Equilibria of Games with Algorithms

Jérôme Renault

Toulouse School of Economics

November 17, 2023

ANITI

Université de Toulouse

Game Theory and AI chair

Investigates possible applications of Game Theory concepts and methods to AI, as well as the consequences of the presence of AI in human interactions.

Thematics include:

1. interactions between algorithms,
2. hybrid games (between algorithms and rational agents),
3. bandits and online learning,
4. decisions and use of information in complex strategic environments,
5. computation of equilibria in min-max problems,
6. zero-order global Lipschitz optimization,
7. approximation and certification

($\sim$ 47 (pre-)publications, including Maths of Operations Research, Maths Programming, Operations Research, Theoretical Economics, NeurIPS, ALT, EC, a reinforcement learning virtual school with 1500 participants, a book)

 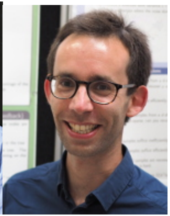    

F. Gensbittel S.Gerchinovitz T. Cesari D. Pizarro M. d'Andrea E. De Montbrun

1. Games played by Exponential Weights Algorithms

(M. d'Andrea, F. Gensbittel, J. Renault)

Machine learning algorithms for decision-making and prediction are mostly designed to optimize the behavior of an agent facing an unknown environment.

As they are more and more widely used, they will widely interact in the future.

$\rightarrow$ Will these interactions lead somewhere ?

Focus here on Exponential Weights Algorithms, used in many domains and contexts.

Exponential weights algorithm

Consider $u : A \times K \to \mathbb{R}$ giving the payoff of an agent when the environment is in state k. At stage $t = 0, 1, 2, ...$, select an action a_t, then observe the environment k_t and get the payoff $u(a_t, k_t)$.

EW algorithm with fixed parameter $\eta > 0$:

- positive weights $w_t(a) > 0$ for all a and t,
- at each stage play proportionally to the current weights,
- weights are initialized (typically $w_0(a) = 1$ for all a) and updated by:

$$w_{t+1}(a) = w_t(a) \exp(\eta u(a, k_t)).$$

$\rightarrow$ No-regret: for all T, k_0,...,k_{T-1},

$$\mathbb{E}\left(\frac{1}{T}\sum_{t=0}^{T-1} u(a_t, k_t)\right) \geq \max_a \left(\frac{1}{T}\sum_{t=0}^{T-1} u(a, k_t)\right) - \underline{\frac{\ln |A|}{\eta T} - \frac{1}{2}\eta \|u\|^2}.$$

Model: Several EW algorithms interact

Base game: each player $i = 1, ..., N$ has a finite action set A_i and receives payoffs according to $u_i : A_1 \times ... \times A_N \to \mathbb{R}$.

Behaviors: each player i uses a EW algo with parameter $\eta_i > 0$ and some initial weights.

Dynamics: at each stage $t = 0, 1, ..,$ each player i plays an action a_i^t with proba proportional to her current weights, then observes a_{-i}^t and updates her weights by:

$$\forall a_i,\ w_i^{t+1}(a_i) = w_i^t(a_i) \exp(\eta_i u_i(a_i, a_{-i}^t)).$$

Define the random variable of normalized weights $p_i^t = (\frac{w_i^t(a_i)}{w_i^t})_{a_i \in A_i} \in \Delta(A_i)$.

$(p^t) = ((p_i^t)_{i \in N})$ is a Markov chain. What are its properties ?

Example

$$\begin{array}{ccc} & L & R \\ T & \multirow{2}{*}{} & \end{array}$$

$$\begin{array}{c} T \\ B \end{array} \begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \left(\begin{array}{cc} (1,1) & (0,0) \\ (0,0) & (1,1) \end{array} \right) \end{array}$$

Proposition: Almost surely, (T,L) is played at all but finitely many stages, or (B,R) is played at all but finitely many stages.

1. not true for all no-regret strategies.
2. generalizes to “strong coordination games”

Weakly Stable Nash Equilibria

Define a Weakly Stable Nash Equilibrium as a NE p such that for each player i, for all a_i, b_i in the support of p_i, for all a_{-i} in the support of p_{-i}: $u_i(a_i, a_{-i}) = u_i(b_i, a_{-i})$.

$$\begin{array}{cc} & \begin{array}{ccc} L & M & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} (3,3) & (0,0) & (2,2) \\ (0,0) & (3,3) & (2,2) \end{pmatrix} \end{array}$$

2 strict NE: (T, L) and (B, M) and a continuum of other WSNE: $(xT + (1-x)B, R)$ with $1/3 \leq x \leq 2/3$.

Theorem 1: Assume (p^t) converges with positive probability to a random variable p^*. Then a.s. p^* is a WSNE.

Corollary: no possible convergence in games without WSNE.

Limit probability of playing a SNE

Define L^t as the probability in $[0,1]$ to play a strict NE at period $t+1$ given the present history at period t.

Theorem 2: $(L^t)_t$ converges to a random variable with values in $\{0,1\}$.

Illustrations:

$$\begin{array}{c c} & \begin{array}{ccc} L & M & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} (3,3) & (0,0) & (2,2) \\ (0,0) & (3,3) & (2,2) \end{pmatrix} \end{array}$$

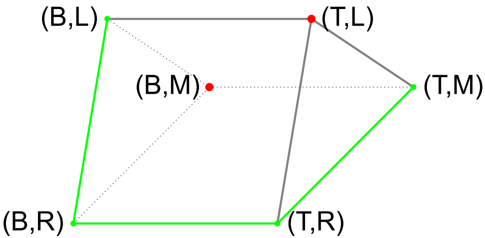

$$\begin{array}{c c} & \begin{array}{ccc} B_1 & B_2 & B_3 \end{array} \\ \begin{array}{c} A_1 \\ A_2 \\ A_3 \end{array} & \begin{pmatrix} (0,0) & (-1,-1) & (-1,-1) \\ (-1,-1) & (1,1) & (1,0) \\ (-1,-1) & (1,0) & (1,1) \end{pmatrix} \end{array}$$

With some proba in $(0,1)$, the play CV to the strict NE (A_1, B_1). Otherwise, it will eventually remain in $\{A_2, A_3\} \times \{B_2, B_3\}$ and diverges.

Limit probability of playing a SNE

Define L^t as the probability in $[0,1]$ to play a strict NE at period $t+1$ given the present history at period t.

Theorem 2: $(L^t)_t$ converges to a random variable with values in $\{0,1\}$.

Illustrations:

$$\begin{array}{c c} & \begin{array}{ccc} L & M & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} (3,3)^{\bullet} & (0,0) & (2,2) \\ (0,0) & (3,3)^{\bullet} & (2,2) \end{pmatrix} \end{array}$$

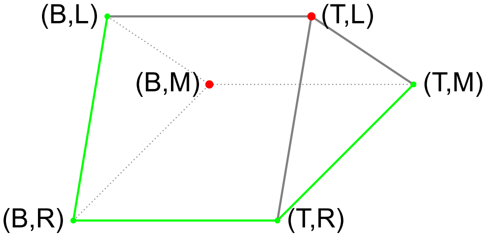

$$\begin{array}{c c} & \begin{array}{ccc} B_1 & B_2 & B_3 \end{array} \\ \begin{array}{c} A_1 \\ A_2 \\ A_3 \end{array} & \begin{pmatrix} (0,0)^{\bullet} & (-1,-1) & (-1,-1) \\ (-1,-1) & (1,1) & (1,0) \\ (-1,-1) & (1,0) & (1,1) \end{pmatrix} \end{array}$$

With some proba in $(0,1)$, the play CV to the strict NE (A_1, B_1). Otherwise, it will eventually remain in $\{A_2, A_3\} \times \{B_2, B_3\}$ and diverges.

2. Optimistic Gradient DescentAscent in Bilinear Games

(E. de Montbrun, J. Renault)

We consider Minmax Optimization:

$$\min_{y \in Y} \sup_{x \in X} g(x, y),$$

for $g : X \times Y \to \mathbb{R}$. Many possibilities

Game-theoretic approach:
Consider the zero-sum game where simultaneously player 1 chooses x in X, player 2 chooses y in Y, and finally player 2 pays $g(x, y)$ to player 1.

Look for a Nash equilibrium (or saddle-point), i.e. for (x^*, y^*) s.t.:

$$\forall x \in X, \forall y \in Y, \; g(x, y^*) \leq g(x^*, y^*) \leq g(x^*, y)$$

$\rightarrow$ Then y^* achieves $\min_{y \in Y} \sup_{x \in X} g(x, y)$.

Example: Generative Adversarial Networks

Goal: simulate a distribution from high-dimensional data: sample elements that mimic the observations (cats, old portraits, people faces,...)(Goodfellow *et al.* 2014)
Can be seen as a Min-Max Optimization problem or a game between a Generator and a Discriminator.

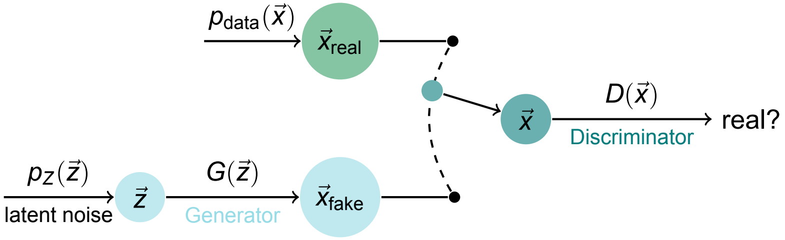

$$\min_{\theta_g} \max_{\theta_d} \left(\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p_Z} \log(1 - D_{\theta_d}(G_{\theta_g}(z))) \right)$$

with $D_{\theta_d}(x) \in [0, 1]$ probability assigned by D that x is true.

Gradient Descent-Ascents

Assuming g is "nice" and a saddle-point exists, how to find a saddle-point of g ?

- Difficult problem ("Non-Concave Games: a challenge for Game Theory's next 100 years", Daskalakis 2022).
- Natural to look at Gradient Descent-Ascents: given $\eta > 0$, iterate:

$$\begin{cases} x_{t+1} = x_t + \eta \frac{\partial g}{\partial x}(x_t, y_t) \\ y_{t+1} = y_t - \eta \frac{\partial g}{\partial y}(x_t, y_t) \end{cases}$$

Does not always converge.

- $\rightarrow$ We consider Optimistic Gradient Descent-Ascents

$$\begin{cases} x_{t+1} = x_t + 2\eta \frac{\partial g}{\partial x}(x_t, y_t) - \eta \frac{\partial g}{\partial x}(x_{t-1}, y_{t-1}) \\ y_{t+1} = y_t - 2\eta \frac{\partial g}{\partial y}(x_t, y_t) + \eta \frac{\partial g}{\partial y}(x_{t-1}, y_{t-1}) \end{cases}$$

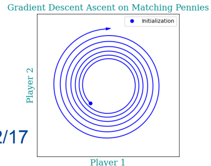

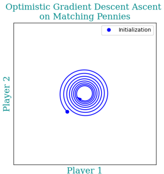

For the presentation: simple unconstrained bilinear case

$$g(x, y) = x^T A y,$$

for some $A \in \mathbb{R}^{n \times p}$, with $x \in \mathbb{R}^n$, $y \in \mathbb{R}^p$.

OGDA for zero-sum games

Theorem 1: Write $\mu_{\max} = \max\{\mu, \mu \in \mathrm{Sp}(AA^T)\}$ and $\mu_{\min} = \min\{\mu > 0, \mu \in \mathrm{Sp}(AA^T)\}$.

1) If $0 < \eta < \frac{1}{\sqrt{3\mu_{\max}}}$, then $(x_t, y_t)_t$ CV to the Nash equilibrium (x_∞, y_∞), where

$$x_\infty = \mathrm{proj}^{\perp}_{\mathrm{Ker}(A^T)}(x_0) \text{ and } y_\infty = \mathrm{proj}^{\perp}_{\mathrm{Ker}(A)}(y_0)$$

The CV is exponential, and we compute the exact rate $\lambda(\eta)$.

2) We compute the optimal η, and the corresponding rate λ^* in $[\frac{\sqrt{2}}{2}, 1)$ as a function of $\mu_{\max}$ and $\mu_{\min}$.

(Improvement on the litterature)

OGDA for general-sum games

We introduce OGDA for the general-sum game $(x^T Ay, x^T By)$:

$$\begin{cases} x_{t+1} &= x_t + 2\eta A y_t - \eta A y_{t-1} \\ y_{t+1} &= y_t + 2\eta B x_t - \eta B x_{t-1} \end{cases}$$

Theorem 2:

- CV of OGDA may fail in the general-sum context
- "Nice" sufficient conditions for convergence
- Coordination in a "large" class of games: either $(x_t, y_t)_t$ CV to a Nash equilibrium, or both payoffs $x_t^T A y_t$ and $x_t^T B y_t$ CV to $+\infty$

Theorem: Let A, B be in $\mathbb{R}^{n\times p}$ with $\mathrm{Sp}(B^T A) \subset \mathbb{R}$. Let $\eta \in (0, \frac{1}{2\sqrt{\mu_{\max}}})$, with $\mu_{\max}$ the largest eigenvalue of $B^T A$. Assume $B^T A$ is diagonalizable, $\mathrm{Ker}(A) \oplus \mathrm{Im}(B^T) = \mathbb{C}^p$ and $\mathrm{Ker}(B^T) \oplus \mathrm{Im}(A) = \mathbb{C}^n$. Then either $(x_t, y_t)_t$ CV to a Nash equilibrium, or both payoffs $x_t^T A y_t$ and $x_t^T B y_t$ CV to $+\infty$.

Back to the MinSup problem: $\min_y \sup_x x^T A y$

Option 1: Run OGDA for the zero-sum game for some step-size η, with a CV rate of $\lambda(\eta)$.

Option 2: Optimize the step-size η to get a rate $\lambda^\star = \sqrt{\frac{1}{2}\left(1+\sqrt{1-\frac{\beta}{8}}\right)} \in [\frac{\sqrt{2}}{2}, 1)$, better when $\frac{\min\{\mu>0, \mu\in \mathrm{Sp}(A^T A)\}}{\max\{\mu, \mu\in \mathrm{Sp}(A^T A)\}}$ is large.

Option 3: Run OGDA for the general-sum game $(x^T A y, x^T B y)$ with $B = -(A^+)^T$, and $\eta = \frac{1}{2}$. Theorem 3: Then $(y_t)_t$ converges to the same limit, with exponential rate arbitrarily close to the optimal value $\frac{\sqrt{2}}{2}$.

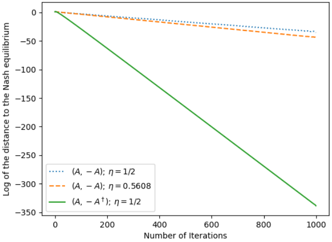

Going further: extension to polynomial games

Assume $g(x,y) = \psi(x)^T A\phi(y)$, with ψ, ϕ smooth. We can run the general-sum extension of OGDA on the game with payoffs $g_1 = g$ and $g_2(x,y) = -\psi(x)^T (A^+)^T \phi(y)$.

Example : $g(x,y) = -2x^2 + \frac{1}{3}y^2 + xy$, with x, $y \in \mathbb{R}$. Here $\psi(x) = (-x^2, x, 1)^T$, $\phi(y) = (y^2, y, 1)^T$ and $A = \begin{pmatrix} 0 & 0 & 2 \\ 0 & 1 & 0 \\ \frac{1}{3} & 0 & 0 \end{pmatrix}$. Then $g_2(x,y) = 0.5x^2 - 3y^2 - xy$. $NE = \{(0,0)\}$ in both games.

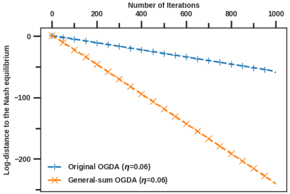

→ Running OGDA on an appropriate general-sum game improves the CV to equilibria of a MinSup problem.