Mean-Field Optimal Control

M. Fornasier, F. Solombrino, ESAIM: COCV, 2013

Gianluca Covini

Presentation for the exam of *Statistical Mechanics and Methods for Complex Systems*

Table of contents

Introduction

Mean Field Theory

Mean Field Theory

Study the collective behavior of multi-agent systems.

In particular, *state evolution* of particle-like systems with social interaction forces.

Mean Field Theory

Study the collective behavior of multi-agent systems.

In particular, *state evolution* of particle-like systems with social interaction forces.

Applications

- Particles with position and velocity which interact through fundamental forces;

Mean Field Theory

Study the collective behavior of multi-agent systems.

In particular, *state evolution* of particle-like systems with social interaction forces.

Applications

- Particles with position and velocity which interact through fundamental forces;
- animal collective behavior: emergence of flocks;

Mean Field Theory

Study the collective behavior of multi-agent systems.

In particular, *state evolution* of particle-like systems with social interaction forces.

Applications

- Particles with position and velocity which interact through fundamental forces;
- animal collective behavior: emergence of flocks;
- opinion dynamics: emergence of consensus.

Mean Field Theory

Mean Field Theory

Study the collective behavior of multi-agent systems.

In particular, *state evolution* of particle-like systems with social interaction forces.

Applications

- Particles with position and velocity which interact through fundamental forces;
- animal collective behavior: emergence of flocks;
- opinion dynamics: emergence of consensus.

Goal

Study the emergence of global behavior under the mean-field limit, i.e. for the number of particles/individuals $N \to \infty$.

An example

Cucker-Smale Model

The CS model deals with flocking (or consensus) emergence

$$(CS) \quad \begin{cases} \dot{x}_i &= v_i \\ \dot{v}_i &= \frac{1}{N}\sum_{j=1}^{N} a(|x_j - x_i|)(v_j - v_i) \end{cases} \qquad i = 1, \ldots, N, \quad t \in [0, T]$$

An example

Cucker-Smale Model

The CS model deals with flocking (or consensus) emergence

$$(CS) \quad \begin{cases} \dot{x}_i &= v_i \\ \dot{v}_i &= \frac{1}{N}\sum_{j=1}^{N} a(|x_j - x_i|)(v_j - v_i) \end{cases} \qquad i = 1, \dots, N, \quad t \in [0, T]$$

Question

Can a policy-maker intervene towards pattern formation?

CS with control

Cucker-Smale Model with Control

Given $f_N : [0, T] \times \mathbb{R}^{2d} \to \mathbb{R}^d$

$$(\text{CS}) \quad \begin{cases} \dot{x}_i = v_i \\ \dot{v}_i = \dfrac{1}{N} \displaystyle\sum_{j=1}^{N} a(|x_j - x_i|)(v_j - v_i) + f_N(t, x_i, v_i) \end{cases} \quad \begin{array}{l} i = 1, \dots, N, \\ t \in [0, T] \end{array}$$

CS with control

Cucker-Smale Model with Control

Given $f_N : [0,T] \times \mathbb{R}^{2d} \to \mathbb{R}^d$

$$(CS) \quad \begin{cases} \dot{x}_i = v_i & i = 1, \dots, N, \\ \dot{v}_i = \frac{1}{N} \sum_{j=1}^{N} a(|x_j - x_i|)(v_j - v_i) + f_N(t, x_i, v_i) & t \in [0,T] \end{cases}$$

Control Cost

$$\int_0^T \frac{1}{N} \sum_{i=1}^{N} (|v_i - \frac{1}{N} \sum_{j=1}^{N} v_j|^2 + \gamma |f_N(t, x_i, v_i)|) dt, \quad \gamma > 0$$

- The first term gives cost to the discrepancy to the flocking;
- the second term entails sparsity of the control.

Finite-Dimensional Problem

Optimization Problem

We represent the system as $\mu_N(t) = \frac{1}{N}\sum_{i=1}^{N}\delta_{(x_i(t),v_i(t))}$, with $(x_i, v_i) = (x_i(t), v_i(t)) \in \mathbb{R}^{2d}$.

General Finite-dimensional Problem

$$\min_{f_N \in \mathcal{F}_\ell} \quad \int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, \mu_N) + \psi(f_N(t, x, v))d\mu_N(x, v)dt$$

$$\text{s.t.} \quad \begin{cases} \dot{x}_i = v_i & i = 1, \dots, N, \\ \dot{v}_i = (H \star \mu_N)(x_i, v_i) + f_N(t, x_i, v_i) & t \in [0, T] \end{cases}$$

$$\text{where } (H \star \mu)(x, v) = \int_{\mathbb{R}^{2d}} H((x, v) - (\xi, \nu))d\mu(\xi, \nu)$$

$$\text{if } \mu = \mu_N, \text{ then} \qquad = \frac{1}{N}\sum_{j=1}^{N} H((x, y) - (x_j, v_j))$$

Optimization Problem

General Finite-dimensional Problem

$$\min_{f_N \in \mathcal{F}_\ell} \quad \int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, \mu_N) + \psi(f_N(t, x, v)) d\mu_N(x, v) dt$$

$$\text{s.t.} \quad \begin{cases} \dot{x}_i = v_i & i = 1, \dots, N, \\ \dot{v}_i = (H \star \mu_N)(x_i, v_i) + f_N(t, x_i, v_i) & t \in [0, T] \end{cases}$$

Admissible Controls

We can imagine $\mathcal{F}_\ell \ni f$ such as

- $f : [0, T] \to W^{1,p}(\Omega \subseteq \mathbb{R}^{2d}, \mathbb{R}^d) \in L^q$;
- $|f(t, 0)| + \mathrm{Lip}(f(t, \cdot), \mathbb{R}^d) \leq \ell(t) \in L^q$ for a.e. $t \in [0, T]$.

Note: (Under assumptions on L, ψ, H) the problem always has a solution.

Questions

Question

What happens under the mean-field limit ($N \to \infty$)?

In particular,

1. Which pattern does the system reach? $\mu_N \to \mu_\infty$?
2. Which is the right control for $N \to \infty$? $f_N \to f_\infty$?
3. What is the relation between μ_∞ and f_∞?

Mean-Field Limit

Limit for μ_N

Question

$\mu_N(t) \in \mathcal{P}_1(\mathbb{R}^{2d}) = \{\pi$ probability measure with $\int_{\mathbb{R}^{2d}} |x| d\pi(x) < \infty\}$.

Which topology on $\mathcal{P}_1(\mathbb{R}^{2d})$?

Limit for μ_N

Question

$\mu_N(t) \in \mathcal{P}_1(\mathbb{R}^{2d}) = \{\pi \text{ probability measure with } \int_{\mathbb{R}^{2d}} |x| d\pi(x) < \infty\}$.

Which topology on $\mathcal{P}_1(\mathbb{R}^{2d})$?

Wasserstein Distance

$$\mathcal{W}_1(\mu, \nu) := \min_{\gamma \in \Gamma(\mu,\nu)} \int_{\mathbb{R}^{2d}} |(\xi, s) - (\eta, u)| d\gamma((\xi, s), (\eta, u))$$

where $\Gamma(\mu, \nu) := \{\gamma \in \mathcal{P}_1(\mathbb{R}^{2d} \times \mathbb{R}^{2d}) : \pi^X_\# \gamma = \mu, \pi^Y_\# \gamma = \nu\}$.

Limit for μ_N

Question

$\mu_N(t) \in \mathcal{P}_1(\mathbb{R}^{2d}) = \{\pi \text{ probability measure with } \int_{\mathbb{R}^{2d}} |x| d\pi(x) < \infty\}$.

Which topology on $\mathcal{P}_1(\mathbb{R}^{2d})$?

Wasserstein Distance

$$\mathcal{W}_1(\mu, \nu) := \min_{\gamma \in \Gamma(\mu,\nu)} \int_{\mathbb{R}^{2d}} |(\xi, s) - (\eta, u)| d\gamma((\xi, s), (\eta, u))$$

where $\Gamma(\mu, \nu) := \{\gamma \in \mathcal{P}_1(\mathbb{R}^{2d} \times \mathbb{R}^{2d}) : \pi^X_\# \gamma = \mu, \pi^Y_\# \gamma = \nu\}$.

Idea

Find μ_∞ such that

$$\mathcal{W}_1(\mu_N(t), \mu_\infty(t)) \to 0 \iff \begin{array}{c} \mu_N(t) \rightharpoonup \mu_\infty(t) \\ \displaystyle\int_{\mathbb{R}^{2d}} |x| d\mu_N(x) \to \int_{\mathbb{R}^{2d}} |x| d\mu(x) \end{array}$$

Limit for f_N

Question

$$f_N \in L^q((0,T), W^{1,p}(\Omega \subseteq \mathbb{R}^{2d}, \mathbb{R}^d))$$

Which notion of convergence?

Limit for f_N

Question

$$f_N \in L^q((0,T), W^{1,p}(\Omega \subseteq \mathbb{R}^{2d}, \mathbb{R}^d))$$

Which notion of convergence?

Weak Convergence

$f_N \rightharpoonup f_\infty$ if and only if

$$\int_0^T \langle \phi(t), f_N(t,\cdot) - f_\infty(t,\cdot)\rangle dt \to 0 \quad \forall \phi \in L^{q'}((0,T), H^{-1,p}(\mathbb{R}^{2d}, \mathbb{R}^d))$$

(We actually consider only ϕ with compact support in Ω.)

Infinite-Dimensional Problem

Question

Which is the relation between μ_∞ and f_∞ we are looking for?

Infinite-Dimensional Problem

Question

Which is the relation between μ_∞ and f_∞ we are looking for?

Main Result

(Under assumptions of continuity of L and lipschitzianity on ψ and H) the (μ_N, f_N) solution of the finite-dimensional problem converge *up to a subsequence* to (μ_∞, f_∞), solutions of the infinite-dimensional problem

$$\min_{f_\infty \in \mathcal{F}_\ell} \quad \int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, \mu_\infty) + \psi(f_\infty(t, x, v))d\mu_\infty(x, v)dt$$

$$\text{s.t.} \quad \frac{\partial \mu_\infty}{\partial t} + v \cdot \nabla_x \mu_\infty = \nabla_v \cdot [(H \star \mu_\infty + f_\infty)\mu_\infty] \quad \forall t \in [0, T]$$

Note: μ_∞ is a weak solution of the PDE in the sense of distributions.

Comparison of the Problems

Finite-Dimensional Problem

$$\min_{f_N \in \mathcal{F}_\ell} \int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, \mu_N) + \psi(f_N(t, x, v)) d\mu_N(x, v) dt$$

$$\text{s.t.} \quad \begin{cases} \dot{x}_i = v_i & i = 1, \ldots, N, \\ \dot{v}_i = (H \star \mu_N)(x_i, v_i) + f_N(t, x_i, v_i) & t \in [0, T] \end{cases}$$

Infinite-Dimensional Problem

$$\min_{f_\infty \in \mathcal{F}_\ell} \int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, \mu_\infty) + \psi(f_\infty(t, x, v)) d\mu_\infty(x, v) dt$$

$$\text{s.t.} \quad \frac{\partial \mu_\infty}{\partial t} + v \cdot \nabla_x \mu_\infty = \nabla_v \cdot [(H \star \mu_\infty + f_\infty) \mu_\infty] \quad t \in [0, T]$$

Sketch of the Proof

The proof develops the limit for μ_N and the one for f_N together.

How to find μ_∞

- As solution of the ODE system, μ_N has regularity properties; in particular, it is equi-bounded and equi-Lipschitz.
- We can apply Ascoli-Arzelà theorem, and say that, up to a subsequence, $\exists \mu^{(lim)}$ such that $\mu_N \to \mu^{(lim)}$ with respect to $\mathcal{W}_1$.

Sketch of the Proof

The proof develops the limit for μ_N and the one for f_N together.

How to find μ_∞

- As solution of the ODE system, μ_N has regularity properties; in particular, it is equi-bounded and equi-Lipschitz.
- We can apply Ascoli-Arzelà theorem, and say that, up to a subsequence, $\exists \mu^{(lim)}$ such that $\mu_N \to \mu^{(lim)}$ with respect to $\mathcal{W}_1$.

How to find f_∞

- The space of admissible controls $\mathcal{F}_\ell$ can be shown to be weakly compact.
- Therefore, $\exists f^{(lim)} \in \mathcal{F}_\ell$ such that, up to a subsequence, $f_N \rightharpoonup f^{(lim)}$.

Sketch of the Proof

Solution of the Infinite-Dimensional PDE

- It can be shown that $\mu^{(lim)}$ solves the infinite-dimensional PDE by
  - showing that each μ_N is weak solution of the PDE in the sense of distributions, in particular $\forall \zeta \in C_c^\infty$,

$$\int_{\mathbb{R}^{2d}} \zeta(x,v) d\mu_N(t)(x,v) - \int_{\mathbb{R}^{2d}} \zeta(x,v) d\mu_N(0)(x,v) =$$
$$= \int_{\mathbb{R}^{2d}} (\nabla_x \zeta(x,v) \cdot v + \nabla_v \zeta(x,v) \cdot (H \star \mu_N(t))(x,v) +$$
$$+ \nabla_v \zeta(x,v) \cdot f_N(t,x,v)) d\mu_N(t)(x,v)$$

  - and then computing the limit for $N \to \infty$.

Optimality of $f^{(lim)}$

Lastly, one has to show that $f^{(lim)}$ is optimal with respect to the infinite-dimensional problem. This limit of optimization problems is called Γ-limit. In particular, given $g \in \mathcal{F}_\ell$, we have that

$$\begin{aligned}
&\int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, \mu^{(lim)}) + \psi(f^{(lim)}(t, x, v)) d\mu^{(lim)}(t, x, v) dt \\
&\leq \liminf_{N\to\infty} \int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, \mu_N) + \psi(f_N(t, x, v)) d\mu_N(t, x, v) dt \\
&\leq \liminf_{N\to\infty} \int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, (\mu_g)_N) + \psi(g(t, x, v)) d(\mu_g)_N(t, x, v) dt \\
&= \int_0^T \int_{\mathbb{R}^{2d}} (L(x, v, \mu_g) + \psi(g(t, x, v)) d\mu_g(t, x, v) dt
\end{aligned}$$

Sketch of the Proof

Optimality of f_∞

$$
\begin{aligned}
&\int_0^T \int_{\mathbb{R}^{2d}} (L(x,v,\mu^{(lim)}) + \psi(f^{(lim)}(t,x,v)) d\mu^{(lim)}(t,x,v) dt \\
&\leq \liminf_{N\to\infty} \int_0^T \int_{\mathbb{R}^{2d}} (L(x,v,\mu_N) + \psi(f_N(t,x,v)) d\mu_N(t,x,v) dt \\
&\leq \liminf_{N\to\infty} \int_0^T \int_{\mathbb{R}^{2d}} (L(x,v,(\mu_g)_N) + \psi(g(t,x,v)) d(\mu_g)_N(t,x,v) dt \\
&= \int_0^T \int_{\mathbb{R}^{2d}} (L(x,v,\mu_g) + \psi(g(t,x,v)) d\mu_g(t,x,v) dt
\end{aligned}
$$

- The first inequality is consequence of properties of semicontinuity of the controls and dominated convergence;
- the second inequality is due to optimality of f_N;
- the last equality is due to dominated convergence.

Sketch of the Proof

Conclusion

We can therefore conclude that $\exists(\mu_\infty, f_\infty)$ and, in particular, $\mu_\infty := \mu^{(lim)}$ and $f_\infty := f^{(lim)}$.

Conclusion

Conclusion

Take-home

(Under regularity assumptions) an optimal control problem with "mean-field" ODE contraints of the form we defined converges weakly and up to a subsequence under mean-field limit to an optimal control problem with McKean-Vlasov PDE contstraint.

Conclusion

Take-home

(Under regularity assumptions) an optimal control problem with "mean-field" ODE contraints of the form we defined converges weakly and up to a subsequence under mean-field limit to an optimal control problem with McKean-Vlasov PDE contstraint.

Other Observations

We had no time to focus on:

- The fact that the finite-dimensional problem (and therefore the infinite-dimensional) has a solution for every admissible control;
- the assumptions on the functions describing the system (which may be considered too strong, in particular the linear bound on interaction term H);
- the fact that the convergence is only weak and up to a subsequence.

Conclusion

Strong and weak point of the paper

- It proves an important existence result;
- it does not provide characterisation of the mean-field control neither numerical ways to approximate it.

Conclusion

Strong and weak point of the paper

- It proves an important existence result;
- it does not provide characterisation of the mean-field control neither numerical ways to approximate it.

Consequences

The results can be used to prove consistency to a wide class of problems of practical interest under the mean-field limit, for example the flocking emergence in Cucker-Smale model we cited at the beginning.

Conclusion

Strong and weak point of the paper

- It proves an important existence result;
- it does not provide characterisation of the mean-field control neither numerical ways to approximate it.

Consequences

The results can be used to prove consistency to a wide class of problems of practical interest under the mean-field limit, for example the flocking emergence in Cucker-Smale model we cited at the beginning.

Thank you for your attention!