Mean Field Games and Gradient Flows

H. Mete Soner, ORFE, Princeton

MAFE Workshop Risk Mitigation

September 3, 2024

This talk is based on joint projects with :

Rene Carmona (Princeton),

Quentin Cormier (INRIA and Ecole Polytechnique),

Felix Höfer (Princeton).

I want to establish connections between several mean-field models with the following characteristics :

- Agents or particles or players are identical and are subject to independent idiosyhncrotic noise.
- We assume that the system is large and take the infinite limit.
- By law of large numbers, the collective behavior of the agents are described their distribution. Hence, the state is the set of probability measures.

Models are :

- Classical dynamical systems assuming that the energy and or the entropy is given.
- Mean Field Games that have similar qualitative behavior.
- The related Mean Field Control.
- How do we construct one from the other systematically ?

Outline

Kuramoto

Kuramoto (1975) considered a population of N coupled phase oscillators θ_t^k having natural frequencies ω^k distributed with a given density, and whose dynamics are governed by

$$\frac{\mathrm{d}}{\mathrm{d}t} X_t^k = \omega^k - \frac{\kappa}{N} \sum_{j=1}^{N} \sin(X_t^k - X_t^j) \approx \omega^k - \kappa(X_t^k - \bar{X}_t), \qquad k = 1, \dots, N,$$

where $\bar{X}_t$ is the mean location. For large κ values they attract each other.

The following “energy” is related to this system :

$$\mathcal{E} := \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} \sin^2((x-y)/2)\ \mu_t^N(\mathrm{d}x)\mu_t^N(\mathrm{d}y), \quad \text{where} \quad \mu_t^N(\mathrm{d}x) := \frac{1}{N} \sum_{k=1}^{N} \delta_{X_t^k}(\mathrm{d}x).$$

This is from the github page of Helge Dietert from Paris.

`https://hdietert.github.io/static/kuramoto-animation/kuramoto.html`

We write the energy as

$$\mathcal{E} = \mathcal{F}(\mu_t^N), \quad \text{where} \quad \mathcal{F}(\mu) := \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} \sin^2((x-y)/2)\ \mu(\mathrm{d}x)\mu(\mathrm{d}y).$$

Then, the linear derivative is given by,

$$\delta_\mu \mathcal{F}(\mu)(x) = 2 \int_{-\pi}^{\pi} \sin^2((x-y)/2)\ \mu(\mathrm{d}y) = \int_{-\pi}^{\pi} (1 - \cos(x-y))\ \mu(\mathrm{d}y).$$

We directly calculate that Lions derivative is given by,

$$\partial_\mu \mathcal{F}(\mu)(c) := \nabla_x(\delta_\mu \mathcal{F}(\mu)(x)) = \int_{-\pi}^{\pi} \sin(x-y)\ \mu(\mathrm{d}y).$$

Hence, the Kuramoto equation with $\omega^k = 0$ can be written as

$$\frac{\mathrm{d}}{\mathrm{d}t} X_t^k = -\frac{\kappa}{N} \sum_{j=1}^{N} \sin(X_t^k - X_t^j) = -\int_{-\pi}^{\pi} \sin(X_t^k - y)\ \mu_t^N(\mathrm{d}y) = -\partial_\mu \mathcal{F}(\mu_t^N)(X_t^k).$$

We add Brownian motion and write the Kuramoto equation as

$$\mathrm{d}X_t^k = -\partial_\mu \mathcal{F}(\mu_t^N)(X_t)\mathrm{d}t + \sigma \mathrm{d}W_t = -\nabla_x(\delta_\mu \mathcal{F}(\mu_t^N)(X_t^k))\mathrm{d}t + \sigma \mathrm{d}W_t^k.$$

- Particles are identical with independent idiosyncratic noise. By Law of Large Numbers,

$$\mu_t^N(\mathrm{d}x) := \frac{1}{N}\sum_{k=1}^{N} \delta_{X_t^k}(\mathrm{d}x) \rightharpoonup \mu_t(\mathrm{d}x), \qquad \text{as } N \to \infty.$$

  where μ_t is the law of the 'representative' particle.
- Then, the equation for the representative particle is the following McKean-Vlasov equation,

$$\mathrm{d}X_t = -\partial_\mu \mathcal{F}(\mu_t)(X_t)\mathrm{d}t + \sigma \mathrm{d}W_t, \quad \text{and} \quad \mu_t = Law(X_t).$$

- Hence, Kuramoto equation is a Langevin flow in $\mathcal{L}^2$.

Mean Field Game

Kuramoto Mean Field Game

1. Start with a deterministic flow of probability measures $\boldsymbol{\mu} = (\mu_t)_{t\geq 0}$ with $\mu_0 = \mu$.
2. Find the optimal control $\boldsymbol{\alpha}^{*,\boldsymbol{\mu}} = (\alpha_t^{*,\boldsymbol{\mu}})_{t\geq 0}$ minimizing,

$$\boldsymbol{\alpha} = (\alpha_t)_{t\geq 0} \quad\mapsto\quad J(\boldsymbol{\alpha}\,;\,\boldsymbol{\mu}) := \mathbb{E}\int_0^\infty e^{-\beta t}\,[\kappa\, L(X_t^{\boldsymbol{\alpha}},\mu_t) + \tfrac{1}{2}(\alpha_t)^2]\,\mathrm{d}t,$$

where $\mathrm{d}X_t^{\boldsymbol{\alpha}} = \alpha_t \mathrm{d}t + \sigma \mathrm{d}B_t$, $Law(X_0) = \mu_0$, and

$$L(x,\mu) := 2\int_{-\pi}^{\pi} \sin^2((x-y)/2)\;\mu(\mathrm{d}y) = \delta_\mu \mathcal{F}(\mu)(x).$$

3. Find a fixed point $\mu_t = Law(X_t^{\boldsymbol{\alpha}^{*,\boldsymbol{\mu}}})$.

Synchronization of coupled oscillators is a game, by Yin, Mehta, Meyn, Shanbhag, IEEE (2011).

Synchronization in a Kuramoto Mean Field Game, Carmona, Cormier, Soner, CPDE (2023).

A Remark

- In finite player games, knowing Nash equilibria are characterized by the strategies.
- In the mean field limit, representative agent's action do not impact the location of the other players.
- Hence, the distribution of the 'other' players suffice to describe the minimization problem of the representative agent.
- However, we could also focus on the feedback controls as the feedback controls determine the distribution.
- Technically, working with probability distributions has many advantages.

Let $U(\mathrm{d}x) := \frac{\mathrm{d}x}{2\pi}$ be the uniform measure on the circle. Then,

$$L(x, U) = \int_{-\pi}^{\pi} 2\sin^2\left(\tfrac{x-y}{2}\right)\, U(\mathrm{d}x) \equiv 1.$$

Then, the control problem corresponding to the stationary flow U is

$$\text{minimize} \quad \boldsymbol{\alpha} = (\alpha_t)_{t\geq 0} \;\mapsto\; J(\boldsymbol{\alpha}\,;\; U) := \mathbb{E}\int_0^\infty e^{-\beta t}\,[\kappa + \tfrac{1}{2}(\alpha_t)^2]\,\mathrm{d}t.$$

Cleary the optimal solution is $\alpha^* \equiv 0$, and the optimal state is $\mathrm{d}X_t^* = 0\,\mathrm{d}t + \sigma \mathrm{d}B_t$. Hence, $X_t^* = X_0^* + \sigma B_t$ and as $Law(X_0^*) = U$, we have $Law(X_t^*) = U$ as well. Hence,

The uniform measure U is a stationary Nash equilibrium for every parameter.

The uniform distribution represents incoherence or lack of synchronization.

Critical interaction parameter is $\kappa_c := \beta\sigma^2 + \sigma^4/2$.

Theorem (Sub-critical interaction : incoherence)

(Carmona, Cormier, S.)(2023) For $\kappa < \kappa_c$, the uniform measure is locally stable.
Namely, there exist a positive constant $\rho > 0$ depending on β, σ, κ such that for any μ_0 satisfying $d(\mu_0 - U) \leq \rho$, there exists a solution $\boldsymbol{\mu} = (\mu_t)_{t\geq 0}$ of the Kuramoto mean field game with interaction parameter κ with $\mu_0 = \nu$ and μ_t converges in law to the uniform distribution as t tends to infinity.

Theorem (Super-critical interaction : synchronization)

(Carmona, Cormier, S.)(2023) For $\kappa > \kappa_c$, there are non-trivial stationary Nash equilibria.

Potential Structure

Outline

Central Planner Problem

- In Mean Field Games, we start with a flow of probability measures $\boldsymbol{\mu} = (\mu_t)_{t\geq 0}$, and find the optimal response $\boldsymbol{\alpha}^{*,\boldsymbol{\mu}}$ by minimizing,

$$\boldsymbol{\alpha} = (\alpha_t)_{t\geq 0} \quad \mapsto \quad J_g(\boldsymbol{\alpha}\,;\; \boldsymbol{\mu}) := \mathbb{E}\int_0^\infty e^{-\beta t}\left[L(X_t^{\boldsymbol{\alpha}}, \mu_t) + \tfrac{1}{2}(\alpha_t)^2\right]\mathrm{d}t,$$

  where $\mathrm{d}X_t^{\boldsymbol{\alpha}} = \alpha_t \mathrm{d}t + \sigma \mathrm{d}B_t$. Then, look for a fixed point $\mu_t = Law(X_t^*)$.

- In the Central Planner problem, the representative agent minimizes

$$\boldsymbol{\alpha} = (\alpha_t)_{t\geq 0} \quad \mapsto \quad J_p(\boldsymbol{\alpha}) := \mathbb{E}\int_0^\infty e^{-\beta t}\left[L(X_t^{\boldsymbol{\alpha}}, \mathcal{L}_t^{\boldsymbol{\alpha}}) + \tfrac{1}{2}(\alpha_t)^2\right]\mathrm{d}t,$$

  where $\mathrm{d}X_t^{\boldsymbol{\alpha}} = \alpha_t \mathrm{d}t + \sigma \mathrm{d}B_t$, and $\mathcal{L}_t^{\boldsymbol{\alpha}} = Law(X_t^{\boldsymbol{\alpha}})$.

- *In general, they are two different problems, and the difference is the price of anarchy.*

- Note that $\mathbb{E}[L(X_t^{\boldsymbol{\alpha}}, \mathcal{L}_t^{\boldsymbol{\alpha}})] = \int L(x, \mathcal{L}_t^{\boldsymbol{\alpha}})\; \mathcal{L}_t^{\boldsymbol{\alpha}}(\mathrm{d}x) =: L_p(\mathcal{L}_t^{\boldsymbol{\alpha}})$.

- In potential games, the running cost is given by $L(x,\mu) = \delta_\mu \mathcal{F}(\mu)(x)$ for some $\mathcal{F}(\mu)$.
- The running cost of the central planner is

$$L_p(\mu) = \int \delta_\mu \mathcal{F}(\mu)(x)\ \mu(\mathrm{d}x).$$

- In the Mean Field Control, we consider the problem of minimizing

$$\boldsymbol{\alpha} = (\alpha_t)_{t\geq 0} \quad \mapsto \quad J_c(\boldsymbol{\alpha}) := \mathbb{E}\int_0^\infty e^{-\beta t}\left[\mathcal{F}(\mathcal{L}_t^{\boldsymbol{\alpha}}) + \tfrac{1}{2}(\alpha_t)^2\right] \mathrm{d}t.$$

- In all problems, $\mathrm{d}X_t^{\boldsymbol{\alpha}} = \alpha_t \mathrm{d}t + \sigma \mathrm{d}B_t$, and $\mathcal{L}_t^{\boldsymbol{\alpha}} = Law(X_t^{\boldsymbol{\alpha}})$.
- In general,

$$L_p(\mu) = \int \delta_\mu \mathcal{F}(\mu)(x)\ \mu(\mathrm{d}x) \neq \mathcal{F}(\mu).$$

- In the Kuramoto problem

$$L_p(\mu) = 2\int_{-\pi}^{\pi} \sin^2((x-y)/2)\mu(\mathrm{d}y)\mu(\mathrm{d}x) = 2\mathcal{F}(\mu).$$

Potential Case

- In potential games, the running cost is given by $L(x,\mu) = \delta_\mu \mathcal{F}(\mu)(x)$ for some $\mathcal{F}(\mu)$.
- The running cost of the central planner is

$$L_p(\mu) = \int \delta_\mu \mathcal{F}(\mu)(x)\ \mu(\mathrm{d}x).$$

- In the Mean Field Control, we consider the problem of minimizing

$$\boldsymbol{\alpha} = (\alpha_t)_{t\geq 0} \quad \mapsto \quad J_c(\boldsymbol{\alpha}) := \mathbb{E}\int_0^\infty e^{-\beta t}\left[\mathcal{F}(\mathcal{L}_t^{\boldsymbol{\alpha}}) + \tfrac{1}{2}(\alpha_t)^2\right] \mathrm{d}t.$$

- In all problems, $\mathrm{d}X_t^{\boldsymbol{\alpha}} = \alpha_t \mathrm{d}t + \sigma \mathrm{d}B_t$, and $\mathcal{L}_t^{\boldsymbol{\alpha}} = Law(X_t^{\boldsymbol{\alpha}})$.
- In general,

$$L_p(\mu) = \int \delta_\mu \mathcal{F}(\mu)(x)\ \mu(\mathrm{d}x) \neq \mathcal{F}(\mu).$$

- In the Kuramoto problem

$$L_p(\mu) = 2\int_{-\pi}^{\pi} \sin^2((x-y)/2)\mu(\mathrm{d}y)\mu(\mathrm{d}x) = 2\mathcal{F}(\mu).$$

In all problems, $\mathrm{d}X_t^{\boldsymbol{\alpha}} = \alpha_t \mathrm{d}t + \sigma \mathrm{d}B_t$, and $\mathcal{L}_t^{\boldsymbol{\alpha}} = Law(X_t^{\boldsymbol{\alpha}})$.

- Mean Field game. Given μ_t, minimize

$$\mathbb{E}\int_0^\infty e^{-\beta t}\left[L_g(X_t^{\boldsymbol{\alpha}}, \mu_t) + \tfrac{1}{2}(\alpha_t)^2\right] \mathrm{d}t, \qquad L_g(x,\mu) = \delta_\mu \mathcal{F}(\mu)(x),$$

  and find the fixed point $\mu_t = Law(X_t^*)$.

- Central Planner Problem is to minimize

$$J_c(\boldsymbol{\alpha}) := \mathbb{E}\int_0^\infty e^{-\beta t}\left[L_p(\mathcal{L}_t^{\alpha}) + \tfrac{1}{2}(\alpha_t)^2\right] \mathrm{d}t, \qquad L_p(\mu) = \int \delta_\mu \mathcal{F}(\mu)(x)\ \mu(\mathrm{d}x).$$

- Mean Field Control is to minimize

$$J_c(\boldsymbol{\alpha}) := \mathbb{E}\int_0^\infty e^{-\beta t}\left[L_c(\mathcal{L}_t^{\alpha}) + \tfrac{1}{2}(\alpha_t)^2\right] \mathrm{d}t, \qquad L_c(\mu) = \mathcal{F}(\mu).$$

Theorem

Suppose that $L(x, \mu) = \delta_\mu \mathcal{F}(\mu)(x)$. Then, any minimizer of the Mean Field Control problem is a Nash equilibrium of the Mean Field game problem.

- Some suggest this connection as a selection mechanism when there are multiple Nash equilibria.
- Although the minimizer of the Mean Field Control problem is a Nash equilibrium of the Mean Field Game, the value functions are not equal as the running costs are different.

Stable Solutions in Potential Mean Field Game Systems, by Briani, Cardaliaguet, NoDEA (2015).

Potential Mean-Field Games and Gradient Flows, Höfer, Soner, archiv (2024).

Mean Field Control

The general problem is

$$v(\mu) := \inf_{\boldsymbol{\alpha}} J_c(\mu, \boldsymbol{\alpha}) \ = \mathbb{E} \int_0^\infty e^{-\beta t} \left[\mathcal{F}(\mathcal{L}_t^{\boldsymbol{\alpha}}) + \tfrac{1}{2}(\alpha_t)^2\right] \mathrm{d}t,$$

where $\mathrm{d}X_t^{\boldsymbol{\alpha}} = \alpha_t \mathrm{d}t + \sigma \mathrm{d}B_t \in \mathcal{X}$, $\mathcal{L}_t^{\boldsymbol{\alpha}} = Law(X_t^{\boldsymbol{\alpha}})$, and $\mathcal{L}_0^{\boldsymbol{\alpha}} = \mu$.

By dynamic programming, we see that the value function v solves,

$$\boxed{\beta v(\mu) = \mathcal{H}(\mu, \delta_\mu v(\mu)) + \mathcal{F}(\mu), \qquad \mu \in \mathcal{P}(\mathcal{X}),}$$

where $\mathcal{P}(\mathcal{X})$ is the set of probability measures on the state space $\mathcal{X}$ and

$$\mathcal{H}(\mu, \varphi) = \int_{\mathcal{X}} H(\nabla_x \varphi(x), D^2\varphi(x)) \ \mu(\mathrm{d}x),$$

$$H(p, A) := \inf_{\alpha}(\alpha \cdot p + \frac{1}{2}|\alpha|^2) + \frac{1}{2}\mathrm{trace}(\sigma\sigma^T A), \quad \Rightarrow \quad \alpha^* = -p.$$

Since

$$\beta v(\mu) = \mathcal{H}(\mu, \delta_\mu v(\mu)) + \mathcal{F}(\mu), \qquad \mu \in \mathcal{P}(\mathcal{X}),$$

and

$$\mathcal{H}(\mu, \varphi) = \int_{\mathcal{X}} H(\nabla_x \varphi(x), D^2\varphi(x)) \, \mu(\mathrm{d}x),$$

$$H(p, A) := \inf_\alpha(\alpha \cdot p + \frac{1}{2}|\alpha|^2) + \frac{1}{2}\mathrm{trace}(\sigma\sigma^T A), \quad \Rightarrow \quad \alpha^* = -p,$$

we have

$$\alpha^*(\mu)(x) = -\nabla_x(\delta_\mu v(\mu)(x)) = -\partial_\mu v(\mu)(x), \qquad x \in \mathcal{X}, \ \mu \in \mathcal{P}(\mathcal{X}).$$

So the optimally controlled state equation. is

$$\mathrm{d}X_t^* = \alpha^*(\mathcal{L}_t^*)(X_t^*) + \sigma \mathrm{d}W_t = -\partial_\mu v(\mathcal{L}_t^*)(X_t^*) + \sigma \mathrm{d}W_t.$$

- The value function is given by,

$$v(\mu) = \mathbb{E}\int_0^\infty e^{-\beta t}\left[\mathcal{F}(\mathcal{L}_t^{\boldsymbol{\alpha}}) + \tfrac{1}{2}(\alpha_t)^2\right] \mathrm{d}t.$$

- The optimally controlled state solves,

$$\boxed{\mathrm{d}X_t^* = -\partial_\mu v(\mu_t^*)(X_t^*) + \sigma \mathrm{d}W_t, \quad \text{and} \quad \mu_t^* = Law(X_t^*).}$$

- Compare it to the original Langevin equation,

$$\boxed{\mathrm{d}X_t = -\partial_\mu \mathcal{F}(\mu_t)(X_t) + \sigma \mathrm{d}W_t, \quad \text{and} \quad \mu_t = Law(X_t).}$$

- In most cases, v is similar to the original energy functional $\mathcal{F}$.

Concluding

- Mean Field Control can be used to construct Nash equilibria for the Mean Field Games.
- Mean Field formalism produce models that are analogous to gradient flows.

THANK YOU FOR YOUR ATTENTION.

Synchronization in a Kuramoto Mean Field Game

with Rene Carmona and Quentin Cormier,
Communications in Partial Differential Equations (2023).

Potential Mean-Field games and gradient flows, with Felix Höfer, preprint (2024).

- Mean Field Control can be used to construct Nash equilibria for the Mean Field Games.
- Mean Field formalism produce models that are analogous to gradient flows.

THANK YOU FOR YOUR ATTENTION.

Synchronization in a Kuramoto Mean Field Game

with Rene Carmona and Quentin Cormier,
Communications in Partial Differential Equations (2023).

Potential Mean-Field games and gradient flows, with Felix Höfer, preprint (2024).

- $\mathrm{d}X_t = \alpha(t, X_t)\,\mathrm{d}t + \sigma \mathrm{d}W_t$.
- The Kolmogrov equation for the distribution μ_t is

$$\frac{\mathrm{d}}{\mathrm{d}t}\int \varphi(x)\mu_t(\mathrm{d}x) = \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[\varphi(X_t)] = \mathbb{E}[\alpha(t, X_t)\cdot\nabla\varphi(X_t) + \tfrac{1}{2}\mathrm{trace}(\sigma\sigma^T D^2\varphi(X_t))]$$
$$= \int (\alpha(t,x)\cdot\nabla\varphi(x) + \tfrac{1}{2}\mathrm{trace}(\sigma\sigma^T D^2\varphi(x)))\;\mathrm{d}\mu_t(\mathrm{d}x).$$

- Therefore, the Hamiltonian (with $\varphi = \delta_\mu v$), is given by,

$$\mathcal{H}(\mu,\varphi) = \inf_{\alpha(t,\cdot)}\int (\alpha(t,x)\cdot\nabla\varphi(x) + \tfrac{1}{2}|\alpha(t,x)|^2 + \tfrac{1}{2}\mathrm{trace}(\sigma\sigma^T D^2\varphi(x)))\;\mathrm{d}\mu(\mathrm{d}x)$$
$$= \int [\inf_{\alpha\in\mathbb{R}^d}\left(\alpha\cdot\nabla\varphi(x) + \tfrac{1}{2}|\alpha(t,x)|^2\right) + \tfrac{1}{2}\mathrm{trace}(\sigma\sigma^T D^2\varphi(x))]\;\mathrm{d}\mu(\mathrm{d}x)$$
$$= \int -\tfrac{1}{2}|\nabla\varphi(x)|^2 + \tfrac{1}{2}\mathrm{trace}(\sigma\sigma^T D^2\varphi(x))]\;\mathrm{d}\mu(\mathrm{d}x).$$