Generalized conditional gradient method for potential mean field games

Laurent Pfeiffer
Inria and CentraleSupélec, Université Paris-Saclay

Joint work with Pierre Lavigne (Institut Louis Bachelier)

Séminaire français d'optimisation
February 3, 2022

Introduction

1. Description of the **generalized conditional gradient** (GCG) algorithm (an extension of the Frank-Wolfe algorithm).

   Linear convergence for adaptive stepsize rules (in a simple setting).

2. For a simple class of potential games, the GCG algorithm is a **best-response procedure**.

3. Heuristic derivation of the **mean-field game** (MFG) of interest.

4. Application of the GCG method to MFGs, **interpretation and convergence results**.

General setting

Consider the following problem:

$$\inf_{x \in \mathbb{R}^n} f(x) := f_1(x) + f_2(x), \quad \text{subject to: } x \in K. \qquad (\mathcal{P})$$

Assumptions:

- $K \subseteq \mathbb{R}^n$ is convex
- $f_1 \colon K \to \mathbb{R}$ and $f_2 \colon K \to \mathbb{R}$ are convex
- f_1 is lower semi-continuous
- f_2 has a Lipschitz-continuous gradient
- K is non-empty and compact.

Let $\bar{x}$ denote a solution to the problem.

Subproblem

- Given $x \in \mathbb{R}^n$, we denote by $f_{\text{lin}}[x] \colon \mathbb{R}^n \to \mathbb{R}$ the (partial) **linearization** of f at x, defined by:

$$f_{\text{lin}}[x](y) = f_1(y) + f_2(x) + \langle \nabla f_2(x), y - x \rangle.$$

  Since f_2 is convex, $f_{\text{lin}}[x](y) \leq f(y)$ for all $x \in \mathbb{R}^n$.

- We consider the **linearized problem** at x, defined by

$$\inf_{y \in \mathbb{R}^n} f_{\text{lin}}[x](y), \quad \text{subject to: } y \in K. \qquad (\mathcal{P}_{\text{lin}}(x))$$

  Assumption: $\mathcal{P}_{\text{lin}}(x)$ is **numerically easy to solve**, for any $x \in K$.

- We call **primal-dual gap** the number $\sigma(x)$ defined by

$$0 \leq \sigma(x) = f_{\text{lin}}[x](x) - \Big(\inf_{y \in K} f_{\text{lin}}[x](y) \Big).$$

Subproblem

- Given $x \in \mathbb{R}^n$, we denote by $f_{\text{lin}}[x] \colon \mathbb{R}^n \to \mathbb{R}$ the (partial) **linearization** of f at x, defined by:

$$f_{\text{lin}}[x](y) = f_1(y) + f_2(x) + \langle \nabla f_2(x), y - x \rangle.$$

  Since f_2 is convex, $f_{\text{lin}}[x](y) \leq f(y)$ for all $x \in \mathbb{R}^n$.

- We consider the **linearized problem** at x, defined by

$$\inf_{y \in \mathbb{R}^n} f_{\text{lin}}[x](y), \quad \text{subject to: } y \in K. \tag{$\mathcal{P}_{\text{lin}}(x)$}$$

  Assumption: $\mathcal{P}_{\text{lin}}(x)$ is **numerically easy to solve**, for any $x \in K$.

- We call **primal-dual gap** the number $\sigma(x)$ defined by

$$0 \leq \sigma(x) = f_{\text{lin}}[x](x) - \left(\inf_{y \in K} f_{\text{lin}}[x](y) \right).$$

Subproblem

- Given $x \in \mathbb{R}^n$, we denote by $f_{\text{lin}}[x] \colon \mathbb{R}^n \to \mathbb{R}$ the (partial) **linearization** of f at x, defined by:

$$f_{\text{lin}}[x](y) = f_1(y) + f_2(x) + \langle \nabla f_2(x), y - x \rangle.$$

  Since f_2 is convex, $f_{\text{lin}}[x](y) \leq f(y)$ for all $x \in \mathbb{R}^n$.

- We consider the **linearized problem** at x, defined by

$$\inf_{y \in \mathbb{R}^n} f_{\text{lin}}[x](y), \quad \text{subject to: } y \in K. \qquad (\mathcal{P}_{\text{lin}}(x))$$

  Assumption: $\mathcal{P}_{\text{lin}}(x)$ is **numerically easy to solve**, for any $x \in K$.

- We call **primal-dual gap** the number $\sigma(x)$ defined by

$$0 \leq \sigma(x) = f_{\text{lin}}[x](x) - \left(\inf_{y \in K} f_{\text{lin}}[x](y) \right).$$

Optimality certificate

Lemma

Let $x \in K$. Then, x is $\sigma(x)$-optimal, that is to say,

$$f(x) \leq f(\bar{x}) + \sigma(x).$$

Proof. By definition, we have

$$-\sigma(x) = -f_{\text{lin}}[x](x) + \inf_{y \in K} f_{\text{lin}}[x](y) \leq -f(x) + f_{\text{lin}}[x](\bar{x}).$$

Finally, we have $f_{\text{lin}}[x](\bar{x}) \leq f(\bar{x})$. Therefore,

$$-\sigma(x) \leq -f(x) + f(\bar{x}).$$

Remark. The condition $\sigma(x) = 0$ is also a **necessary** condition of optimality.

Algorithm

Algorithm 1: Generalized conditional gradient algorithm

Input: $\bar{x}_0 \in K$;

for $k = 0, 1, \ldots$ **do**

Find a solution x_k to $\mathcal{P}_{\mathrm{lin}}(\bar{x}_k)$;

Choose a stepsize $\delta_k \in [0, 1]$;

Set $\bar{x}_{k+1} = (1 - \delta_k)\bar{x}_k + \delta_k x_k$;

end

Theorem

There exists a constant $C > 0$ such that the following holds true.

- If $\delta_k = \dfrac{1}{k+1}$, then $f(\bar{x}_k) \leq f(\bar{x}) + \dfrac{C \ln(k)}{k}$, $\forall k > 1$.
- If $\delta_k = \dfrac{2}{k+2}$, then $f(\bar{x}_k) \leq f(\bar{x}) + \dfrac{C}{k}$, $\forall k > 0$.

References

When $f_1 = 0$, the GCG algorithm coincides with the well-known **Frank-Wolfe** algorithm.

Frank, Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 1956.

Earliest reference about the GCG algorithm:

Bredies, Lorenz, Maass. A generalized conditional gradient method and its connection to an iterative shrinkage method. *Computational Optim. and App.*, 2009.

The **linear convergence** rate exhibited next is adapted from

Kunisch, Walter. On fast convergence rates for generalized conditional gradient methods with backtracking stepsize, *ArXiv preprint*, 2021.

An adaptative stepsize rule

Theorem

Assume the following:

- The set K is non-empty, convex, and closed.
- The function f_1 is l.s.c. and α-**strongly** convex over K.
- The function f_2 is convex with an L-Lipschitz gradient.

Consider the **adaptative** stepsize rule

$$\delta_k = \min\left(\frac{\sigma_k}{LD_k^2}, 1\right),$$

where $\sigma_k = \sigma(\bar{x}_k)$ and $D_k = \|x_k - \bar{x}_k\|$.

Then there exists $\lambda \in [0, 1)$ such that

$$f(\bar{x}_k) - f(\bar{x}) \leq \lambda^k \big(f(\bar{x}_0) - f(\bar{x})\big), \quad \forall k \in \mathbb{N}.$$

Motivation

Let $k \in \mathbb{N}$. For $\delta \in [0,1]$, we set $x_\delta = (1-\delta)\bar{x}_k + \delta x_k$. We have the following upper bound:

$$\begin{aligned} f(x_\delta) &= f_1(x_\delta) + f_2(x_\delta) \\ &\leq \big[(1-\delta) f_1(\bar{x}_k) + \delta f_1(x_k)\big] \\ &\qquad + \left[f_2(\bar{x}_k) + \delta \langle \nabla f_2(\bar{x}_k), x_k - \bar{x}_k \rangle + \frac{L\delta^2}{2} D_k^2 \right]. \end{aligned}$$

Re-arraging:

$$f(x_\delta) - f(\bar{x}_k) \leq h(\delta) := -\delta \sigma_k + \frac{L\delta^2}{2} D_k^2.$$

The chosen stepsize $\delta_k = \min\left(\frac{\sigma_k}{LD_k^2}, 1\right)$, minimizes h over $[0,1]$.

Proof

Proof of the theorem.

Step 1. A bound of D_k.

- By construction x_k minimizes $f_{\text{lin}}[\bar{x}_k](\cdot)$ over K.
- The point $\bar{x}_k$ is σ_k-optimal for this minimization problem.
- Moreover, $f_{\text{lin}}[\bar{x}_k](\cdot)$ is α-strongly convex (since f_1 is α-strongly convex).

Therefore,

$$D_k^2 = \|x_k - \bar{x}_k\|^2 \leq \frac{2\sigma_k}{\alpha}.$$

Remark. The strong convexity of f_1 is only used at this step of the proof.

Proof

Step 2. A bound of $h(\delta_k)$.

- Case 1: $\sigma_k \geq LD_k^2$. Then $\delta_k = 1$ and

$$h(\delta_k) = -\sigma_k + \frac{L}{2} D_k^2 \leq -\frac{1}{2}\sigma_k.$$

- Case 2: $\sigma_k < LD_k^2$. Then $\delta_k = \frac{\sigma_k}{LD_k^2}$ and

$$h(\delta_k) = -\frac{\sigma_k^2}{2LD_k^2} \leq -\frac{\sigma_k \alpha}{4L}.$$

Therefore, $h(\delta_k) \leq -\omega\sigma_k$, where $\omega = \min\left(\frac{1}{2}, \frac{\alpha}{4L}\right) > 0$.

Proof

Step 3. Conclusion.
We deduce that

$$\begin{aligned} f(\bar{x}_{k+1}) - f(\bar{x}) &= f(x_{\delta_k}) - f(\bar{x}) \\ &\leq (f(\bar{x}_k) - f(\bar{x})) + h(\delta_k) \\ &\leq (f(\bar{x}_k) - f(\bar{x})) - \omega\sigma_k \\ &\leq (f(\bar{x}_k) - f(\bar{x})) - \omega(f(\bar{x}_k) - f(\bar{x})) \\ &= (1-\omega)(f(\bar{x}_k) - f(\bar{x})). \end{aligned}$$

Variants

Some other adaptative stepsize rules can be considered.

- **Exact minimization:**

$$\delta_k \in \underset{\delta \in [0,1]}{\operatorname{argmin}} \; f(x_\delta).$$

- **Armijo-Goldstein rule:** given $\gamma \in (0,1)$ and $\eta \in (0,1)$,

$$\delta_k \in \operatorname{argmax}\Big\{ \delta \;\Big|\; f(x_\delta) \leq f(\bar{x}_k) - \eta \delta \sigma_k, \; \delta = \gamma^j, \; j = 0, 1, \ldots \Big\}$$

The result of the theorem remains true for these choices of rules. They do not require the knowledge of L.

Model

- Let $(Y, \mathcal{Y}, \mu)$ be a probability space. We consider a continuum of agents, characterized by a **parameter** $y \in Y$, with probability distribution μ.

- The game involves two variables:
  - the **decisions** $X \in L^\infty(Y; \mathbb{R}^d)$; $X(y)$ is the decision of the agents with parameter y
  - a **price** variable $\lambda \in \mathbb{R}^d$, common to all agents.

$$\begin{tikzcd} \boxed{\begin{array}{c}\text{Decision } X(\cdot)\\ \text{of the agents}\end{array}} \arrow[r, shift left, "\text{Coupling function}"'] & \boxed{\begin{array}{c}\text{Price } \lambda\\ \text{(common to all)}\end{array}} \arrow[l, shift left, "\text{Optim. problem}"'] \end{tikzcd}$$

- **Non-atomic** agents: they do not take into account their own impact on λ in the optimization problem.

Model

- The decision variables X satisfy

$$X(y) \in \underset{x \in X_{\text{ad}}}{\operatorname{argmin}}\; f_{\lambda,y}(x) := \ell(x,y) + \langle \lambda, x \rangle, \qquad (\mathcal{P}_{\lambda,y})$$

  where $X_{\text{ad}} \subseteq \mathbb{R}^d$ and $\ell \colon X_{\text{ad}} \times Y \to \mathbb{R}$.

- The price $\lambda \in \mathbb{R}^d$ is deduced from $X \in L^\infty(Y, \mathbb{R}^d)$ through

$$\lambda = \psi\Big(\int_Y X(y) \, \mathrm{d}\mu(y) \Big),$$

  where the price function $\psi \colon \mathbb{R}^d \to \mathbb{R}^d$ is given.

Interpretation: Cournot equilibrium, $\int_Y X(y) \, \mathrm{d}\mu(y)$ is a demand of some product.

Example

Theorem

Assume the following:

- The set X_{ad} is convex and closed.
- The cost $\ell(\cdot, y)$ is α-strongly convex over X_{ad}, for any $y \in Y$.
- There exists $C > 0$ and $x_0 \in X_{\text{ad}}$ such that for all $x \in X_{ad}$ and for all $y \in Y$,

$$\ell(x, y) \geq \frac{1}{C}\|x\|^2 - C \quad \text{and} \quad \ell(x_0, y) \leq C.$$

- The function ψ is Lipschitz-continuous and bounded over X_{ad}.

Then **there exists** a pair $(X, \lambda) \in L^\infty(Y; \mathbb{R}^d) \times \mathbb{R}^d$ which is **solution** to the game.

Proof

Proof.

- It is easy to verify that for all $\lambda \in \mathbb{R}^d$, for all $y \in Y$, the problem $\mathcal{P}_{\lambda,y}$ has a unique solution $X_\lambda(y)$.
- Moreover, $X_\lambda(\cdot) \in L^\infty(Y;\mathbb{R}^d)$ and the mapping

$$\lambda \in \mathbb{R}^d \mapsto X_\lambda \in L^\infty(Y;\mathbb{R}^d),$$

  called **best-response** function, is Lipschitz-continuous.
- The game boils down to the fixpoint relation

$$\lambda = \theta(\lambda) := \psi\Big(\int_Y X_\lambda(y) \, \mathrm{d}\mu(y) \Big).$$

- Let $C > 0$ denote a bound of $\|\psi\|$ over X_{ad}. The mapping θ is continuous from $\bar{B}_{\mathbb{R}^d}(C)$ to $\bar{B}_{\mathbb{R}^d}(C)$. Thus by the **Schauder fixpoint theorem**, there exists λ such that $\lambda = \theta(\lambda)$.

Potential formulation

Theorem

Consider the assumptions of the previous theorem. Assume moreover that there exists a convex function $\phi \colon \mathbb{R}^d \to \mathbb{R}$ such that

$$\psi = \nabla \phi.$$

Then there exists a **unique** solution $(\bar{X}, \bar{\lambda})$ to the game. Moreover, $\bar{X}$ is the unique solution to the following **potential problem**:

$$\min_{X \in L^\infty(Y; X_{\mathrm{ad}})} F(X) := \int_Y \ell(X(y), y) \,\mathrm{d}\mu(y) + \phi\Big(\int_Y X(y) \,\mathrm{d}\mu(y) \Big).$$

Proof

Let $(\bar{X}, \bar{\lambda})$ be a solution. Then, for any $X \in L^\infty(Y; X_{\text{ad}})$,

$$
\begin{aligned}
F(X) - F(\bar{X}) &= \int_Y \left(\ell(X(y), y) - \ell(\bar{X}(y), y)\right) \mathrm{d}\mu(x) \\
&\qquad + \phi\Big(\int_Y X(y)\, \mathrm{d}\mu(y)\Big) - \phi\Big(\int_Y \bar{X}(y)\, \mathrm{d}\mu(y)\Big) \\
&\geq \int_Y \left(\ell(X(y), y) - \ell(\bar{X}(y), y)\right) \mathrm{d}\mu(x) \\
&\qquad + \Big\langle \underbrace{\nabla\phi\big(\textstyle\int_Y \bar{X}(y)\, \mathrm{d}\mu(y)\big)}_{=\bar{\lambda}}, \displaystyle\int_Y X(y)\, \mathrm{d}\mu(y) - \int_Y \bar{X}(y)\, \mathrm{d}\mu(y) \Big\rangle \\
&= \int_Y \Big[\big(\ell(X(y), y) + \langle \bar{\lambda}, X(y)\rangle\big) - \big(\ell(\bar{X}(y), y) + \langle \bar{\lambda}, \bar{X}(y)\rangle\big)\Big] \mathrm{d}\mu(y) \\
&= \int_Y \Big[f_{\bar{\lambda}, y}(X(y)) - f_{\bar{\lambda}, y}(\bar{X}(y))\Big] \mathrm{d}\mu(y) \geq 0.
\end{aligned}
$$

Proof

Conclusion.

- It follows that $\bar{X}$ minimizes F.
- It is easy to verify that F is α-strongly convex (in $L^2_\mu(Y;\mathbb{R}^d)$), thus F has a unique minimizer.
- As a consequence, the solution to the game is unique.

Application of GCG

The proof of the potential formulation reveals a **suitable decomposition** $F = F_1 + F_2$ for the application of the generalized conditional gradient method! We define

$$F_1(X) = \int_Y \ell(X(y), y)\, \mathrm{d}\mu(x) \quad \text{and} \quad F_2(x) = \phi\Big(\int_Y X(y)\, \mathrm{d}\mu(y) \Big).$$

Let $\bar{X}_k$ and $X \in L^\infty(Y; \mathbb{R}^d)$. Let $\lambda_k = \nabla\phi\big(\int_Y \bar{X}_k(y)\, \mathrm{d}\mu(y) \big)$. We have

$$F_{\text{lin}}[\bar{X}_k](X) = \int_Y \Big(\underbrace{\ell(X(y), y) + \langle \lambda_k, X(y) \rangle}_{f_{\lambda_k, y}(X(y))} \Big)\, \mathrm{d}\mu(y) + \text{Constant}.$$

The unique minimizer of $F_{\text{lin}}[\bar{X}_k]$ is the **best-response** function $X_k = X_{\lambda_k}$.

Application of GCG

Algorithm 2: Fictitious play

Input: $\bar{X}_0 \in L^\infty(Y, X_{\mathrm{ad}})$;

for $k = 0, 1, \ldots$ **do**

- [Prediction] Compute $\lambda_k = \nabla\phi\big(\int_Y \bar{X}_k(y)\,\mathrm{d}\mu(y)\big)$.
- [Best-response] Compute $X_k = X_{\lambda_k}(\cdot)$.
- [Learning] Set $\bar{X}_{k+1} = (1-\delta_k)\bar{X}_k + \delta_k X_k$, for some $\delta_k \in [0,1]$.

end

The GCG algo. (with $\delta_k = \frac{1}{k+1}$) coincides with the **fictitious play**.

The contribution F_1 of the potential cost is strongly convex in $L^2_\mu(Y;\mathbb{R}^d)$, thus **linear convergence** can be achieved.

Exploitability

The primal-dual gap is given by

$$\begin{aligned}\sigma(\bar{X}_k) &= F_{\text{lin}}[\bar{X}_k](\bar{X}_k) - F_{\text{lin}}[\bar{X}_k](X_{\lambda_k}) \\ &= \int_Y \underbrace{\left[f_{\lambda_k,y}(\bar{X}_k(y)) - \inf_{x \in X_{\text{ad}}} f_{\lambda_k,y}(x) \right]}_{\substack{\text{Best possible improvement for agent } y, \\ \text{assuming } \lambda_k \text{ fixed}}} \mathrm{d}\mu(y) \geq 0.\end{aligned}$$

In the present context, $\sigma(\bar{X}_k)$ is referred to as **exploitability**.

References

Connexion "best-reply" and Frank-Wolfe in a continuous-time setting:

Sorin. Continuous Time Learning Algorithms in Optimization and Game Theory, *Dynamic Games and Applications*, 2022.

Applying Frank-Wolfe to potential games is an old idea:

Fukushima. A modified Frank-Wolfe algorithm for solving the traffic assignment problem, *Transportation research*, 1984.

N-player differential game

We begin with a differential game with N players.

Data:

- N i.i.d. random variables $(\bar{X}_0^i)_{i=1,\dots,N}$ in $\mathbb{R}^d$, with probability distribution $m_0 \in \mathcal{P}(\mathbb{R}^n)$
- N independent Brownian motions $(W_t^i)_{t\in[0,T],\, i=1,\dots,N}$
- a running cost $L\colon \mathbb{R}^d \to \mathbb{R}$
- a terminal cost $g\colon \mathbb{R}^d \to \mathbb{R}$
- a price function $\psi\colon \mathbb{R}^d \to \mathbb{R}^d$.

Decision variables of the agent i:

- a control A^i (an adapted stochastic process)
- the associated state X^i, solution to:

$$\mathrm{d}X_t^i = A_t^i\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}W_t^i, \quad X_0^i = \bar{X}_0^i.$$

N-player differential game

Equilibrium problem: find $N+1$ stochastic processes $(\bar{A}^1, ..., \bar{A}^N)$ and λ such that

$$\begin{cases} \bar{A}^i \in \underset{A^i \in \mathbb{L}^2(0,T;\mathbb{R}^d)}{\operatorname{argmin}} \ J[\lambda](A^i) \\ \text{where } J[\lambda](A^i) = \mathbb{E}\Big[\int_0^T \Big(L(A^i_t) + \langle \lambda_t, A^i_t \rangle \Big) \, \mathrm{d}t + g(X^i_T) \Big] \end{cases}$$

and

$$\lambda_t = \psi\Big(\frac{1}{N} \sum_{j=1}^N \bar{A}^j_t \Big).$$

Decision $(\bar{A}^1, ..., \bar{A}^N)$ of the agents $\xleftarrow{\text{Stochastic control pb.}}$ $\xrightarrow[\text{Coupling function}]{}$ Price λ (common to all)

MFG

Mean-field game (MFG): a **limit** model for the above game, as N goes to infinity. At the limit, we "expect":

- the price λ to be deterministic
- the controls of the agents to have the same distribution and to be independent.

The MFG can be posed as an equilibrium problem involving a single pair $(\bar{X}, \bar{A})$ (for a "representative agent") and λ:

$$\begin{cases} \bar{A} \in \underset{A \in \mathbb{L}^2(0,T;\mathbb{R}^d)}{\operatorname{argmin}} \ J[\lambda](A) \\ \text{where } J[\lambda](A) = \mathbb{E}\Big[\int_0^T \Big(L(X_t, A_t) + \langle \lambda_t, A_t \rangle \Big) \, \mathrm{d}t + g(X_T) \Big] \end{cases}$$

and

$$\lambda_t = \psi(\mathbb{E}[\bar{A}_t]).$$

PDE formulation

In the PDE formulation of the problem, the optimal control is characterized via a **feedback function** v which is such that $\bar{A}_t = v(\bar{X}_t, t)$, almost surely.

$$\boxed{\begin{array}{c}\text{Feedback } v\\ \text{of the}\\ \text{representative agent}\end{array}} \quad \underset{\text{Coupling function}}{\overset{\text{Stochastic control pb.}}{\leftrightarrows}} \quad \boxed{\text{Price } \lambda}$$

PDE formulation

Our MFG model involves two additional variables:

- The feedback v is deduced from λ via the **value function** u.
- The price λ is deduced from v via the **distribution** m.

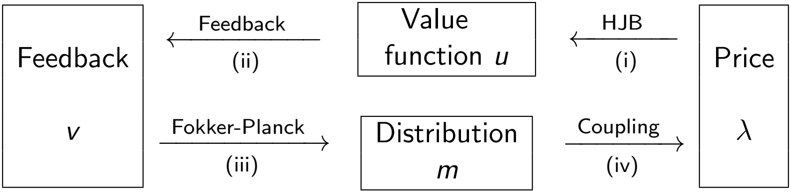

Remark. Our model is a mean field game of controls, since λ depends on m and v.

PDE formulation

From λ to v.

- The value function u is the solution to the **Hamilton-Jacobi-Bellman** (HJB) equation:

$$\begin{cases} -\partial_t u - \nabla u + H(\nabla u + \lambda) = & 0 \\ u(T,x) = & g(x), \end{cases} \tag{i}$$

  where $H(p) = \sup_\alpha \big(-\langle p, \alpha\rangle - L(\alpha)\big)$.

  Notation: $u = \mathrm{HJB}(\lambda)$.

- The **optimal feedback** is given by

$$v(t,x) = -\nabla H(\nabla u(t,x) + \lambda(t)). \tag{ii}$$

PDE formulation

From v to λ.

- Let m denote the probability distribution of X (when v is used). Then m is the solution to the **Fokker-Planck equation**:

$$\begin{cases} \partial_t m - \Delta m + \mathrm{div}(mv) = & 0, \\ m(0,x) = & m_0(x). \end{cases} \tag{iii}$$

  Notation: $m = \mathrm{FP}(v)$.

- Finally, λ can be described by

$$\lambda(t) = \psi\Big(\int v(t,x) m(t,x) \,\mathrm{d}x \Big). \tag{iv}$$

MFG: the coupled system (i)-(iv) with unknown (m, v, u, P).

References

Pioneering works:

Lasry, Lions. Mean field games, *Japanese Journal of Maths.*, 2007.

Huang, Malhamé, Caines. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle, *Communications in Info. and Systems*, 2006.

Our work related to the model:

Bonnans, Hadikhanloo, Pfeiffer. Schauder estimates for a class of potential mean field games of controls., *Applied Maths. Optim.*, 2021.

Assumptions

Periodicity:

- $g(x+y) = g(x)$ for all $y \in \mathbb{Z}^d$, so that the PDEs of the MFG can be considered on $Q := \mathbb{T}^d \times [0, T]$ with periodic boundary conditions.

Monotonicity assumptions:

- $\psi = \nabla\phi$, where ϕ is convex
- L is strongly convex.

Regularity assumptions:

- $L(v) \leq C(1 + \|v\|^2)$
- $H \in C^2(\mathbb{R}^d)$, H, ∇H, $\nabla^2 H$ are locally Hölder continuous
- ψ is Lipschitz continuous and bounded
- $m_0 \in C^3(\mathbb{T}^d)$, $m_0 \geq 0$, $\int_{\mathbb{T}^d} m_0(x)\,\mathrm{d}x = 1$, $g \in C^3(\mathbb{T}^d)$.

Main result

Theorem

There exists a unique **classical solution** $(\bar{m}, \bar{v}, \bar{u}, \bar{P})$ to the MFG system *(i)-(ii)-(iii)-(iv)*, with

$$\begin{aligned} &\bar{u} \in C^{2+\beta,1+\beta/2}(Q), && \bar{m} \in C^{2+\beta,1+\beta/2}(Q), \\ &\bar{v} \in C^{\beta}(Q), D_x v \in C^{\beta}(Q), && \bar{P} \in C^{\beta}(0,T), \end{aligned}$$

for some $\beta \in (0,1)$.

Notation:

$$\begin{aligned} C^{2+\beta,1+\beta}(Q) := \Big\{ u \in C^{\beta}(Q) \,|\, \partial_t u \in C^{\beta}(Q), \\ \nabla u \in C^{\beta}(Q),\, \nabla^2 u \in C^{\beta}(Q) \Big\}. \end{aligned}$$

Potential formulation

Consider the cost function $\mathcal{J} \colon W^{2,1,p}(Q) \times L^\infty(Q) \to \mathbb{R}$,

$$\mathcal{J}(m,v) = \iint_Q L(v(x,t))m(x,t)\,\mathrm{d}x\,\mathrm{d}t + \int_0^T \phi\Big(\int_{\mathbb{T}^d} v(x,t)m(x,t)\,\mathrm{d}x \Big)\,\mathrm{d}t.$$

Lemma (Potential formulation)

Let $(\bar{u}, \bar{m}, \bar{v}, \bar{P})$ be the solution to (MFG). Then, $(\bar{m}, \bar{v})$ is a solution *to:*

$$\min_{\substack{m \in W^{2,1,p}(Q) \\ v \in L^\infty(Q,\mathbb{R}^k)}} \mathcal{J}(m,v) \quad s.t.: \begin{cases} \partial_t m - \Delta m + div(vm) = & 0, \\ m(x,0) = & m_0(x). \end{cases}$$

Convexity of the potential problem

Reformulate the potential problem:

- Change of variable $(m, v) \to (m, w) := (m, mv)$.

This yields an equivalent convex problem:

$$\left\{ \begin{array}{l} \min\limits_{(m,w)} \ \tilde{\mathcal{J}}(m,w) := \underbrace{\iint_Q L\left(\frac{w}{m}\right) m \,\mathrm{d}x\,\mathrm{d}t + \int_{\mathbb{T}^d} g m(\cdot, T)\mathrm{d}x}_{=:\tilde{\mathcal{J}}_1(m,w)} \\ \qquad\qquad\qquad\qquad + \underbrace{\int_0^T \phi\left(\int_{\mathbb{T}^d} w dx\right) \mathrm{d}t}_{=:\tilde{\mathcal{J}}_2(m,w)} \\ \text{s.t.:} \ \left\{ \begin{aligned} \partial_t m - \sigma \Delta m + \operatorname{div}(w) &= 0, \\ m(x,0) &= m_0(x). \end{aligned} \right. \end{array} \right.$$

GCG

The linearized problem at $(\bar{m}_k, \bar{w}_k)$ reads:

$$\begin{cases} \min\limits_{(m,w)} \displaystyle\iint_Q L\left(\tfrac{w}{m}\right) m \,\mathrm{d}x\,\mathrm{d}t + \int_{\mathbb{T}^d} g m(\cdot, T)\mathrm{d}x \\ \qquad\qquad + \displaystyle\int_0^T \left\langle \psi\left(\textstyle\int_{\mathbb{T}^d} \bar{w}_k dx \right), \textstyle\int_{\mathbb{T}^d} w dx \right\rangle \mathrm{d}t \\ \text{s.t.:} \begin{cases} \partial_t m - \sigma \Delta m + \mathrm{div}(w) = & 0, \\ m(x,0) = & m_0(x). \end{cases} \end{cases}$$

Let us set:

$$\lambda_k = \psi\left(\textstyle\int_{\mathbb{T}^d} \bar{w}_k dx \right)$$

GCG

After change of variable $(m, v) = (m, w/m)$, we obtain the following linearized problem:

$$\begin{cases} \min\limits_{(m,v)} \displaystyle\iint_Q \big(L(v) + \langle \lambda_k, v \rangle\big) m \, \mathrm{d}x \, \mathrm{d}t + \int_{\mathbb{T}^d} g m(\cdot, T) \mathrm{d}x \\ \text{s.t.:} \quad \begin{cases} \partial_t m - \sigma \Delta m + \operatorname{div}(mv) = & 0, \\ m(x, 0) = & m_0(x). \end{cases} \end{cases}$$

Observation: the **linearized problem** is the potential formulation of the **stochastic control problem** of the representative agent, for $\lambda = \lambda_k$.

A solution (m_k, w_k) is found as follows:

- Compute $u_k = \mathrm{HJB}(\lambda_k)$, $v_k = -\nabla H(\nabla u_k + \lambda_k)$.
- Compute $m_k = \mathrm{FP}(v_k)$, $w_k = m_k v_k$.

Application of GCG

Algorithm 3: Fictitious play

Input: $\bar{v}_0$; Compute: $\bar{m}_0 = \text{FP}(\bar{v}_0)$, $\bar{w}_0 = \bar{m}_0 \bar{v}_0$;

for $k = 0, 1, \dots$ **do**

[Prediction] Compute $\lambda_k = \psi\big(\int_{\mathbb{T}^d} \bar{w}_k \, \mathrm{d}\mu(y) \big)$. (iv)

[Best-resp.] Compute $u_k = \text{HJB}(\lambda_k)$. (i)

Set $v_k = -\nabla H(\nabla u_k + \lambda_k)$. (ii)

Compute $m_k = \text{FP}(v_k)$. (iii)

Set $w_k = m_k v_k$.

[Learning] Choose $\delta_k \in [0, 1]$.

Set $\bar{m}_{k+1} = (1 - \delta_k)\bar{m}_k + \delta_k m_k$.

Set $\bar{w}_{k+1} = (1 - \delta_k)\bar{w}_k + \delta_k w_k$.

end

Theorem

- The GCG algorithm is **well-posed**. It generates sequences $(\bar{m}_k, \bar{w}_k)$ and (m_k, w_k) in $(C^{2,1}(Q) \times C^{1,0}(Q))$, $v_k \in C^{1,0}(Q)$, $u_k \in C^{2,1}(Q)$, and $P_k \in C^0([0,T])$.
- Let $\varepsilon_k = \tilde{\mathcal{J}}(\bar{m}_k, \bar{w}_k) - \tilde{\mathcal{J}}(\bar{m}, \bar{w})$. There exist constants $C > 0$ and $\gamma \in (0,1)$ such that
  - If $\delta_k = \dfrac{1}{k+1}$, then $\varepsilon_k \leq \dfrac{C \ln(k)}{k}$.
  - If $\delta_k = \dfrac{2}{k+2}$, then $\varepsilon_k \leq \dfrac{C}{k}$.
  - If δ_k is determined by an **adaptative** rule, then $\varepsilon_k \leq C\gamma^k$.
- Moreover,

$$
\begin{aligned}
&\|\bar{m}_k - \bar{m}\|_{L^\infty(0,T;L^2(\mathbb{T}^d))} + \|\bar{w}_k - \bar{w}\|_{L^2(Q)} \\
&\quad + \|P_k - \bar{P}\|_{L^2(0,T)} + \|u_k - \bar{u}\|_{L^\infty(Q)} \leq C\sqrt{\varepsilon_k}.
\end{aligned}
$$

Elements of proof

- Well-posedness: based on estimates for parabolic PDEs.
- Linear convergence: the cost $\tilde{J}$ is not strongly convex $\rightarrow$ difficulty.

Let $k \in \mathbb{N}$. The challenge is to prove estimates of the form:

$$\|\bar{m}_k - m_k\| = O(\sqrt{\sigma_k}) \quad \text{and} \quad \|\bar{w}_k - w_k\| = O(\sqrt{\sigma_k}).$$

Let $\bar{v}_k = \bar{w}_k / \bar{m}_k$. By construction:

- The feedback v_k is optimal for the stochastic optimal control problem with $\lambda = \lambda_k$.
- The feedback $\bar{v}_k$ is σ_k-optimal for this problem.

Elements of proof

A standard calculation (involving some integration by parts yields)

$$\sigma_k = \mathcal{J}_{\text{lin}}(\bar{v}_k) - \mathcal{J}_{\text{lin}}(v_k) \geq \iint_Q \bar{m}_k \|\bar{v}_k - v_k\|^2.$$

Let $\zeta_k = \bar{m}_k(\bar{v}_k - v_k)$. We have $\|\zeta_k\|_{L^2(Q)} \leq \sqrt{\sigma_k}$.
Let $\mu = \bar{m}_k - m_k$. It is the solution to the PDE

$$\mu_t - \Delta\mu + \text{div}(\mu v_k) = -\text{div}(\zeta_k).$$

The classical theory of parabolic PDEs yields the estimate

$$\|\mu\|_{L^\infty(0,T;L^2(\mathbb{T}^d))} \leq C\|\zeta\|_{L^2(Q)}.$$

We finally obtain

$$\|\bar{w}_k - w_k\|_{L^2(Q)} = \|\zeta + \mu v_k\|_{L^2(Q)} \leq C\sqrt{\sigma_k}.$$

Convergence results

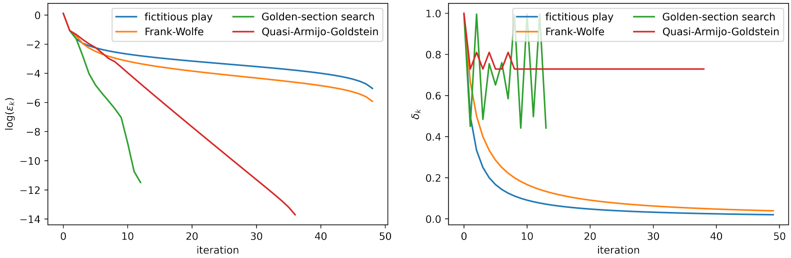

(a) Gap $\ln(\varepsilon_k)$ (b) Stepsize δ_k

Figure: Convergence results for an MFG with price term

Convergence results

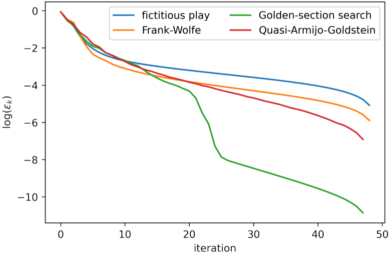

(a) Gap $\ln(\varepsilon_k)$

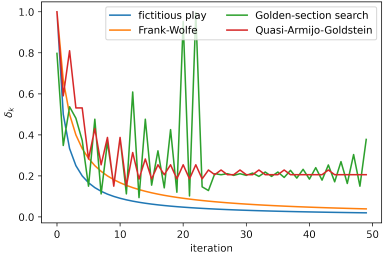

(b) Stepsize δ_k

Figure: Convergence results for an MFG with congestion term

References

Fictitious play for MFGs (non-convex potential setting):

- Cardaliaguet, Hadikhanloo. Learning in mean field games: the fictitious play, *ESAIM: Control, Optimisation and Calculus of Variations, 2017.*

Frank-Wolfe for MFGs (discrete setting):

- Elie, Piérolat, Laurière, Geist, Pietquin. Approximate fictitious play for mean field games, *ArXiv preprint*, 2020.

Preliminary version of our work:

- Bonnans, Lavigne, Pfeiffer. Generalized conditional gradient and learning in potential mean field games, *ArXiv preprint*, 2021.

Thank you for your attention!