Long-Time Behavior of Optimal Particle Systems in Mean-Field Control

Presented by Wenbin Yan[1]

with Songbo Wang[2]

Université Paris Dauphine-PSL[1,2]

Young Researchers' Days

Outline

Mean-Field Control (MFC) study

- Mean-field control = optimal control of large systems of weakly interacting agents acting cooperatively under a central planner.
- Microscopic description = an N-particle stochastic control problem, where the planner controls exchangeable diffusions interacting through their empirical distribution.
- Macroscopic description = a representative-agent control problem whose dynamics and cost depend on the law of the population.
- Applications: economics (EDF), finance, crowd motion, large-scale machine learning models, etc. Guéant–Lasry–Lions ('11), Chizat–Bach ('18), Achdou–Buera–Lasry–Lions–Moll ('14), Lachapelle–Wolfram ('11), Lachapelle–Salomon–Turinici ('10)

Some references:

- McKean–Vlasov control and mean-field control: Budhiraja–Dupuis–Fischer (2012), Lacker (2015), Carmona–Delarue (2018).
- Quantitative convergence and propagation of chaos: Germain–Pham–Warin (2022), Cardaliaguet–Souganidis (2022), Cardaliaguet–Jackson–Mimikos-Stamatopoulos–Souganidis (2023).
- Long-time behavior and turnpike phenomena: Trélat–Zuazua (2015), Cardaliaguet–Porretta (2020), Cardaliaguet–Maillet–Yan (2025), Yan(2026).

Outline

Let $\mathcal{X} = \mathbb{R}^d$ or $\mathbb{T}^d$. For $x = (x_1, \ldots, x_N) \in \mathcal{X}^N$, set

$$\mu_x^N := \frac{1}{N} \sum_{i=1}^{N} \delta_{x_i}.$$

A central planner controls the whole system

$$dX_t^i = \alpha_t^i dt + \sqrt{2}\, dB_t^i, \qquad X_0^i = x_0^i, \qquad i = 1, \ldots, N,$$

and minimizes the social cost

$$J^N(\boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^T \left(\sum_{i=1}^{N} \ell(\alpha_t^i, X_t^i) + NF(\mu_{\boldsymbol{X}_t}^N)\right) dt + NG(\mu_{\boldsymbol{X}_T}^N)\right].$$

$\boldsymbol{X}_t := (X_t^1, \cdots, X_t^N)$,
$\ell : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ strongly convex for the first variable while regular enough for the second,
$F, G : \mathcal{P}(\mathcal{X}) \to \mathbb{R}$ convex functions regular enough for intrinsic derivative/Lions derivative.

A central planner controls the whole system

$$dX_t^i = \alpha_t^i dt + \sqrt{2}\, dB_t^i, \qquad X_0^i = x_0^i, \qquad i = 1, \dots, N,$$

and minimizes the social cost

$$J^N(\boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^T \left(\sum_{i=1}^N \ell(\alpha_t^i, X_t^i) + NF(\mu^N_{\boldsymbol{X}_t})\right) dt + NG(\mu^N_{\boldsymbol{X}_T})\right].$$

The value function of this optimal control problem is

$$U_t^N(x) = \inf_{\boldsymbol{\alpha}} \mathbb{E}\left[\int_t^T \left(\sum_{i=1}^N \ell(\alpha_s^i, X_s^i) + NF(\mu^N_{\boldsymbol{X}_s})\right) ds + NG(\mu^N_{\boldsymbol{X}_T})\right],$$

and solves the finite-dimensional HJB equation

$$\begin{cases} \partial_t U_t^N + \Delta_x U_t^N - h^N(\nabla_x U_t^N, x) + NF(\mu_x^N) = 0, \\ U_T^N(x) = NG(\mu_x^N), \end{cases}$$

where

$$h^N(p, x) := \sum_{i=1}^N h(p_i, x_i).$$

N-particle Control Problem

A central planner controls the whole system

$$dX_t^i = \alpha_t^i dt + \sqrt{2}\, dB_t^i, \qquad X_0^i = x_0^i, \qquad i = 1, \ldots, N,$$

and minimizes the social cost

$$J^N(\boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^T \left(\sum_{i=1}^N \ell(\alpha_t^i, X_t^i) + NF(\mu_{\boldsymbol{X}_t}^N)\right) dt + NG(\mu_{\boldsymbol{X}_T}^N)\right].$$

The optimal feedback is

$$\bar{\alpha}_t^i = -\nabla_p h(\nabla_i U_t^N(\boldsymbol{X}_t), X_t^i),$$

so the optimal particle system satisfies

$$dX_t^i = -\nabla_p h(\nabla_i U_t^N(\boldsymbol{X}_t), X_t^i) dt + \sqrt{2}\, dB_t^i.$$

A central planner controls the whole system

$$dX_t^i = \alpha_t^i dt + \sqrt{2}\, dB_t^i, \qquad X_0^i = x_0^i, \qquad i = 1, \dots, N,$$

and minimizes the social cost

$$J^N(\boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^T \left(\sum_{i=1}^N \ell(\alpha_t^i, X_t^i) + NF(\mu^N_{\boldsymbol{X}_t})\right) dt + NG(\mu^N_{\boldsymbol{X}_T})\right].$$

The optimal feedback is

$$\bar{\alpha}_t^i = -\nabla_p h(\nabla_i U_t^N(\boldsymbol{X}_t), X_t^i),$$

so the optimal particle system satisfies

$$dX_t^i = -\nabla_p h(\nabla_i U_t^N(\boldsymbol{X}_t), X_t^i) dt + \sqrt{2}\, dB_t^i.$$

If $m_t^N \in \mathcal{P}(X^N)$ denotes the law of $\boldsymbol{X}_t$, then

$$\begin{cases} \partial_t m_t^N = \Delta_{\mathbf{x}} m_t^N + \nabla_x \cdot \left(m_t^N \nabla_p h^N (\nabla_{\mathbf{x}} U_t^N, \mathbf{x})\right), \\ \quad m_0^N = \mathcal{L}(\boldsymbol{X}_0). \end{cases}$$

N-particle Control Problem

A central planner controls the whole system

$$dX_t^i = \alpha_t^i dt + \sqrt{2}\, dB_t^i, \qquad X_0^i = x_0^i, \qquad i = 1, \dots, N,$$

and minimizes the social cost

$$J^N(\boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^T \left(\sum_{i=1}^N \ell(\alpha_t^i, X_t^i) + NF(\mu_{\boldsymbol{X}_t}^N)\right) dt + NG(\mu_{\boldsymbol{X}_T}^N)\right].$$

So the analytical description of the N-particle optimal system is

$$\begin{cases} \partial_t U_t^N &= -\Delta_x U_t^N + h^N(\nabla_x U_t^N, x) - NF(\mu_x^N) \quad , U_T^N(x) = NG(\mu_x^N), \\ \partial_t m_t^N &= \Delta_x m_t^N + \nabla_x \cdot \left(m_t^N \nabla_p h^N(\nabla_x U_t^N, x)\right) \quad , m_0^N = \mathcal{L}(\boldsymbol{X}_0). \end{cases}$$

So the analytical description of the N-particle optimal system is

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) - NF(\mu_{\mathbf{x}}^N), \qquad U_T^N = NG(\mu_{\mathbf{x}}^N), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot \left(m_t^N \nabla_p h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x})\right), \quad m_0^N = \mathcal{L}(\mathbf{X}_0). \end{cases}$$

So the analytical description of the N-particle optimal system is

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) - NF(\mu_{\mathbf{x}}^N), \qquad U_T^N = NG(\mu_{\mathbf{x}}^N), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot \left(m_t^N \nabla_p h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) \right), \quad m_0^N = \mathcal{L}(\mathbf{X}_0). \end{cases}$$

Mean-field limit:

When $N \to \infty$, one expects

$$\frac{1}{N} U_t^N(x_1, \ldots, x_N) \approx U_t(\mu_{\mathbf{x}}^N),$$

where

$$U_t : \mathcal{P}(\mathcal{X}) \to \mathbb{R}$$

is the mean-field value function.

So the analytical description of the N-particle optimal system is

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) - NF(\mu_{\mathbf{x}}^N), \qquad U_T^N = NG(\mu_{\mathbf{x}}^N), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot \left(m_t^N \nabla_p h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) \right), \quad m_0^N = \mathcal{L}(\mathbf{X}_0). \end{cases}$$

Mean-field limit:

The analogue of the finite-dimensional HJB equation is the HJB equation on $\mathcal{P}(\mathcal{X})$:

$$\begin{cases} \partial_t U_t(m) = -\int_{\mathcal{X}} \Delta_x \delta_m U_t(m, x)\, m(dx) + \int_{\mathcal{X}} h(D_m U_t(m, x), x)\, m(dx) - F(m), \\ U_T(m) = G(m). \end{cases}$$

where D_m denotes the intrinsic derivative/Lions derivative defined as follows,

$$D_m U_t(m, x) := \nabla_x \delta_m U_t(m, x).$$

So the analytical description of the N-particle optimal system is

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) - NF(\mu_{\mathbf{x}}^N), \qquad U_T^N = NG(\mu_{\mathbf{x}}^N), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot \left(m_t^N \nabla_p h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) \right), \quad m_0^N = \mathcal{L}(\mathbf{X}_0). \end{cases}$$

Mean-field limit:

The optimal mean-field flow satisfies the Mckean-Vlasov equation:

$$\begin{cases} \partial_t m_t = \Delta m_t + \mathrm{div} \left(m_t \nabla_p h(D_m U_t(m_t, \cdot), \cdot) \right), \\ m_{t=0} = m_0. \end{cases}$$

So the analytical description of the N-particle optimal system is

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) - NF(\mu_{\mathbf{x}}^N), \qquad U_T^N = NG(\mu_{\mathbf{x}}^N), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot \left(m_t^N \nabla_p h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) \right), \quad m_0^N = \mathcal{L}(\mathbf{X}_0). \end{cases}$$

Mean-field limit:

Thus the mean-field control system is

$$\begin{cases} \partial_t U_t(m) &= -\int_{\mathcal{X}} \Delta_x \delta_m U_t(m, x)\, m(dx) + \int_{\mathcal{X}} h(D_m U_t(m, x), x)\, m(dx) - F(m) \quad , U_T = G, \\ \partial_t m_t &= \Delta m_t + \operatorname{div}\left(m_t \nabla_p h(D_m U_t(m_t, \cdot), \cdot) \right) \quad , m_0 = \hat{m}. \end{cases}$$

So the analytical description of the N-particle optimal system is

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) - NF(\mu_{\mathbf{x}}^N), \qquad U_T^N = NG(\mu_{\mathbf{x}}^N), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot \left(m_t^N \nabla_p h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) \right), \quad m_0^N = \mathcal{L}(\mathbf{X}_0). \end{cases}$$

Mean-field limit:

Thus the mean-field control system is

$$\begin{cases} \partial_t U_t(m) &= -\int_{\mathcal{X}} \Delta_x \delta_m U_t(m,x)\, m(dx) + \int_{\mathcal{X}} h(D_m U_t(m,x),x)\, m(dx) - F(m) \quad , U_T = G, \\ \partial_t m_t &= \Delta m_t + \mathrm{div}\left(m_t \nabla_p h(D_m U_t(m_t,\cdot),\cdot) \right) \quad , m_0 = \hat{m}. \end{cases}$$

The idea of propagation of chaos (POC) in MFC is to justify the passage from the N-particle optimal control problem to its mean-field limit:

$$\frac{1}{N} U_t^N(x_1,\dots,x_N) \longrightarrow U_t(\mu) \qquad \text{whenever } \mu_{\mathbf{x}}^N \to \mu.$$

Moreover, if $m_t^N = \mathcal{L}(\boldsymbol{X}_t)$ is the law of the optimal particle system, then

$$m_t^{N:k} \longrightarrow m_t^{\otimes k}, \qquad \text{for every fixed } k \geq 1,$$

or equivalently,

$$\mu_{\boldsymbol{X}_t}^N \longrightarrow m_t \qquad \text{as } N \to \infty.$$

So the analytical description of the N-particle optimal system is

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) - NF(\mu_{\mathbf{x}}^N), \qquad U_T^N = NG(\mu_{\mathbf{x}}^N), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot \left(m_t^N \nabla_p h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) \right), \quad m_0^N = \mathcal{L}(\mathbf{X}_0) = \hat{m}^{\otimes N}. \end{cases}$$

Mean-field limit:

Thus the mean-field control system is

$$\begin{cases} \partial_t U_t(m) &= -\int_{\mathcal{X}} \Delta_x \delta_m U_t(m,x)\, m(dx) + \int_{\mathcal{X}} h(D_m U_t(m,x), x)\, m(dx) - F(m) \quad , U_T = G, \\ \partial_t m_t &= \Delta m_t + \operatorname{div}\left(m_t \nabla_p h(D_m U_t(m_t, \cdot), \cdot) \right) \quad , m_0 = \hat{m}. \end{cases}$$

The idea of propagation of chaos (POC) in MFC is to justify the passage from the N-particle optimal control problem to its mean-field limit:

$$\frac{1}{N} U_t^N(x_1, \ldots, x_N) \longrightarrow U_t(\mu) \qquad \text{whenever } \mu_{\mathbf{x}}^N \to \mu.$$

Moreover, if $m_t^N = \mathcal{L}(\boldsymbol{X}_t)$ is the law of the optimal particle system, then

$$m_t^{N:k} \longrightarrow m_t^{\otimes k}, \qquad \text{for every fixed } k \geq 1,$$

or equivalently,

$$\mu_{\boldsymbol{X}_t}^N \longrightarrow m_t \qquad \text{as } N \to \infty.$$

Mean-field limit:

The idea of propagation of chaos (POC) in MFC is to justify the passage from the *N*-particle optimal control problem to its mean-field limit:

$$\frac{1}{N}U_t^N(x_1,\dots,x_N) \longrightarrow U_t(\mu) \qquad \text{whenever } \mu_{\mathbf{x}}^N \to \mu.$$

Moreover, if $m_t^N = \mathcal{L}(\boldsymbol{X}_t)$ is the law of the optimal particle system, then

$$m_t^{N:k} \longrightarrow m_t^{\otimes k}, \qquad \text{for every fixed } k \geq 1,$$

or equivalently,

$$\mu_{\boldsymbol{X}_t}^N \longrightarrow m_t \qquad \text{as } N \to \infty.$$

Note that here the convergence at $t = 0$ relies on the I.I.D. assumption.

- At what speed with respect to N? (Smoothness of F and G)
- At what speed with respect to T? (Uniform POC problem)
- Can the convergence happen without the I.I.D. assumption? (Turnpike-type POC)
- Can the convergence happen without assuming the same terminal cost?

Outline

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) - NF(\mu_{\mathbf{x}}^N), \qquad U_T^N = G^N(\mathbf{x}), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot \left(m_t^N \nabla_p h^N(\nabla_{\mathbf{x}} U_t^N, \mathbf{x}) \right), \quad m_0^N = \hat{m}_0^N. \end{cases}$$

$$\begin{cases} \partial_t U_t(m) &= -\int_{\mathcal{X}} \Delta_x \delta_m U_t(m,x)\, m(dx) + \int_{\mathcal{X}} h(D_m U_t(m,x), x)\, m(dx) - F(m) \quad , U_T = G, \\ \partial_t m_t &= \Delta m_t + \mathrm{div}\left(m_t \nabla_p h(D_m U_t(m_t, \cdot), \cdot) \right) \quad , m_0 = \hat{m}. \end{cases}$$

Main Result: Turnpike-type POC of MFC

Under suitable convexity and regularity assumptions, the optimal particle system satisfies

$$\mathbb{E}\left[\mathbf{d}_2(\mu_{\mathbf{X}_t}^N, m_t) \right] \le C_1 \cdot \varepsilon_N + C_2 \cdot \left(e^{-\lambda t} + e^{-\lambda(T-t)} \right), \qquad t \in [0, T],$$

where λ depends on F and d, C_1 depends on d, and C_2 depends on $H(\hat{m}_0^N | \hat{m}_0^{\otimes N}) + \|G - G^N\|$.

Here, H denotes the relative entropy and ε_N^2 is the usual empirical-measure approximation error:

$$\varepsilon_N^2 \simeq \begin{cases} N^{-1}, & d = 1, \\ \frac{\log N}{N}, & d = 2, \\ N^{-2/d}, & d \ge 3. \end{cases}$$

Proof of the Main Result

True by the axiom of *It is obvious, apparently.*

Thank you!

- Toy Model;
- Entropy and Fisher Information.

Proof of the Main Result

$$\begin{cases} \partial_t U_t^N &= -\Delta_{\mathbf{x}} U_t^N + \frac{1}{2}|\nabla_{\mathbf{x}} U_t^N|^2 - NF(\mu_{\mathbf{x}}^N), \quad U_T^N = NG(\mu_{\mathbf{x}}^N), \\ \partial_t m_t^N &= \Delta_{\mathbf{x}} m_t^N + \nabla_{\mathbf{x}} \cdot (m_t^N \nabla_{\mathbf{x}} U_t^N)\,, \qquad m_{t_0}^N = m_0^N. \end{cases}$$

$$\begin{cases} \partial_t U_t(m) &= -\int_{\mathcal{X}} \Delta_x \delta_m U_t(m,x)\, m(dx) + \int_{\mathcal{X}} \frac{1}{2}|D_m U_t(m,x)|^2\, m(dx) - F(m) \quad , U_T = G, \\ \partial_t m_t &= \Delta m_t + \mathrm{div}\,(m_t D_m U_t(m_t,\cdot))) \quad , m_{t_0} = m_0. \end{cases}$$

Thank you!