a)$
- $u_i(a)$ = payoff of player $i \in \mathcal{N}$ at pure strategy profile $a \in \mathcal{A}$

- think of $\mathcal{A}$ as the space of states of the game; an element is a tuple that contains one strategy for each player
- given a strategy profile - a state - each player gets some payoff...
- ...and putting these together we get the global payoff
- this is the object we're interested in decomposing

Mixed Extension

Mixed strategy: probability distribution over pure strategies

$$\text{for each player } i \in \mathcal{N}, \quad x_i \in \mathcal{X}_i = \Delta(\mathcal{A}_i) \subset \mathcal{V}_i = \mathbb{R}^{A_i}$$

Expected payoff

$$\bar{u}_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}, \quad \underbrace{(x_1, \dots, x_N)}_{\text{mixed strategy profile}} \mapsto \mathbb{E}_{a \sim x}[u_i(a)]$$

Simultaneous gradient $v_i : \mathcal{X} \to \mathcal{V}_i^* = \mathbb{R}^{A_i}$

Gradient of exp. payoff $\bar{u}_i$ w.r.t. mixed strategy x_i of player i

$$v_i(x) := \left(\frac{\partial \bar{u}_i(x)}{\partial x_{i,a_i}} \right)_{a_i \in \mathcal{A}_i} \in \mathbb{R}^{A_i}$$

$$v_i(x) \cdot x_i = \bar{u}_i(x) \in \mathbb{R}$$

- **Expected payoff**: expectation value of $u_i(a)$ where the pure strategy profile a is drawn according to the probability distribution x

$$u_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}, \quad \underbrace{(x_1, \dots, x_N)}_{\text{mixed strategy profile}} \mapsto \mathbb{E}_{a \sim x}[u_i(a)] = \sum_{a \in \mathcal{A}} u_i(a) \underbrace{\prod_{j \in \mathcal{N}} x_{j,a_j}}_{\mathrm{P}_x(a)}$$

- Take $\mathcal{V}_i^*$ as notation to distinguish strategy space from payoff space; deeper reason why $v_i \in \mathcal{V}_i^*$ is that v_i is actually a differential, not a gradient, and as such it lives in the dual space

Mixed Extension

Mixed strategy: probability distribution over pure strategies

$$\text{for each player } i \in \mathcal{N}, \quad x_i \in \mathcal{X}_i = \Delta(\mathcal{A}_i) \subset \mathcal{V}_i = \mathbb{R}^{A_i}$$

Expected payoff

$$\bar{u}_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}, \quad \underbrace{(x_1, \ldots, x_N)}_{\text{mixed strategy profile}} \quad \mapsto \mathbb{E}_{a \sim x}[u_i(a)]$$

Simultaneous gradient $v_i : \mathcal{X} \to \mathcal{V}_i^* = \mathbb{R}^{A_i}$

Gradient of exp. payoff $\bar{u}_i$ w.r.t. mixed strategy x_i of player i

$$v_i(x) := \left(\frac{\partial \bar{u}_i(x)}{\partial x_{i,a_i}} \right)_{a_i \in \mathcal{A}_i} \in \mathbb{R}^{A_i}$$

$$v_i(x) \cdot x_i = \bar{u}_i(x) \in \mathbb{R}$$

- **Expected payoff**: expectation value of $u_i(a)$ where the pure strategy profile a is drawn according to the probability distribution x

$$u_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}, \quad \underbrace{(x_1, \ldots, x_N)}_{\text{mixed strategy profile}} \quad \mapsto \mathbb{E}_{a \sim x}[u_i(a)] = \sum_{a \in \mathcal{A}} u_i(a) \underbrace{\prod_{j \in \mathcal{N}} x_{j,a_j}}_{P_x(a)}$$

- Take $\mathcal{V}_i^*$ as notation to distinguish strategy space from payoff space; deeper reason why $v_i \in \mathcal{V}_i^*$ is that v_i is actually a differential, not a gradient, and as such it lives in the dual space

Mixed Extension

Mixed strategy: probability distribution over pure strategies

$$\text{for each player } i \in \mathcal{N}, \quad x_i \in \mathcal{X}_i = \Delta(\mathcal{A}_i) \subset \mathcal{V}_i = \mathbb{R}^{A_i}$$

Expected payoff

$$\bar{u}_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}, \quad \underbrace{(x_1, \ldots, x_N)}_{\text{mixed strategy profile}} \mapsto \mathbb{E}_{a \sim x}[u_i(a)]$$

Simultaneous gradient $v_i : \mathcal{X} \to \mathcal{V}_i^* = \mathbb{R}^{A_i}$

Gradient of exp. payoff $\bar{u}_i$ w.r.t. mixed strategy x_i of player i

$$v_i(x) := \left(\frac{\partial \bar{u}_i(x)}{\partial x_{i,a_i}} \right)_{a_i \in \mathcal{A}_i} \in \mathbb{R}^{A_i}$$

$$v_i(x) \cdot x_i = \bar{u}_i(x) \in \mathbb{R}$$

Mixed Extension
Mixed strategy: probability distribution over pure strategies
for each player $i \in \mathcal{N}$, $x_i \in \mathcal{X}_i = \Delta(\mathcal{A}_i) \subset \mathcal{V}_i = \mathbb{R}^{A_i}$
Expected payoff
$\bar{u}_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}$, $(x_1, \ldots, x_N) \mapsto \mathbb{E}_{a \sim x}[u_i(a)]$
Simultaneous gradient $v_i : \mathcal{X} \to \mathcal{V}_i^* = \mathbb{R}^{A_i}$
Gradient of exp. payoff $\bar{u}_i$ w.r.t. mixed strategy x_i of player i
$v_i(x) := \left(\frac{\partial \bar{u}_i(x)}{\partial x_{i,a_i}} \right)_{a_i \in \mathcal{A}_i} \in \mathbb{R}^{A_i}$

- **Expected payoff**: expectation value of $u_i(a)$ where the pure strategy profile a is drawn according to the probability distribution x

$$u_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}, \quad \underbrace{(x_1, \ldots, x_N)}_{\text{mixed strategy profile}} \mapsto \mathbb{E}_{a \sim x}[u_i(a)] = \sum_{a \in \mathcal{A}} u_i(a) \underbrace{\prod_{j \in \mathcal{N}} x_{j,a_j}}_{P_x(a)}$$

- Take $\mathcal{V}_i^*$ as notation to distinguish strategy space from payoff space; deeper reason why $v_i \in \mathcal{V}_i^*$ is that v_i is actually a differential, not a gradient, and as such it lives in the dual space

Mixed Extension

Mixed strategy: probability distribution over pure strategies

$$\text{for each player } i \in \mathcal{N}, \quad x_i \in \mathcal{X}_i = \Delta(\mathcal{A}_i) \subset \mathcal{V}_i = \mathbb{R}^{A_i}$$

Expected payoff

$$\bar{u}_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}, \quad \underbrace{(x_1, \ldots, x_N)}_{\text{mixed strategy profile}} \mapsto \mathbb{E}_{a \sim x}[u_i(a)]$$

Simultaneous gradient $v_i : \mathcal{X} \to \mathcal{V}_i^* = \mathbb{R}^{A_i}$

Gradient of exp. payoff $\bar{u}_i$ w.r.t. mixed strategy x_i of player i

$$v_i(x) := \left(\frac{\partial \bar{u}_i(x)}{\partial x_{i,a_i}} \right)_{a_i \in \mathcal{A}_i} \in \mathbb{R}^{A_i}$$

$$v_i(x) \cdot x_i = \bar{u}_i(x) \in \mathbb{R}$$

Mixed Extension
Mixed strategy: probability distribution over pure strategies
for each player $i \in \mathcal{N}$, $x_i \in \mathcal{X}_i = \Delta(\mathcal{A}_i) \subset \mathcal{V}_i = \mathbb{R}^{A_i}$
Expected payoff
$\bar{u}_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}$, $(x_1, \ldots, x_N) \mapsto \mathbb{E}_{a \sim x}[u_i(a)]$ mixed strategy profile
Simultaneous gradient $v_i : \mathcal{X} \to \mathcal{V}_i^* = \mathbb{R}^{A_i}$
Gradient of exp. payoff $\bar{u}_i$ w.r.t. mixed strategy x_i of player i
$v_i(x) := \left(\frac{\partial \bar{u}_i(x)}{\partial x_{i,a_i}} \right)_{a_i \in \mathcal{A}_i} \in \mathbb{R}^{A_i}$
$v_i(x) \cdot x_i = \bar{u}_i(x) \in \mathbb{R}$

- **Expected payoff**: expectation value of $u_i(a)$ where the pure strategy profile a is drawn according to the probability distribution x

$$u_i : \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}, \quad \underbrace{(x_1, \ldots, x_N)}_{\text{mixed strategy profile}} \mapsto \mathbb{E}_{a \sim x}[u_i(a)] = \sum_{a \in \mathcal{A}} u_i(a) \underbrace{\prod_{j \in \mathcal{N}} x_{j,a_j}}_{\mathrm{P}_x(a)}$$

- Take $\mathcal{V}_i^*$ as notation to distinguish strategy space from payoff space; deeper reason why $v_i \in \mathcal{V}_i^*$ is that v_i is actually a differential, not a gradient, and as such it lives in the dual space

Example - (2×2) game

Strategies $\mathcal{A}_1 = \{A, B\}, \mathcal{A}_2 = \{a, b\}$

Mixed strategy

$$x = (x_1, x_2) \in \mathcal{X}_1 \times \mathcal{X}_2$$

$$x_1 = (x_A, x_B), \quad x_2 = (x_a, x_b)$$

Expected payoff for player 1

$$\bar{u}_1(x) = u_1(A, a)\, x_A x_a + u_1(A, b)\, x_A x_b + u_1(B, a)\, x_B x_a + u_1(B, b)\, x_B x_b \in \mathbb{R}$$

Simultaneous gradient first player $v_1(x) = (\partial_{x_A} \bar{u}_1, \partial_{x_B} \bar{u}_1)$

$$v_1(x) = \Big(u_1(A, a)\, x_a + u_1(A, b)\, x_b,\; u_1(B, a)\, x_a + u_1(B, b)\, x_b\Big)$$

$$v_1(x) \cdot x_1 = \bar{u}_1(x)$$

Example - (2×2) game
Strategies $\mathcal{A}_1 = \{A, B\}, \mathcal{A}_2 = \{a, b\}$
Mixed strategy
$$x = (x_1, x_2) \in \mathcal{X}_1 \times \mathcal{X}_2$$
$$x_1 = (x_A, x_B), \quad x_2 = (x_a, x_b)$$
Expected payoff for player 1
$$\bar{u}_1(x) = u_1(A, a)\, x_A x_a + u_1(A, b)\, x_A x_b + u_1(B, a)\, x_B x_a + u_1(B, b)\, x_B x_b \in \mathbb{R}$$
Simultaneous gradient first player $v_1(x) = (\partial_{x_A} \bar{u}_1, \partial_{x_B} \bar{u}_1)$
$$v_1(x) = \Big(u_1(A, a)\, x_a + u_1(A, b)\, x_b,\; u_1(B, a)\, x_a + u_1(B, b)\, x_b\Big)$$
$$v_1(x) \cdot x_1 = \bar{u}_1(x)$$

Recap

Given finite normal form game $\Gamma = (\mathcal{N}, \mathcal{A}, u)$

Payoff $u : \mathcal{A} \to \mathbb{R}^N$

- Object of combinatorial decomposition

Simultaneous gradient $v_i : \mathcal{X} \to \mathbb{R}^{A_i}$

- Used to define learning dynamics
- Object of smooth decomposition

Recap

Given finite normal form game $\Gamma = (\mathcal{N}, \mathcal{A}, u)$

Payoff $u : \mathcal{A} \to \mathbb{R}^N$

- Object of combinatorial decomposition

Simultaneous gradient $v_i : \mathcal{X} \to \mathbb{R}^{A_i}$

- Used to define learning dynamics
- Object of smooth decomposition

- payoff takes pure strategy profile and gives corresponding payoff for each player
- sim gradient takes mixed strategy profile x and gives A_i numbers that dotted against $x_i \in \mathcal{X}_i$ (which also contains A_i entries) gives the number $\bar{u}_i(x)$

Continuous Time Mirror Descent

- Individual payoff $\bar{u}_i(x)$ depends on strategy of all agents
- Continuous-time, deterministic, multi-agent decision processes
- Agents aim at maximizing their payoff
- Future strategy depends on cumulative incurred payoff
- Choice map : cumulative payoff $\mapsto$ next mixed strategy

Continuous Time Mirror Descent
- Individual payoff $\bar{u}_i(x)$ depends on strategy of all agents
- Continuous-time, deterministic, multi-agent decision processes
- Agents aim at maximizing their payoff
- Future strategy depends on cumulative incurred payoff
- Choice map : cumulative payoff $\mapsto$ next mixed strategy

Continuous Time Mirror Descent

- Individual payoff $\bar{u}_i(x)$ depends on strategy of all agents
- Continuous-time, deterministic, multi-agent decision processes
- Agents aim at maximizing their payoff
- Future strategy depends on cumulative incurred payoff
- Choice map : cumulative payoff $\mapsto$ next mixed strategy

Continuous Time Mirror Descent

- Individual payoff $\bar{u}_i(x)$ depends on strategy of all agents
- Continuous-time, deterministic, multi-agent decision processes
- Agents aim at maximizing their payoff
- Future strategy depends on cumulative incurred payoff
- Choice map : cumulative payoff $\mapsto$ next mixed strategy

Continuous Time Mirror Descent

$$\begin{cases} y_i(t) &= \overbrace{\int_0^t v_i\left(x(s)\right) \mathrm{d}s}^{\text{cumulative payoff of player } i} \\ x_i(t) &= \underbrace{Q_i\left(y_i(t)\right)}_{\text{Choice map, to be defined}} \end{cases}$$

$$\boxed{\begin{cases} \dot{y}_i(t) &= v_i\left(x(t)\right) \\ x_i(t) &= Q_i\left(y_i(t)\right) \end{cases}} \tag{MD}$$

- $x_i(t) \in \mathcal{X}_i$ is mixed strategy of player i at time t
- $y_i(t) \in \mathcal{V}_i^*$ aggregates payoffs of player i until time t
- Aggregate payoff used to update strategy via choice map Q

Continuous Time Mirror Descent

$$\begin{cases} y_i(t) &= \overbrace{\int_0^t v_i\left(x(s)\right) \mathrm{d}s}^{\text{cumulative payoff of player } i} \\ x_i(t) &= \underbrace{Q_i\left(y_i(t)\right)}_{\text{Choice map, to be defined}} \end{cases}$$

$$\boxed{\begin{cases} \dot{y}_i(t) &= v_i\left(x(t)\right) \\ x_i(t) &= Q_i\left(y_i(t)\right) \end{cases}} \tag{MD}$$

- $x_i(t) \in \mathcal{X}_i$ is mixed strategy of player i at time t
- $y_i(t) \in \mathcal{V}_i^*$ aggregates payoffs of player i until time t
- Aggregate payoff used to update strategy via choice map Q

- Given $(\mathcal{N}, \mathcal{A}, u)$
- ...

Choice map and regularizer

Player's set of optimal strategies given mixed strategy profile $x \in \mathcal{X}$

$$\arg\max_{x_i \in \mathcal{X}_i} \{v_i(x) \cdot x_i\}$$

Introduce nice[1] regularizer

$$h : \mathcal{X} \to \mathbb{R}$$

so that choice map is well-defined

$$\begin{aligned} Q_i : \mathcal{V}_i^* &\to \mathcal{X}_i \\ v_i &\longmapsto \arg\max_{x_i \in \mathcal{X}_i} \{v_i \cdot x_i - h(x)\} \end{aligned}$$

[1] smooth, strongly convex, steep; for the non-steep case see [11]

Choice map and regularizer
Player's set of optimal strategies given mixed strategy profile $x \in \mathcal{X}$
$$\arg\max_{x_i \in \mathcal{X}_i} \{v_i(x) \cdot x_i\}$$
Introduce nice[1] regularizer
$$h : \mathcal{X} \to \mathbb{R}$$
so that choice map is well-defined
$$\begin{aligned} Q_i : \mathcal{V}_i^* &\to \mathcal{X}_i \\ v_i &\longmapsto \arg\max_{x_i \in \mathcal{X}_i} \{v_i \cdot x_i - h(x)\} \end{aligned}$$
[1] smooth, strongly convex, steep; for the non-steep case see [11]

- Steep: $\|dh(x_n)\| \to \infty$ at the boundary of $\mathcal{X}$

Example - Exponential MD and Replicator Dynamics [19, 22, 1]

- Entropic regularizer $h(x) = x \cdot \log(x)$
- Induces **logit** choice map

$$Q(y) = \frac{e^y}{e^y \cdot \mathbf{1}}$$

For each player $\dot{y} = v(x)$ and $x = Q(y)$ gives

$$\boxed{\dot{x}_{i,a_i} = x_{i,a_i}\Big(v_{i,a_i}(x) - \bar{u}_i(x)\Big)} \qquad \text{(RE)}$$

Example - Exponential MD and Replicator Dynamics [19, 22, 1]
- Entropic regularizer $h(x) = x \cdot \log(x)$
- Induces logit choice map
$Q(y) = \frac{e^y}{e^y \cdot \mathbf{1}}$
For each player $\dot{y} = v(x)$ and $x = Q(y)$ gives
$\dot{x}_{i,a_i} = x_{i,a_i}\Big(v_{i,a_i}(x) - \bar{u}_i(x)\Big)$ (RE)

- Taylor Jonker 1978
- e.g. prisoner's dilemma
- replicator name: the probability to use a pure strategy grows if at the current game state the payoff of using such pure strategy is higher than the expected payoff
- next we look at some convergence and non-convergence known properties of (MD) on finite normal form games

Convergence in potential games under (MD)

Theorem (Mertikopoulos and Sandholm [11])

If $x(t) \to x^*$ as $t \to \infty$ under (MD), then x^* is Nash equilibrium

Archetypal example: potential games [14, 20]

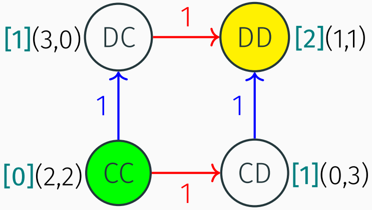

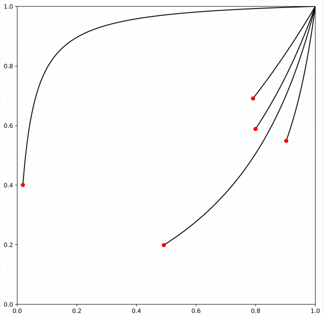

$$u_i(b) - u_i(a) = \phi(b) - \phi(a)$$

for each player and each unilateral deviation

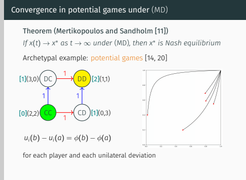

- Monderer Shapley 1996
- Left: RG. Nodes = pure st, edges = unil dev, on edge = payoff diff of deviating player. Exact potential game: there exists scalar function...
- right: dynamics in mixes strategy space for different initial conditions converges to same pure strategy. On the axes the probability of each player to play D, so that the mixed strategy $x = (x_1, x_2) = ((x_{1,D}, x_{1,c}), (x_{2,D}, x_{2,c}))$ converges to $((1,0),(1,0))$
- **convergence to pure NE, max of potential**
- yellow NE, green Pareto Efficient (there is no pareto improvement. PO = str change (not unilateral) st at least one is better and noone is worse

Cycles in zero-sum games under (MD)

Theorem (Mertikopoulos, Papadimitriou, and Piliouras [10])

Almost every solution trajectory $x(t)$ under (MD) is Poincaré recurrent on 2-player zero-sum games with an interior NE.

Theorem (L., Benedetti, Alishah, Mertikopoulos (wp))

(MD) *dynamics on 2-player zero-sum games with an interior NE are Hamiltonian.*

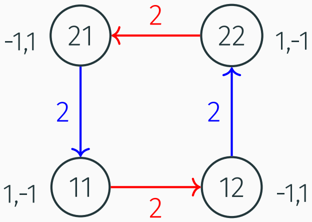

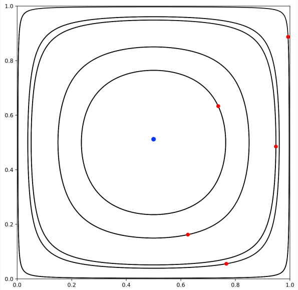

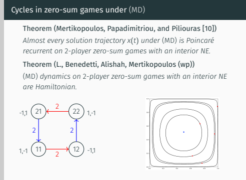

- e.g. matching pennies
- right: again different trajectories for different initial conditions
- NE uniformly mixed $(0.5, 0.5), (0.5, 0.5)$

Questions

- How "close" a generic game is to a potential game?
- Does this measure say anything about convergence?
- Which is the key property making a game "hard to learn", i.e. displaying cycles?

$\Rightarrow$ Decomposition of games

potential component + cycling component

Questions
- How "close" a generic game is to a potential game?
- Does this measure say anything about convergence?
- Which is the key property making a game "hard to learn", i.e. displaying cycles?
$\Rightarrow$ Decomposition of games
potential component + cycling component

Questions

- How “close” a generic game is to a potential game?
- Does this measure say anything about convergence?
- Which is the key property making a game “hard to learn”, i.e. displaying cycles?

$\Rightarrow$ Decomposition of games

$$\text{potential component} + \text{cycling component}$$

Questions
- How “close” a generic game is to a potential game?
- Does this measure say anything about convergence?
- Which is the key property making a game “hard to learn”, i.e. displaying cycles?
$\Rightarrow$ Decomposition of games
potential component + cycling component

Mirror descent overview

Combinatorial decomposition for finite normal form games

Application: Two-players first-price sealed-bid auction

Decomposition for general games

Mirror descent overview
Combinatorial decomposition for finite normal form games
Application: Two-players first-price sealed-bid auction
Decomposition for general games

Theorem (Candogan et al. [3])

Any finite normal form game $(\mathcal{N}, \mathcal{A}, u)$ admits an orthogonal decomposition

$$u = u_{\mathcal{K}} + u_{\mathcal{P}} + u_{\mathcal{H}}$$

- $u_{\mathcal{K}}$ is a normalization component that does not affect the dynamics
- $u_{\mathcal{P}}$ is a potential game
- $u_{\mathcal{H}}$ is an harmonic game

$\implies$ candidate obstacle to convergence

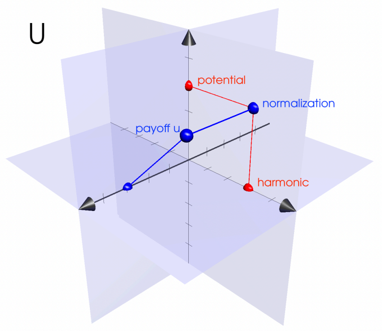

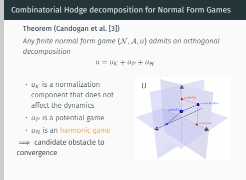

- rather than formal def let me give you intuition of what harmonic games are by example on auctions

Application: Two-players first-price sealed-bid auction

- Two bidders assign a value to a good and place a bid
- Higher bidder wins and pays their bid
- Possibly different values
- Discretization of continuous bids interval

$$u_i(x_i, x_j) = \begin{cases} v_i - x_i & \text{if } x_i > x_j \\ \frac{v_i - x_i}{2} & \text{if } x_i = x_j \\ 0 & \text{else} \end{cases}$$

Toy example

- bids interval $[0, 1]$ discretized in $\{0, 0.5, 1\}$
- $v_1 = 0.8$ and $v_2 = 1$
- No negative payoff: $\mathcal{A}_1 = \{0, 0.5\}$ and $\mathcal{A}_2 = \{0, 0.5, 1\}$

Application: Two-players first-price sealed-bid auction
- Two bidders assign a value to a good and place a bid
- Higher bidder wins and pays their bid
- Possibly different values
- Discretization of continuous bids interval

Application: Two-players first-price sealed-bid auction

- Two bidders assign a value to a good and place a bid
- Higher bidder wins and pays their bid
- Possibly different values
- Discretization of continuous bids interval

$$u_i(x_i, x_j) = \begin{cases} v_i - x_i & \text{if } x_i > x_j \\ \frac{v_i - x_i}{2} & \text{if } x_i = x_j \\ 0 & \text{else} \end{cases}$$

Toy example

- bids interval $[0, 1]$ discretized in $\{0, 0.5, 1\}$
- $v_1 = 0.8$ and $v_2 = 1$
- No negative payoff: $\mathcal{A}_1 = \{0, 0.5\}$ and $\mathcal{A}_2 = \{0, 0.5, 1\}$

Application: Two-players first-price sealed-bid auction
- Two bidders assign a value to a good and place a bid
- Higher bidder wins and pays their bid
- Possibly different values
- Discretization of continuous bids interval
$$u_i(x_i, x_j) = \begin{cases} v_i - x_i & \text{if } x_i > x_j \\ \frac{v_i - x_i}{2} & \text{if } x_i = x_j \\ 0 & \text{else} \end{cases}$$
Toy example
- bids interval $[0, 1]$ discretized in $\{0, 0.5, 1\}$
- $v_1 = 0.8$ and $v_2 = 1$
- No negative payoff: $\mathcal{A}_1 = \{0, 0.5\}$ and $\mathcal{A}_2 = \{0, 0.5, 1\}$

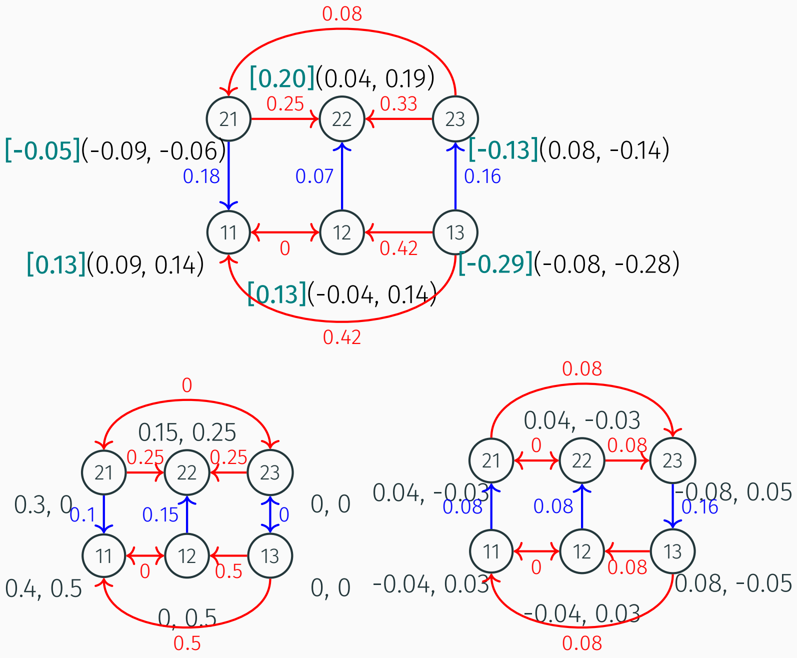

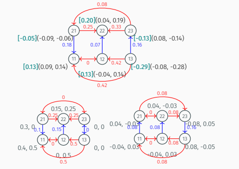

- full game bottom left; normalization not shown
- harmonic: net payoff flow at each node is zero
- strong correlation between harmonic and cycling
- harmonic always admits interior NE and never have pure NE
- space of harmonic and zero-sum has big non-trivial intersection (e.g. harmonic games where players have equal number of strategies are zero-sum games
- empirically, MD seems to cycle in harmonic games
- may be correct ingredient for non-convergence, more general than zero sum!

Decomposition of auctions and MD - Research questions

- (MD) empirically converges to BNE in many continuous auctions[2]
- Discretize and decompose
- **Potentialness** and convergence

$$p = \frac{\|u_{\mathcal{P}}\|}{\|u_{\mathcal{P}}\| + \|u_{\mathcal{H}}\|}$$

- **Convex combination** and convergence threshold

$$u(\alpha) = \alpha u_{\mathcal{P}} + (1-\alpha) u_{\mathcal{H}}$$

[2] Bichler, Fichtl, and Oberlechner [2]

Decomposition of auctions and MD - Research questions
- (MD) empirically converges to BNE in many continuous auctions[2]
- Discretize and decompose
- Potentialness and convergence
- Convex combination and convergence threshold
[2] Bichler, Fichtl, and Oberlechner [2]

- currently joint work with Bary Pradelski, Martin Bichler, Matthias Oberlechner, Panayotis Mertikopoulos
- **Potentialness** $p = \frac{\|u_{\mathcal{P}}\|}{\|u_{\mathcal{P}}\| + \|u_{\mathcal{H}}\|}$. How does high p correlate with convergence (necessary, sufficient)?
- **Perturbation** of the potential component building a new game as a convex combination of the potential and the harmonic components: $u(\alpha) = \alpha u_{\mathcal{P}} + (1-\alpha) u_{\mathcal{H}}$. Is there a threshold of harmonic perturbation at which convergence breaks?

Decomposition of auctions and MD - Research questions

- (MD) empirically converges to BNE in many continuous auctions[2]
- Discretize and decompose
- **Potentialness** and convergence

$$p = \frac{\|u_{\mathcal{P}}\|}{\|u_{\mathcal{P}}\| + \|u_{\mathcal{H}}\|}$$

- **Convex combination** and convergence threshold

$$u(\alpha) = \alpha u_{\mathcal{P}} + (1 - \alpha) u_{\mathcal{H}}$$

[2] Bichler, Fichtl, and Oberlechner [2]

Decomposition of auctions and MD - Research questions
- (MD) empirically converges to BNE in many continuous auctions[2]
- Discretize and decompose
- Potentialness and convergence
$p = \frac{\|u_{\mathcal{P}}\|}{\|u_{\mathcal{P}}\| + \|u_{\mathcal{H}}\|}$
- Convex combination and convergence threshold
$u(\alpha) = \alpha u_{\mathcal{P}} + (1 - \alpha) u_{\mathcal{H}}$
[2] Bichler, Fichtl, and Oberlechner [2]

- currently joint work with Bary Pradelski, Martin Bichler, Matthias Oberlechner, Panayotis Mertikopoulos
- **Potentialness** $p = \frac{\|u_{\mathcal{P}}\|}{\|u_{\mathcal{P}}\| + \|u_{\mathcal{H}}\|}$. How does high p correlate with convergence (necessary, sufficient)?
- **Perturbation** of the potential component building a new game as a convex combination of the potential and the harmonic components: $u(\alpha) = \alpha u_{\mathcal{P}} + (1 - \alpha) u_{\mathcal{H}}$. Is there a threshold of harmonic perturbation at which convergence breaks?

Decomposition of auctions and MD - Experiments[3]

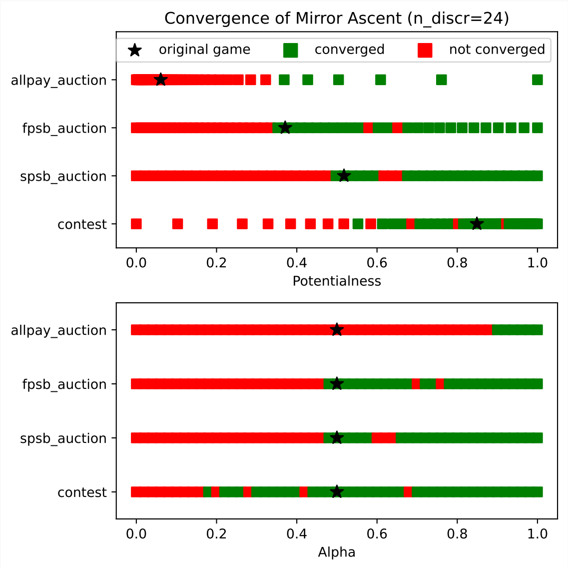

[3]Credits to Matthias Oberlechner for the great image!

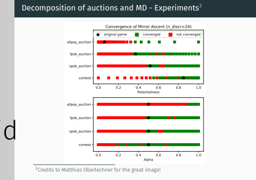

- first: star = origianl game; right and left games built as convex combination; plotted potentialness vs convergence for various types of auctions
- second: same plot changing coordinates (original game fixed in the middle; growing parameter alpha = higher potentialness
- windows of non-covergence at higher potentialness probably numerical (related to learning step)
- MD on the unperturbed game almost always converges, but high potentialness does not seem to be necessary for convergence but it is sufficient, as expected)
- open direction!

Mirror descent overview

Combinatorial decomposition for finite normal form games

Decomposition for general games

Mirror descent overview
Combinatorial decomposition for finite normal form games
Decomposition for general games

Limits of the combinatorial decomposition

- applies only to **finite normal form games**
- **inner product / regularizer**
- **normalization** choice

Smooth decomposition

- applies to **any game** (in the following, population games)
- **same regularizer** for decomposition and dynamics
- decompose directly **simultaneous gradient**, no normalization step

Limits of the combinatorial decomposition
- applies only to finite normal form games
- inner product / regularizer
- normalization choice
Smooth decomposition
- applies to any game (in the following, population games)
- same regularizer for decomposition and dynamics
- decompose directly simultaneous gradient, no normalization step

- decomposition relies on two choices: inner product and normalisation
- dynamics relies on choice of regularizer
- these choices are unrelated, so one can question whether there should be any relation between dynamics and decomposition
- **KEY IDEA reason to decompose sim. gradient** is that it naturally lives in a space where a Hodge decomposition theorem applies, for free! Technically speaking, 1-form. Intuitively, derivatives of payoff. In simple case of pop game, sim. gradient = payoff. In finite nfg as we saw, sim gradient is more complicated. For simpliciti focus on single pop game; wlog.

Limits of the combinatorial decomposition

- applies only to **finite normal form games**
- **inner product / regularizer**
- **normalization** choice

Smooth decomposition

- applies to **any game** (in the following, population games)
- **same regularizer** for decomposition and dynamics
- decompose directly **simultaneous gradient**, no normalization step

Limits of the combinatorial decomposition
· applies only to finite normal form games
· inner product / regularizer
· normalization choice
Smooth decomposition
· applies to any game (in the following, population games)
· same regularizer for decomposition and dynamics
· decompose directly simultaneous gradient, no normalization step

- decomposition relies on two choices: inner product and normalisation
- dynamics relies on choice of regularizer
- these choices are unrelated, so one can question whether there should be any relation between dynamics and decomposition
- **KEY IDEA reason to decompose sim. gradient** is that it naturally lives in a space where a Hodge decomposition theorem applies, for free! Technically speaking, 1-form. Intuitively, derivatives of payoff. In simple case of pop game, sim. gradient = payoff. In finite nfg as we saw, sim gradient is more complicated. For simpliciti focus on single pop game; wlog.

Single Population Game [20, 12]

- Continuum of agents (population)
- Set of **pure strategies** $\mathcal{A} = \{1, 2, \ldots, A\}$
- $\mathcal{X} = \Delta(A)$, distributions of pure strategies in the population
- State space $\mathcal{X}$, population state $x \in \mathcal{X}$
- x_a = fraction of population playing $a \in \mathcal{A}$

Single Population Game [20, 12]
- Continuum of agents (population)
- Set of pure strategies $\mathcal{A} = \{1, 2, \ldots, A\}$
- $\mathcal{X} = \Delta(A)$, distributions of pure strategies in the population
- State space $\mathcal{X}$, population state $x \in \mathcal{X}$
- x_a = fraction of population playing $a \in \mathcal{A}$

- Storyline: imagine pop. of identical agents meeting and playing a normal form game; this gives a distribution of strategies in the population, changing over time.

Single Population Game [20, 12]

- Payoff $u : \mathcal{X} \to \mathbb{R}^A$

$$u_a(x) = \text{payoff of } a\text{-strategist at state } x$$

- Expected payoff $\bar{u} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$

$$\bar{u}(y, x) = y \cdot u(x) = \text{expected payoff of } y\text{-strategist at state } x$$

- Simultaneous gradient $v : \mathcal{X} \to (\mathbb{R}^A)^*$

Gradient of exp. payoff $\bar{u}(y, x)$ w.r.t. mixed strategy y

$$v(x) := \frac{\partial \bar{u}(y, x)}{\partial y} = u(x)$$

Single Population Game [20, 12]
- Payoff $u : \mathcal{X} \to \mathbb{R}^A$
$u_a(x) =$ payoff of a-strategist at state x
- Expected payoff $\bar{u} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$
$\bar{u}(y, x) = y \cdot u(x) =$ expected payoff of y-strategist at state x
- Simultaneous gradient $v : \mathcal{X} \to (\mathbb{R}^A)^*$
Gradient of exp. payoff $\bar{u}(y, x)$ w.r.t. mixed strategy y
$v(x) := \frac{\partial \bar{u}(y, x)}{\partial y} = u(x)$

- Expected payoff $\bar{u} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$

$$\bar{u}(y, x) = y \cdot u(x) = \text{expected payoff of } y\text{-strategist at state } x$$

$$\bar{u}(y, x) = \mathbb{E}_{a \sim y}[u_a(x)]$$

for fixed state x, expectation value of the number $u_a(x)$ where a is drawn according to the distribution y

- sim. gradient is just payoff for single pop game, simpler to deal with. Wlog, can (probably) do with finite nfg, continuous strategy setting, ... wip

Hodge decomposition in a nutshell [6, 16]

- Choose regularizer $h : \mathcal{X} \to \mathbb{R}$ as in (MD)
- $g = \operatorname{Hess} h$

$$\delta v = \text{co-differential} = g^{ml}\left(\partial_l v_m - \Gamma^j_{ml} v_j\right)$$
$$df = \text{differential} = (\partial_1 f, \dots, \partial_A f)$$

Theorem (L., Mertikopoulos, Pradelski (2023))
Given a choice of regularizer, the simultaneous gradient of a population game with at lest 3 strategies admits a unique orthogonal decomposition

$$\boxed{v = df + \beta} \tag{H}$$

for some potential function f and some β such that $\delta\beta = 0$.

Hodge decomposition in a nutshell [6, 16]
- Choose regularizer $h : \mathcal{X} \to \mathbb{R}$ as in (MD)
- $g = \operatorname{Hess} h$
$\delta v = \text{co-differential} = g^{ml}\left(\partial_l v_m - \Gamma^j_{ml} v_j\right)$
$df = \text{differential} = (\partial_1 f, \dots, \partial_A f)$
Theorem (L., Mertikopoulos, Pradelski (2023))
Given a choice of regularizer, the simultaneous gradient of a population game with at lest 3 strategies admits a unique orthogonal decomposition
$v = df + \beta$ (H)
for some potential function f and some β such that $\delta\beta = 0$.

Under the carpet

- simply connected domain, can identify exact with closed and forget about harmonic
- codifferential is defined classically in terms of Hodge operator on forms
- compactness issue, technical proof based on isometry
- Strong convexity $\Rightarrow$ $\operatorname{Hess} h$ is bilinear, symmetric, positive-definite $\Rightarrow$ metric

Consequences of Theorem (H)

- Potential games are precisely those s.t. $v = df$, i.e. $\beta = 0$
- Games s.t. $v = \beta$ are called co-exact

Proposition (L., Mertikopoulos, Pradelski (wp))

(MD) *is volume-preserving on co-exact population games. In particular, there is no interior attractor. Again in particular, there is no interior ESS.*

Proof.

By standard divergence theorem, the flow of a vector field is volume-preserving iff the vector field is divergence-free. The result holds generalizing the divergence with the δ operator. The rest follows since ESSs are asymptotically stable under (MD) [12]. □

Consequences of Theorem (H)
- Potential games are precisely those s.t. $v = df$, i.e. $\beta = 0$
- Games s.t. $v = \beta$ are called co-exact
Proposition (L., Mertikopoulos, Pradelski (wp))
(MD) is volume-preserving on co-exact population games. In particular, there is no interior attractor. Again in particular, there is no interior ESS.
Proof.
By standard divergence theorem, the flow of a vector field is volume-preserving iff the vector field is divergence-free. The result holds generalizing the divergence with the δ operator. The rest follows since ESSs are asymptotically stable under (MD) [12]. □

- may be surprised by the term coexact and not harmonic. Actually, Hodge dec both in combinatorial and smooth setting has 3 components: exact, coexact, harmonic. Interestingly the coexact components vanishes in the combintorial setting and the harmonic one vanished in the smooth setting. The topological reason for this fact is clear and related to the nummber of holes of the space where the decomposition takes place. Game theorethically this is less clear, especially because the harmonic components and the coexact components seem to embody the same non-convergence nature, as opposed to the potential counterpart.

Proposition (L., Mertikopoulos, Pradelski (2023))

A population game with linear zero-sum payoff $v(x) = Ax$ is co-exact with respect to the entropic regularizer iff the bimatrix normal form game (A, A^T) is harmonic.

Corollary This implies the uniformly mixed strategy is a NE, so on this class of co-exact games (MD) dynamics is recurrent and Hamiltonian.

Co-exact game and harmonic games

Proposition (L., Mertikopoulos, Pradelski (2023))
A population game with linear zero-sum payoff $v(x) = Ax$ is co-exact with respect to the entropic regularizer iff the bimatrix normal form game (A, A^T) is harmonic.

Corollary This implies the uniformly mixed strategy is a NE, so on this class of co-exact games (MD) dynamics is recurrent and Hamiltonian.

Co-exact game and harmonic games

Proposition (L., Mertikopoulos, Pradelski (2023))
A population game with linear zero-sum payoff $v(x) = Ax$ is co-exact with respect to the entropic regularizer iff the bimatrix normal form game (A, A^T) is harmonic.

Corollary This implies the uniformly mixed strategy is a NE, so on this class of co-exact games (MD) dynamics is recurrent and Hamiltonian.

Co-exact game and harmonic games

Proposition (L., Mertikopoulos, Pradelski (2023))
A population game with linear zero-sum payoff $v(x) = Ax$ is co-exact with respect to the entropic regularizer iff the bimatrix normal form game (A, A^T) is harmonic.

Corollary This implies the uniformly mixed strategy is a NE, so on this class of co-exact games (MD) dynamics is recurrent and Hamiltonian.

Reminder - Cycles in zero-sum games under (MD)

Theorem (Mertikopoulos, Papadimitriou, and Piliouras [10])

Almost every solution trajectory $x(t)$ under (MD) is Poincaré recurrent on 2-player zero-sum games with an interior NE.

Theorem (L., Benedetti, Alishah, Mertikopoulos (wp))

(MD) dynamics on 2-player zero-sum games with an interior NE are Hamiltonian.

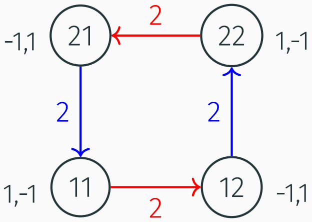

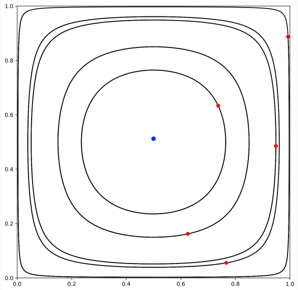

2023-10-13

Recurrence vs Convergence
└Decomposition for general games
└Reminder - Cycles in zero-sum games under (MD)

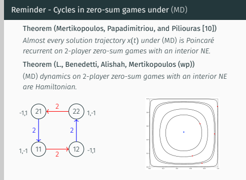

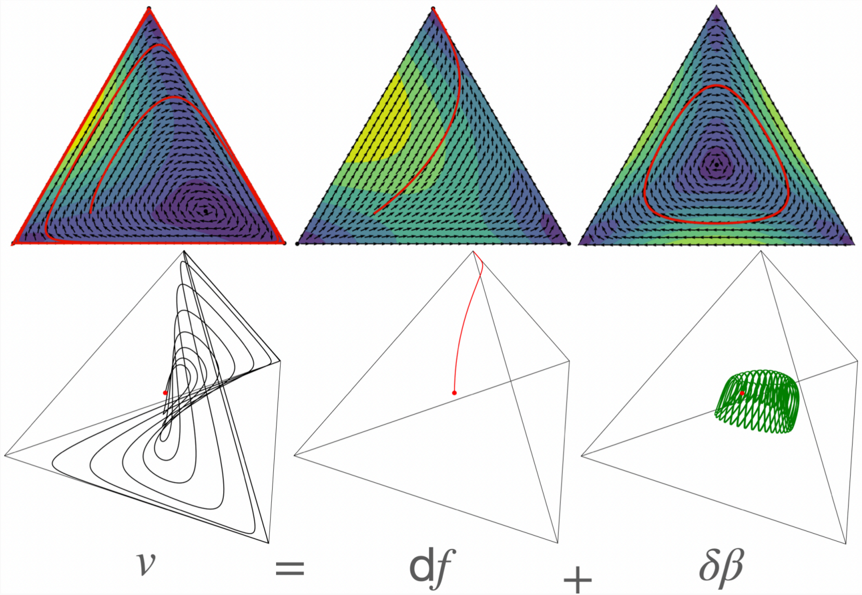

Recurrence vs Convergence
└Decomposition for general games

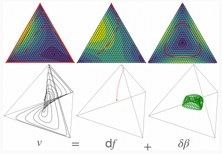

From here

- Characterize co-exact non zero-sum game. **Obstacle**: computation with Christoffel symbols, “derivation” problem
- Perform explicit decomposition $v = df + \beta$. **Obstacle**: solve Laplace equation, integration problem

As in normal form game case

- **Potentialness** and convergence (necessary, sufficient?)
- **Perturbation**. Convergence breaks at constant? *(animation)*

From here

- Characterize co-exact non zero-sum game. **Obstacle**: computation with Christoffel symbols, “derivation” problem
- Perform explicit decomposition $v = df + \beta$. **Obstacle**: solve Laplace equation, integration problem

As in normal form game case

- **Potentialness** and convergence (necessary, sufficient?)
- **Perturbation**. Convergence breaks at constant? (animation)

From here

- Characterize co-exact non zero-sum game. **Obstacle**: computation with Christoffel symbols, “derivation” problem
- Perform explicit decomposition $v = df + \beta$. **Obstacle**: solve Laplace equation, integration problem

As in normal form game case

- **Potentialness** and convergence (necessary, sufficient?)
- **Perturbation**. Convergence breaks at constant? *(animation)*

From here

- Characterize co-exact non zero-sum game. **Obstacle**: computation with Christoffel symbols, “derivation” problem
- Perform explicit decomposition $v = df + \beta$. **Obstacle**: solve Laplace equation, integration problem

As in normal form game case

- **Potentialness** and convergence (necessary, sufficient?)
- **Perturbation**. Convergence breaks at constant? *(animation)*

To recap - Mirror descent and Games Decomposition

- **Decomposition techniques** separate potential-converging component from cycling component $\implies$ characterize games “hard to learn” as co-exact
- **Regularizer** determines both the **learning** dynamics and the geometrical **decomposition** $\implies$ choose regularizer that minimizes co-exact component

Thanks!

To recap - Mirror descent and Games Decomposition

- **Decomposition techniques** separate potential-converging component from cycling component $\implies$ characterize games “hard to learn” as co-exact
- **Regularizer** determines both the **learning** dynamics and the geometrical **decomposition** $\implies$ choose regularizer that minimizes co-exact component

Thanks!

To recap - Mirror descent and Games Decomposition

- **Decomposition techniques** separate potential-converging component from cycling component $\implies$ characterize games "hard to learn" as co-exact
- **Regularizer** determines both the **learning** dynamics and the geometrical **decomposition** $\implies$ choose regularizer that minimizes co-exact component

Thanks!

To recap - Mirror descent and Games Decomposition

- **Decomposition techniques** separate potential-converging component from cycling component $\implies$ characterize games "hard to learn" as co-exact
- **Regularizer** determines both the **learning** dynamics and the geometrical **decomposition** $\implies$ choose regularizer that minimizes co-exact component

Thanks!

Bibliography - i

[1] Sanjeev Arora, Elad Hazan, and Satyen Kale. “The Multiplicative Weights Update Method: A Meta-Algorithm and Applications”. In: *Theory of Computing* 8.6 (May 2012), pp. 121–164. DOI: 10.4086/toc.2012.v008a006. (Visited on 09/19/2023).

[2] Martin Bichler, Maximilian Fichtl, and Matthias Oberlechner. “Computing Bayes Nash Equilibrium Strategies in Auction Games via Simultaneous Online Dual Averaging”. In: *Proceedings of the 24th ACM Conference on Economics and Computation*. EC ’23. New York, NY, USA: Association for Computing Machinery, July 2023, p. 294. ISBN: 9798400701047. DOI: 10.1145/3580507.3597713. (Visited on 08/29/2023).

[3] Ozan Candogan et al. “Flows and Decompositions of Games: Harmonic and Potential Games”. In: *Mathematics of Operations Research* 36.3 (Aug. 2011), pp. 474–503. DOI: 10.1287/moor.1110.0500.

[4] Georges de Rham. *Differentiable Manifolds: Forms, Currents, Harmonic Forms*. Springer Science & Business Media, 1984. ISBN: 978-3-642-61752-2.

[5] Xiaoye Jiang et al. “Statistical Ranking and Combinatorial Hodge Theory”. In: *Mathematical Programming* 127.1 (2011), pp. 203–244.

[6] Jürgen Jost. *Riemannian Geometry and Geometric Analysis*. Universitext. Cham: Springer International Publishing, 2017. ISBN: 978-3-319-61859-3 978-3-319-61860-9. DOI: 10.1007/978-3-319-61860-9. (Visited on 08/30/2023).

[7] John M. Lee. *Introduction to Smooth Manifolds*. 2nd ed. Graduate Texts in Mathematics. Springer-Verlag New York, 2012.

[8] Panayotis Mertikopoulos. “Online Optimization and Learning in Games: Theory and Applications”. 2019.

Bibliography - ii

[9] Panayotis Mertikopoulos and Joon Kwon. *A Continuous-Time Approach to Online Optimization*. 2014.

[10] Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. *Cycles in Adversarial Regularized Learning*. Sept. 2017. DOI: 10.48550/arXiv.1709.02738. arXiv: 1709.02738 [cs].

[11] Panayotis Mertikopoulos and William H Sandholm. "Learning in Games via Reinforcement and Regularization". In: *Mathematics of Operations Research* 41.4 (2016), pp. 1297–1324.

[12] Panayotis Mertikopoulos and William H Sandholm. "Riemannian Game Dynamics". In: *Journal of Economic Theory* 177 (2018), pp. 315–364.

[13] Panayotis Mertikopoulos and Zhengyuan Zhou. "Learning in Games with Continuous Action Sets and Unknown Payoff Functions". In: *Mathematical Programming* 173.1-2 (Jan. 2019), pp. 465–507. ISSN: 0025-5610, 1436-4646. DOI: 10.1007/s10107-018-1254-8. (Visited on 03/23/2023).

[14] Dov Monderer and Lloyd S. Shapley. "Potential Games". In: *Games and Economic Behavior* 14.1 (May 1996), pp. 124–143. ISSN: 0899-8256. DOI: 10.1006/game.1996.0044. (Visited on 02/19/2023).

[15] James R Munkres. *Elements of Algebraic Topology*. Perseus Books, 1984.

[16] Peter Petersen. *Riemannian Geometry*. Vol. 171. Graduate Texts in Mathematics. Cham: Springer International Publishing, 2016. ISBN: 978-3-319-26652-7 978-3-319-26654-1. DOI: 10.1007/978-3-319-26654-1. (Visited on 08/30/2023).

[17] Ralph Tyrell Rockafellar. *Convex Analysis*. Princeton university press, 1970.

[18] Steven Roman, S Axler, and FW Gehring. *Advanced Linear Algebra*. 3rd ed. Springer, 2008.

2023-10-13

Recurrence vs Convergence

└Decomposition for general games

└Bibliography -

Bibliography - iii

[19] Aldo Rustichini. “Optimal Properties of Stimulus—Response Learning Models”. In: *Games and Economic Behavior* 29.1 (Oct. 1999), pp. 244–273. ISSN: 0899-8256. DOI: 10.1006/game.1999.0712. (Visited on 09/01/2023).

[20] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT press, 2010.

[21] Siavash Shahshahani. *A New Mathematical Framework for the Study of Linkage and Selection*. American Mathematical Soc., 1979.

[22] Peter D Taylor and Leo B Jonker. “Evolutionary Stable Strategies and Game Dynamics”. In: *Mathematical biosciences* 40.1-2 (1978), pp. 145–156.

[23] Frank W Warner. *Foundations of Differentiable Manifolds and Lie Groups*. Vol. 94. Springer Science & Business Media, 1983.

Bibliography - iii

[19] Aldo Rustichini. “Optimal Properties of Stimulus—Response Learning Models”. In: *Games and Economic Behavior* 29.1 (Oct. 1999), pp. 244–273. ISSN: 0899-8256. DOI: 10.1006/game.1999.0712. (Visited on 09/01/2023).

[20] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT press, 2010.

[21] Siavash Shahshahani. *A New Mathematical Framework for the Study of Linkage and Selection*. American Mathematical Soc., 1979.

[22] Peter D Taylor and Leo B Jonker. “Evolutionary Stable Strategies and Game Dynamics”. In: *Mathematical biosciences* 40.1-2 (1978), pp. 145–156.

[23] Frank W Warner. *Foundations of Differentiable Manifolds and Lie Groups*. Vol. 94. Springer Science & Business Media, 1983.