Friedrich-Alexander-Universität
Erlangen-Nürnberg

Friedrich-Alexander-Universität
DYNAMICS, CONTROL,
MACHINE LEARNING
AND NUMERICS

How to steer a fleet of agents by controlling only a few of them

Enrico Sartor

14.11.24

Outline

Introduction to the problem

Some motivating examples

Some motivating examples

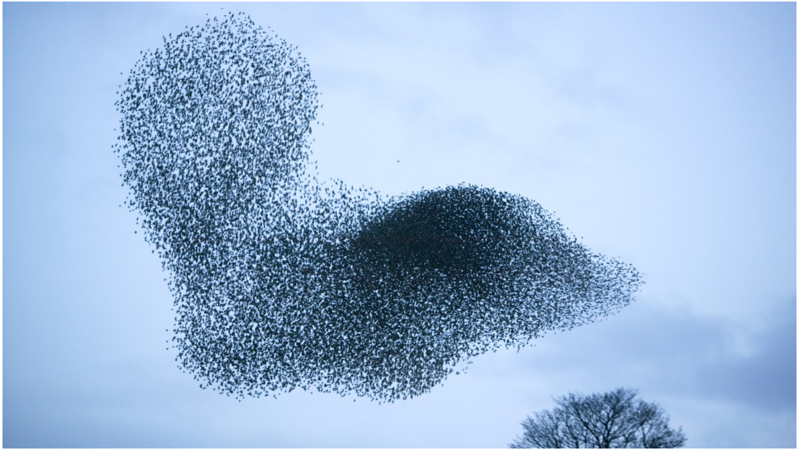

Some motivating examples

Some motivating examples

A possible model

A possible model

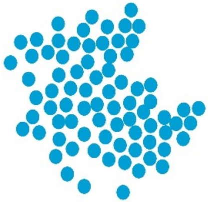

A possible model

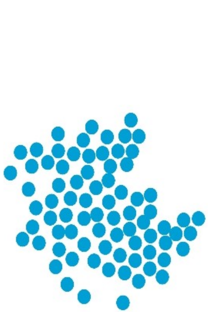

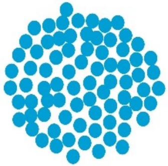

A possible model

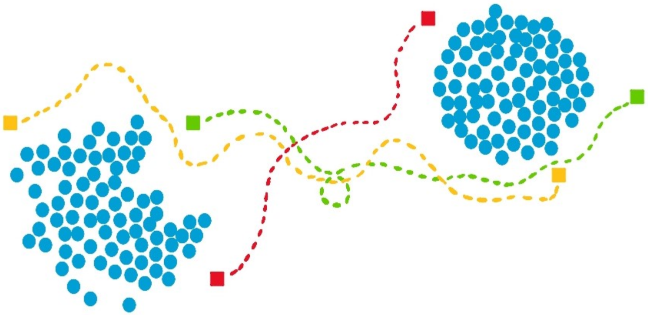

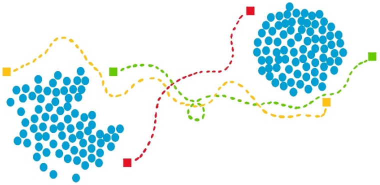

$$
\begin{cases}
\dot{x}_{N,k}(t) = \dfrac{1}{N}\displaystyle\sum_{i=1}^{N} K(x_{N,i}(t) - x_{N,k}(t)) + \dfrac{1}{M}\displaystyle\sum_{j=1}^{M} f(y_j(t) - x_{N,k}(t)) \\
\dot{y}_m(t) = \dfrac{1}{N}\displaystyle\sum_{i=1}^{N} g(x_{N,i}(t) - y_m(t)) + u_m(t)
\end{cases}
$$

A possible model

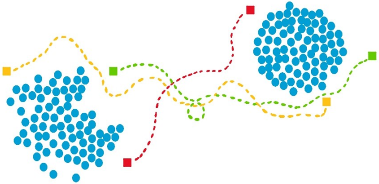

$$\inf_{\boldsymbol{u}} \psi\left(x_{N,1}(T), ..., x_{N,N}(T)\right)$$

The main issues of this kind of models

The main issues of this kind of models

It is sensitive to changes in the number of agents.

The main issues of this kind of models

It is sensitive to changes in the number of agents.

As N grows, solving optimal control problems associated to this models becomes **computationally unfeasible**.

The main issues of this kind of models

It is sensitive to changes in the number of agents.

As N grows, solving optimal control problems associated to this models becomes **computationally unfeasible**.

Non-controllable agents are **indistinguishable** and thus specifying a goal by means of a cost functional to be optimized is not always possible.

The main issues of this kind of models

It is sensitive to changes in the number of agents.

As N grows, solving optimal control problems associated to this models becomes **computationally unfeasible**.

Non-controllable agents are **indistinguishable** and thus specifying a goal by means of a cost functional to be optimized is not always possible.

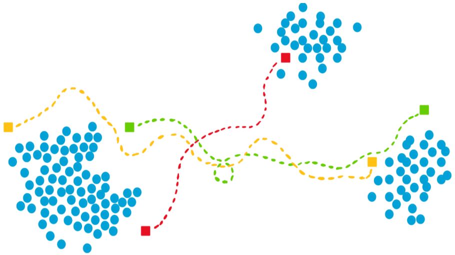

The mean field approach

The mean field approach

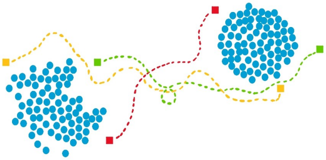

The mean field approach

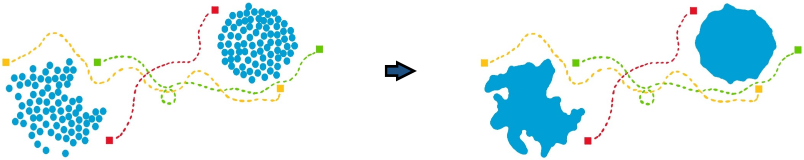

The mean field approach

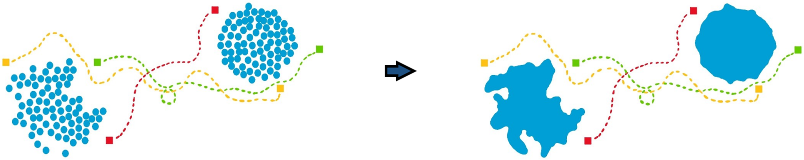

$$\begin{cases} \dot{x}_{N,k}(t) = \hat{F}_N(x_N(t), y(t)) \\ \dot{y}_m(t) = \hat{G}_N(x_N(t), y_m(t)) + u_m(t) \end{cases}$$

The mean field approach

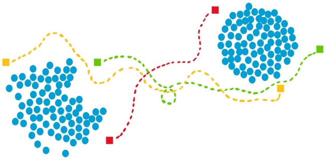

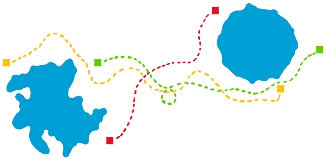

$$\begin{cases} \dot{x}_{N,k}(t) = \hat{F}_N(x_N(t), y(t)) \\ \dot{y}_m(t) = \hat{G}_N(x_N(t), y_m(t)) + u_m(t) \end{cases}$$

$$\begin{cases} \partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](y(t), \cdot)\mu(t)\big) = 0 \\ \dot{y}(t) = G[\mu(t)](y(t)) + u(t) \end{cases}$$

The mean field approach

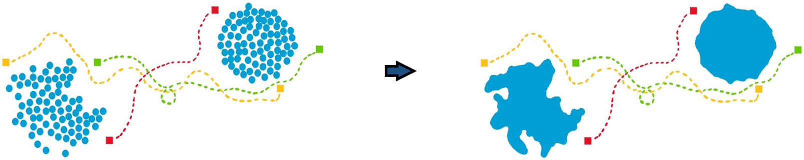

$$\begin{cases} \dot{\boldsymbol{x}}_{N,k}(t) = \hat{\boldsymbol{F}}_N(\boldsymbol{x}_N(t), \boldsymbol{y}(t)) \\ \dot{\boldsymbol{y}}_m(t) = \hat{\boldsymbol{G}}_N(\boldsymbol{x}_N(t), \boldsymbol{y}_m(t)) + \boldsymbol{u}_m(t) \end{cases}$$

$$\begin{cases} \partial_t \boldsymbol{\mu}(t) + \nabla_{\boldsymbol{x}} \cdot \big(\boldsymbol{F}[\boldsymbol{\mu}(t)](\boldsymbol{y}(t), \cdot)\boldsymbol{\mu}(t)\big) = \boldsymbol{0} \\ \dot{\boldsymbol{y}}(t) = \boldsymbol{G}[\boldsymbol{\mu}(t)](\boldsymbol{y}(t)) + \boldsymbol{u}(t) \end{cases}$$

$$F[\mu](y,x) = \int_{\mathbb{R}^d} K(\zeta - x)d\mu(\zeta) + \frac{1}{M}\sum_{m=1}^{M} f(y_m - x) \qquad G[\mu](y) = \int_{\mathbb{R}^d} g(\zeta - y)d\mu(\zeta).$$

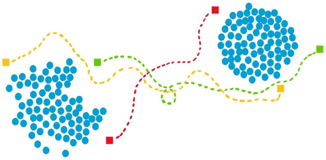

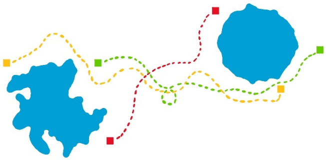

$$\begin{cases} \dot{x}_{N,k}(t) = \hat{F}_N(x_N(t), y(t)) \\ \dot{y}_m(t) = \hat{G}_N(x_N(t), y_m(t)) + u_m(t) \end{cases}$$

$$\begin{cases} \partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](y(t), \cdot)\mu(t)\big) = 0 \\ \dot{y}(t) = G[\mu(t)](y(t)) + u(t) \end{cases}$$

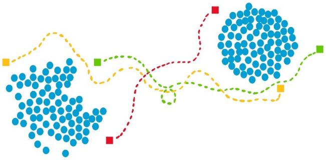

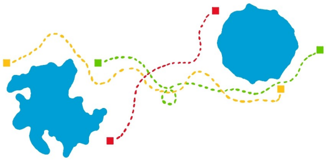

$$\begin{cases} \dot{x}_{N,k}(t) = \hat{F}_N(x_N(t), y(t)) \\ \dot{y}_m(t) = \hat{G}_N(x_N(t), y_m(t)) + u_m(t) \end{cases}$$

$$\begin{cases} \partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](y(t), \cdot)\mu(t)\big) = 0 \\ \dot{y}(t) = G[\mu(t)](y(t)) + u(t) \end{cases}$$

$$\psi\big(x_{N,1}(T), ..., x_{N,N}(T)\big)$$

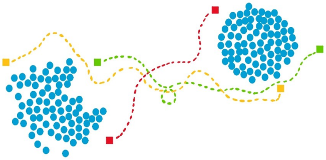

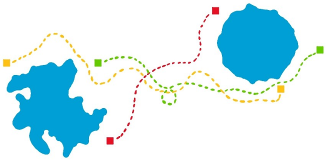

$$\begin{cases} \dot{x}_{N,k}(t) = \hat{F}_N(x_N(t), y(t)) \\ \dot{y}_m(t) = \hat{G}_N(x_N(t), y_m(t)) + u_m(t) \end{cases} \Rightarrow \begin{cases} \partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](y(t), \cdot)\mu(t)\big) = 0 \\ \dot{y}(t) = G[\mu(t)](y(t)) + u(t) \end{cases}$$

$$\psi\big(x_{N,1}(T), ..., x_{N,N}(T)\big) \Rightarrow \psi(\mu(T))$$

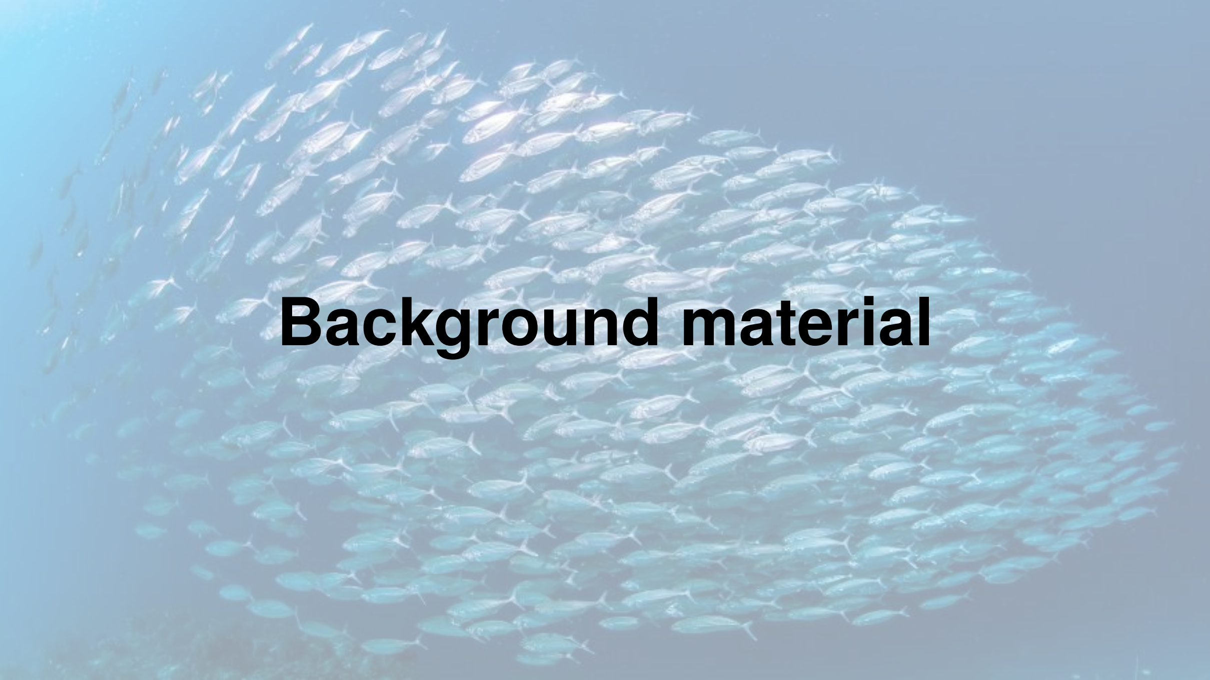

Background material

$\mathcal{P}_2(\mathbb{R}^n)$ and its differentiable structure

$$\mathcal{P}_2(\mathbb{R}^n) := \left\{ \mu \in \mathcal{P}(\mathbb{R}^n) \colon \int_{\mathbb{R}^n} \|x\|^2 d\mu(x) < +\infty \right\}$$

$\mathcal{P}_2(\mathbb{R}^n)$ and its differentiable structure

$$\mathcal{P}_2(\mathbb{R}^n) := \left\{ \mu \in \mathcal{P}(\mathbb{R}^n) : \int_{\mathbb{R}^n} \|x\|^2 d\mu(x) < +\infty \right\}$$

$$W_2(\mu, \nu) := \inf_{\gamma \in \Gamma(\mu,\nu)} \left(\left\{ \int_{\mathbb{R}^n \times \mathbb{R}^n} \|x - y\|^2 d\gamma(x, y) \right\} \right)^{\frac{1}{2}}$$

$$\boldsymbol{\mathcal{P}_2(\mathbb{R}^n) := \left\{ \mu \in \mathcal{P}(\mathbb{R}^n) \colon \int_{\mathbb{R}^n} \|x\|^2 d\mu(x) < +\infty \right\}}$$

$$\boldsymbol{W_2(\mu,\nu) := \inf_{\gamma \in \Gamma(\mu,\nu)} \left(\left\{ \int_{\mathbb{R}^n \times \mathbb{R}^n} \|x-y\|^2 d\gamma(x,y) \right\} \right)^{\frac{1}{2}}}$$

Definition

We say that a function $\varphi \colon \mathcal{P}_2(\mathbb{R}^n) \to \mathbb{R}$ is **Wasserstein differentiable** at μ if there exists $\nabla_\mu \varphi[\mu] \in L^2(\mathbb{R}^n, \mathbb{R}^n, \mu)$ such that

$$\varphi(\nu) = \varphi(\mu) + \int_{\mathbb{R}^n \times \mathbb{R}^n} \langle \nabla_\mu \varphi[\mu](x), y - x \rangle d\gamma(x,y) + o\big(W_2(\mu,\nu)\big) \tag{1}$$

for every $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and every optimal coupling $\gamma \in \Gamma_2^o(\mu,\nu)$.

$$\boldsymbol{\mathcal{P}_2(\mathbb{R}^n) := \left\{\mu \in \mathcal{P}(\mathbb{R}^n) \colon \int_{\mathbb{R}^n} \|x\|^2 d\mu(x) < +\infty\right\}}$$

$$\boldsymbol{W_2(\mu,\nu) := \inf_{\gamma\in\Gamma(\mu,\nu)} \left(\left\{\int_{\mathbb{R}^n\times\mathbb{R}^n} \|x-y\|^2 d\gamma(x,y)\right\}\right)^{\frac{1}{2}}}$$

Definition

We say that a function $\varphi\colon \mathcal{P}_2(\mathbb{R}^n) \to \mathbb{R}$ is **Wasserstein differentiable** at μ if there exists $\nabla_\mu\varphi[\mu] \in L^2(\mathbb{R}^n,\mathbb{R}^n,\mu)$ such that

$$\varphi(\nu) = \varphi(\mu) + \int_{\mathbb{R}^n\times\mathbb{R}^n} \langle \nabla_\mu\varphi[\mu](x), y-x\rangle d\gamma(x,y) + o\big(W_2(\mu,\nu)\big) \qquad (1)$$

for every $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and every optimal coupling $\gamma \in \Gamma_2^o(\mu,\nu)$. In that case we say that $\nabla_\mu\varphi[\mu]$ is a **Wasserstein gradient** of φ at μ.

Two examples of differentiable functionals

Two examples of differentiable functionals

Theorem (Expected values)

If $\hat{\psi}\colon \mathbb{R}^n \to \mathbb{R}$ *is sufficiently regular, then the map*

$$\mu \mapsto \psi[\mu] = \int_{\mathbb{R}^d} \hat{\psi}(x) d\mu(x)$$

is everywhere Wasserstein differentiable with constant differential

$$\nabla_\mu \psi[\mu] = \nabla_x \hat{\psi}.$$

Two examples of differentiable functionals

Two examples of differentiable functionals

Theorem (Wasserstein distance)

Let $\hat{\mu} \in \mathcal{P}_2(\mathbb{R}^d)$ be a fixed reference probability measure. Then, if $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ is absolutely continuous with respect to the Lebesgue measure, the map

$$\frac{1}{2} W_2(\cdot, \hat{\mu})^2 : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$$

is Wasserstein differentiable at μ with Wasserstein differential $\mathrm{Id} - T$, where T is the unique optimal transport map between μ and $\hat{\mu}$.

Non-local continuity equations

Definition

Non-local continuity equations

Definition

A **non-local continuity equation** is a partial differential equation of the form

$$\partial_t \mu(t) + \nabla_x \cdot \big(V[\mu(t)]\mu(t)\big) = 0,$$

Non-local continuity equations

Definition

A **non-local continuity equation** is a partial differential equation of the form

$$\partial_t \mu(t) + \nabla_x \cdot \big(V[\mu(t)]\mu(t)\big) = 0,$$

where the vector field $V\colon [0,T] \times \mathbb{R}^n \times \mathcal{P}_c(\mathbb{R}^n) \to \mathbb{R}^n$ depends also on the distribution μ.

Non-local continuity equations

Definition

A **non-local continuity equation** is a partial differential equation of the form

$$\partial_t \mu(t) + \nabla_x \cdot \big(V[\mu(t)]\mu(t)\big) = 0,$$

where the vector field $V\colon [0,T] \times \mathbb{R}^n \times \mathcal{P}_c(\mathbb{R}^n) \to \mathbb{R}^n$ depends also on the distribution μ. Given $\mu_0 \in \mathcal{P}_c(\mathbb{R}^n)$ we say that a continuous curve $\mu\colon [0,T] \to \mathcal{P}_c(\mathbb{R}^n)$ is a: **distributional** or **Eulerian solution** of the corresponding Cauchy problem if $\mu(0) = \mu_0$ and for every $\xi \in C_c^\infty(\mathbb{R}^n)$ it holds

$$\frac{d}{dt}\int_{\mathbb{R}^n} \xi(x) d\mu(t)(x) = \int_{\mathbb{R}^d} \nabla_x \xi(x) \cdot V[\mu(t)](t,x) d\mu(x);$$

The optimal control problem

The notion of solution

The notion of solution

$$\begin{cases} \partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](t, y(t), \cdot)\mu(t)\big) = 0 \\ \dot{y}(t) = G[\mu(t)](t, y(t)) + u(t) \end{cases}$$

The notion of solution

$$\begin{cases} \partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](t, y(t), \cdot)\mu(t)\big) = 0 \\ \dot{y}(t) = G[\mu(t)](t, y(t)) + u(t) \end{cases}$$

Definition

The notion of solution

$$\begin{cases} \partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](t, y(t), \cdot) \mu(t) \big) = 0 \\ \dot{y}(t) = G[\mu(t)](t, y(t)) + u(t) \end{cases}$$

Definition

Let $u \in L^1([0,T],U)$, $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$ and $y_0 \in \mathbb{R}^{dM}$ be given.

The notion of solution

$$\begin{cases} \partial_t \boldsymbol{\mu}(t) + \nabla_x \cdot \big(F[\boldsymbol{\mu}(t)](t, \boldsymbol{y}(t), \cdot)\boldsymbol{\mu}(t) \big) = \mathbf{0} \\ \dot{\boldsymbol{y}}(t) = G[\boldsymbol{\mu}(t)](t, \boldsymbol{y}(t)) + \boldsymbol{u}(t) \end{cases}$$

Definition

Let $u \in L^1([0,T],U)$, $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$ and $y_0 \in \mathbb{R}^{dM}$ be given. We say that $t \mapsto (\boldsymbol{\mu}(t), \mathbf{y}(t))$ is a *solution* if the following conditions hold:

- $\boldsymbol{\mu}(0) = \mu_0$ and $\mathbf{y}(0) = y_0$;

The notion of solution

$$\begin{cases} \partial_t \boldsymbol{\mu}(t) + \boldsymbol{\nabla}_x \cdot \big(F[\boldsymbol{\mu}(t)](t, \boldsymbol{y}(t), \cdot)\boldsymbol{\mu}(t)\big) = \mathbf{0} \\ \dot{\boldsymbol{y}}(t) = G[\boldsymbol{\mu}(t)](t, \boldsymbol{y}(t)) + \boldsymbol{u}(t) \end{cases}$$

Definition

Let $u \in L^1([0,T],U)$, $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$ and $y_0 \in \mathbb{R}^{dM}$ be given. We say that $t \mapsto (\boldsymbol{\mu}(t), \mathbf{y}(t))$ is a *solution* if the following conditions hold:

- $\boldsymbol{\mu}(0) = \mu_0$ and $\mathbf{y}(0) = y_0$;
- $\boldsymbol{\mu} \colon [0,T] \to \mathcal{P}_c(\mathbb{R}^d)$ is a **distributional solution** of the non-local continuity equation

$$\partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](\mathbf{y}(t), \cdot)\mu(t)\big) = 0,$$

The notion of solution

$$\begin{cases} \partial_t \boldsymbol{\mu}(t) + \nabla_x \cdot \big(\boldsymbol{F}[\boldsymbol{\mu}(t)](\boldsymbol{t}, \boldsymbol{y}(t), \cdot)\boldsymbol{\mu}(t) \big) = \mathbf{0} \\ \dot{\boldsymbol{y}}(t) = \boldsymbol{G}[\boldsymbol{\mu}(t)](\boldsymbol{t}, \boldsymbol{y}(t)) + \boldsymbol{u}(t) \end{cases}$$

Definition

Let $u \in L^1([0,T],U)$, $\mu_0 \in \mathcal{P}_c(\mathbb{R}^d)$ and $y_0 \in \mathbb{R}^{dM}$ be given. We say that $t \mapsto (\boldsymbol{\mu}(t), \mathbf{y}(t))$ is a *solution* if the following conditions hold:

- $\boldsymbol{\mu}(0) = \mu_0$ and $\mathbf{y}(0) = y_0$;
- $\boldsymbol{\mu}\colon [0,T] \to \mathcal{P}_c(\mathbb{R}^d)$ is a **distributional solution** of the non-local continuity equation

$$\partial_t \mu(t) + \nabla_x \cdot \big(F[\mu(t)](\mathbf{y}(t), \cdot)\mu(t) \big) = 0,$$

- $\mathbf{y}\colon [0,T] \to \mathbb{R}^{dM}$ is a **Carathéodory solution** of

$$\dot{y}(t) = G[\boldsymbol{\mu}(t)](y(t)) + u(t)$$

Well-posedness

Well-posedness

Theorem

Well-posedness

Theorem

If we assume that:

Well-posedness

Theorem

If we assume that:

- F, G *are Lipschitz continuous on compact sets;*

Well-posedness

Theorem

If we assume that:

- F, G *are Lipschitz continuous on compact sets;*
- *there exists* $C > 0$ *such that*

$$\|F[\mu](t,x,y)\|_d \leq C(1 + \|x\|_d + \|y\|_c + \mathcal{M}_\infty(\mu)) \quad \text{and} \quad \|G[\mu](t,y)\|_c \leq C(1 + \|y\|_c + \mathcal{M}_\infty(\mu))$$

Well-posedness

Theorem

If we assume that:

- F, G *are Lipschitz continuous on compact sets;*
- *there exists* $C > 0$ *such that*

$$\|F[\mu](t,x,y)\|_d \leq C(1 + \|x\|_d + \|y\|_c + \mathcal{M}_\infty(\mu)) \quad \text{and} \quad \|G[\mu](t,y)\|_c \leq C(1 + \|y\|_c + \mathcal{M}_\infty(\mu))$$

then

Well-posedness

Theorem

If we assume that:

- F, G *are Lipschitz continuous on compact sets;*
- *there exists* $C > 0$ *such that*

$$\|F[\mu](t,x,y)\|_d \leq C(1 + \|x\|_d + \|y\|_c + \mathcal{M}_\infty(\mu)) \quad \textit{and} \quad \|G[\mu](t,y)\|_c \leq C(1 + \|y\|_c + \mathcal{M}_\infty(\mu))$$

then

- *every Cauchy problem admits a unique solution;*

Well-posedness

Theorem

If we assume that:

- F, G *are Lipschitz continuous on compact sets;*
- *there exists* $C > 0$ *such that*

$$\|F[\mu](t,x,y)\|_d \leq C(1 + \|x\|_d + \|y\|_c + \mathcal{M}_\infty(\mu)) \quad \text{and} \quad \|G[\mu](t,y)\|_c \leq C(1 + \|y\|_c + \mathcal{M}_\infty(\mu))$$

then

- *every Cauchy problem admits a unique solution;*
- *solutions depend continuously on initial conditions and control laws;*

Well-posedness

Theorem

If we assume that:

- F, G *are Lipschitz continuous on compact sets;*
- *there exists* $C > 0$ *such that*

$$\|F[\mu](t,x,y)\|_d \leq C(1+\|x\|_d+\|y\|_c+\mathcal{M}_\infty(\mu)) \quad \text{and} \quad \|G[\mu](t,y)\|_c \leq C(1+\|y\|_c+\mathcal{M}_\infty(\mu))$$

then

- *every Cauchy problem admits a unique solution;*
- *solutions depend continuously on initial conditions and control laws;*
- *if moreover* F *and* G *are differentiable with respect to each variable with continuous differentials, we have differentiable dependence on initial conditions.*

Existence of optimal solutions

Existence of optimal solutions

$$\mathcal{U} := L^1\big([0,T],U\big)$$

Existence of optimal solutions

$$\mathcal{U} := L^1\big([0,T],U\big)$$

$$\inf_{u\in\mathcal{U}} J(\mu_0, y_0, u) = \inf_{u\in\mathcal{U}} \Big\{ \psi\big(\boldsymbol{\mu}(\mu_0, y_0, u; T)\big) \Big\} \qquad \textbf{(OCP)}$$

Existence of optimal solutions

$$\mathcal{U} := L^1\big([0,T],U\big)$$

$$\inf_{u\in\mathcal{U}} J(\mu_0,y_0,u) = \inf_{u\in\mathcal{U}} \Big\{\psi\big(\boldsymbol{\mu}(\mu_0,y_0,u;T)\big)\Big\} \qquad \textbf{(OCP)}$$

Theorem (Existence of optimal controls)

If $\psi\colon \mathcal{P}(\mathbb{R}^n) \to \mathbb{R}$ is lower semicontinuous, then ***OCP*** *admits a solution.*

A Pontryagin minimum principle

Theorem

Let u^* be an ***optimal control*** *for our coupled system and $(\boldsymbol{\mu}^*, \mathbf{y}^*)$ be the corresponding* ***optimal trajectory****.*

A Pontryagin minimum principle

Theorem

*Let u^* be an **optimal control** for our coupled system and $(\boldsymbol{\mu}^*, \mathbf{y}^*)$ be the corresponding **optimal trajectory**. If ψ is Wasserstein differentiable at $\boldsymbol{\mu}^*(T)$ with essentially bounded Wasserstein gradient,*

A Pontryagin minimum principle

Theorem

*Let u^* be an **optimal control** for our coupled system and $(\boldsymbol{\mu}^*, \mathbf{y}^*)$ be the corresponding **optimal trajectory**. If ψ is Wasserstein differentiable at $\boldsymbol{\mu}^*(T)$ with essentially bounded Wasserstein gradient, then there exist*

$$\mathbf{q}^* : [0,T] \to \mathbb{R}^c \quad \textit{and} \quad \boldsymbol{\nu}^* : [0,T] \to \mathcal{P}_c(\mathbb{R}^d \times \mathbb{R}^d)$$

such that $(\boldsymbol{\nu}^*, \mathbf{y}^*, \mathbf{q}^*)$ solves the non-local adjoint equation on $\mathbb{R}^d \times \mathbb{R}^d$ with boundary conditions

$$\mathbf{y}^*(0) = y_0 \quad \textit{and} \quad \pi^1_* \boldsymbol{\nu}^*(0) = \mu_0, \quad \mathbf{q}^*(T) = 0 \quad \textit{and} \quad \pi^2_* \boldsymbol{\nu}^*(T) = \nabla_\mu \psi[\boldsymbol{\mu}^*(T)]_* \boldsymbol{\mu}^*(T),$$

A Pontryagin minimum principle

Theorem

*Let u^* be an **optimal control** for our coupled system and $(\boldsymbol{\mu}^*, \mathbf{y}^*)$ be the corresponding **optimal trajectory**. If ψ is Wasserstein differentiable at $\boldsymbol{\mu}^*(T)$ with essentially bounded Wasserstein gradient, then there exist*

$$\mathbf{q}^* : [0,T] \to \mathbb{R}^c \quad \textit{and} \quad \boldsymbol{\nu}^* : [0,T] \to \mathcal{P}_c(\mathbb{R}^d \times \mathbb{R}^d)$$

such that $(\boldsymbol{\nu}^*, \mathbf{y}^*, \mathbf{q}^*)$ solves the non-local adjoint equation on $\mathbb{R}^d \times \mathbb{R}^d$ with boundary conditions

$$\mathbf{y}^*(0) = y_0 \quad \textit{and} \quad \pi^1_* \boldsymbol{\nu}^*(0) = \mu_0, \quad \mathbf{q}^*(T) = 0 \quad \textit{and} \quad \pi^2_* \boldsymbol{\nu}^*(T) = \nabla_\mu \psi[\boldsymbol{\mu}^*(T)]_* \boldsymbol{\mu}^*(T),$$

and

$$\pi^1_* \boldsymbol{\nu}^*(t) = \boldsymbol{\mu}^*(t),$$

A Pontryagin minimum principle

Theorem

*Let u^* be an **optimal control** for our coupled system and $(\boldsymbol{\mu}^*, \mathbf{y}^*)$ be the corresponding **optimal trajectory**. If ψ is Wasserstein differentiable at $\boldsymbol{\mu}^*(T)$ with essentially bounded Wasserstein gradient, then there exist*

$$\mathbf{q}^* : [0,T] \to \mathbb{R}^c \quad \text{and} \quad \boldsymbol{\nu}^* : [0,T] \to \mathcal{P}_c(\mathbb{R}^d \times \mathbb{R}^d)$$

such that $(\boldsymbol{\nu}^*, \mathbf{y}^*, \mathbf{q}^*)$ solves the non-local adjoint equation on $\mathbb{R}^d \times \mathbb{R}^d$ with boundary conditions

$$\mathbf{y}^*(0) = y_0 \quad \text{and} \quad \pi^1_* \boldsymbol{\nu}^*(0) = \mu_0, \quad \mathbf{q}^*(T) = 0 \quad \text{and} \quad \pi^2_* \boldsymbol{\nu}^*(T) = \nabla_\mu \psi[\boldsymbol{\mu}^*(T)]_* \boldsymbol{\mu}^*(T),$$

and

$$\pi^1_* \boldsymbol{\nu}^*(t) = \boldsymbol{\mu}^*(t),$$

*and in order that the **optimality condition***

$$\mathbf{q}^*(t) \cdot u^*(t) = \min_{\omega \in U} \mathbf{q}^*(t) \cdot \omega.$$

Proof sketch

Proof sketch

- **needle variations:**

Proof sketch

- **needle variations:**

$$u^\varepsilon(t) = \begin{cases} \omega & \text{if } t \in [\tau - \varepsilon, \tau] \\ u^*(t) & \text{otherwise} \end{cases}$$

Proof sketch

- **needle variations:**

$$u^{\varepsilon}(t) = \begin{cases} \omega & \text{if } t \in [\tau - \varepsilon, \tau] \\ u^*(t) & \text{otherwise} \end{cases}$$

- **optimality conditions:**

Proof sketch

- **needle variations:**

$$u^\varepsilon(t) = \begin{cases} \omega & \text{if } t \in [\tau - \varepsilon, \tau] \\ u^*(t) & \text{otherwise} \end{cases}$$

- **optimality conditions:**

$$\frac{\psi(\boldsymbol{\mu}^\varepsilon(T)) - \psi(\boldsymbol{\mu}^*(T))}{\varepsilon} \geq 0$$

Proof sketch

- **needle variations:**

$$u^\varepsilon(t) = \begin{cases} \omega & \text{if } t \in [\tau - \varepsilon, \tau] \\ u^*(t) & \text{otherwise} \end{cases}$$

- **optimality conditions:**

$$\frac{\psi(\boldsymbol{\mu}^\varepsilon(T)) - \psi(\boldsymbol{\mu}^*(T))}{\varepsilon} \geq 0$$

$$\int_{\mathbb{R}^d} \nabla_\mu \psi[\boldsymbol{\mu}^*(T)](\boldsymbol{\Phi}^*_\tau(T, x)) \cdot \mathbf{w}(T, x) d\boldsymbol{\mu}^*(\tau)(x) \geq 0$$

Proof sketch

- **needle variations:**

$$u^{\varepsilon}(t) = \begin{cases} \omega & \text{if } t \in [\tau - \varepsilon, \tau] \\ u^*(t) & \text{otherwise} \end{cases}$$

- **optimality conditions:**

$$\frac{\psi(\boldsymbol{\mu}^{\varepsilon}(T)) - \psi(\boldsymbol{\mu}^*(T))}{\varepsilon} \geq 0$$

$$\int_{\mathbb{R}^d} \nabla_{\mu}\psi[\boldsymbol{\mu}^*(T)](\boldsymbol{\Phi}^*_{\tau}(T,x)) \cdot \mathbf{w}(T,x) d\boldsymbol{\mu}^*(\tau)(x) \geq 0$$

- **the adjoint equation property:**

Proof sketch

- **needle variations:**

$$u^\varepsilon(t) = \begin{cases} \omega & \text{if } t \in [\tau - \varepsilon, \tau] \\ u^*(t) & \text{otherwise} \end{cases}$$

- **optimality conditions:**

$$\frac{\psi(\boldsymbol{\mu}^\varepsilon(T)) - \psi(\boldsymbol{\mu}^*(T))}{\varepsilon} \geq 0$$

$$\int_{\mathbb{R}^d} \nabla_\mu \psi[\boldsymbol{\mu}^*(T)](\boldsymbol{\Phi}^*_\tau(T,x)) \cdot \mathbf{w}(T,x) d\boldsymbol{\mu}^*(\tau)(x) \geq 0$$

- **the adjoint equation property:**

$$t \mapsto \int_{\mathbb{R}^{2d}} p \cdot \mathbf{w}(t, \boldsymbol{\Phi}^*_\tau(t,x)) d\boldsymbol{\nu}^*(t)(x,p) + \mathbf{q}^*(t) \cdot \mathbf{v}(t),$$

The mean-field interpretation

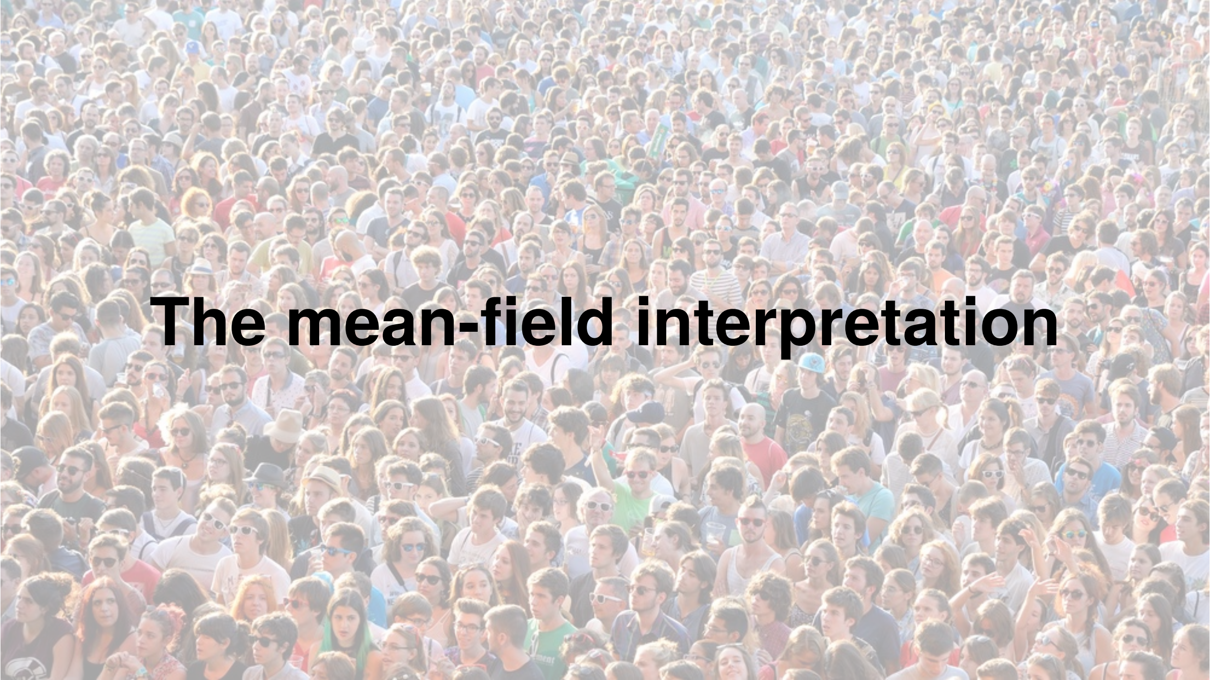

A commutative diagram

A commutative diagram

$$\psi[\mu] := \int_{\mathbb{R}^d} \hat{\psi}(x) d\mu(x),$$

A commutative diagram

$$\psi(\mu_N(T;u)) = \frac{1}{N}\sum_{n=1}^{N}\hat{\psi}(x_{N,n}(T;u)).$$

A commutative diagram

$$\psi(\mu_N(T;u)) = \frac{1}{N}\sum_{n=1}^{N}\hat{\psi}(x_{N,n}(T;u)).$$

finite-dimensional OCP

A commutative diagram

$$\psi(\mu_N(T;u)) = \frac{1}{N}\sum_{n=1}^{N}\hat{\psi}(x_{N,n}(T;u)).$$

$$\textbf{finite-dimensional OCP} \xrightarrow{\textbf{mean-field limit}} \textbf{mean-field OCP}$$

$$\psi(\mu_N(T;u)) = \frac{1}{N}\sum_{n=1}^{N}\hat{\psi}(x_{N,n}(T;u)).$$

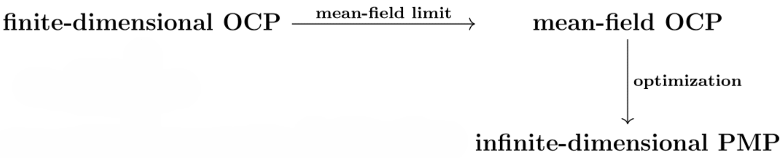

$$\psi(\mu_N(T;u)) = \frac{1}{N}\sum_{n=1}^{N}\hat{\psi}(x_{N,n}(T;u)).$$

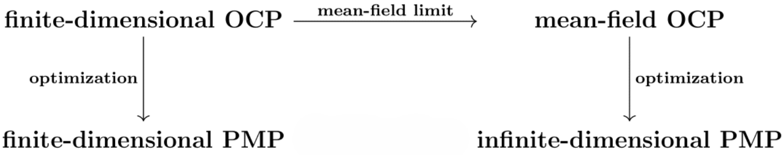

A commutative diagram

$$\psi(\mu_N(T;u)) = \frac{1}{N}\sum_{n=1}^{N}\hat{\psi}(x_{N,n}(T;u)).$$

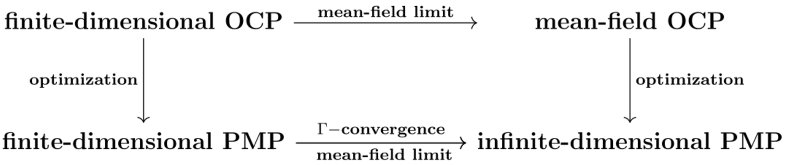

Mean-field regularity

$$\psi(\boldsymbol{\mu}_N(T)) = \inf_{\gamma} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma(x, y) \Big|_{(\mathbf{x}_1(T), \dots, \mathbf{x}_N(T))},$$

Mean-field regularity

$$\psi(\boldsymbol{\mu}_N(T)) = \inf_{\gamma} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma(x, y) \Big|_{(\mathbf{x}_1(T), \ldots, \mathbf{x}_N(T))},$$

with

$$\gamma \in \Gamma\left(\frac{1}{N} \sum_{n=1}^{N} \delta_{x_n}, \hat{\mu}\right).$$

Mean-field regularity

$$\psi(\boldsymbol{\mu}_N(T)) = \inf_{\gamma} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma(x, y) \Big|_{(\mathbf{x}_1(T), \dots, \mathbf{x}_N(T))},$$

with

$$\gamma \in \Gamma\left(\frac{1}{N} \sum_{n=1}^{N} \delta_{x_n}, \hat{\mu}\right).$$

Theorem

If the reference probability measure $\hat{\mu}$ has compact support, u^* is an optimal control for the coupled PDE-ODE system and $(\boldsymbol{\nu}^*, \mathbf{y}*, \mathbf{q}^*)$ is the corresponding optimal trajectory, then

$$\sum_{m=1}^{M} \mathbf{q}^m \cdot u^m(t) = \min_{\omega^m \in U} \sum_{m=1}^{M} \mathbf{q}^m(t) \cdot \omega^m$$

for almost every $t \in [0, T]$.

Performance guarantees

Performance guarantees

Theorem

If there exists $R > 0$ such that $\mathcal{M}_\infty\left(\mu_0^N\right), \mathcal{M}_\infty\left(\mu_0^\infty\right) \leq R$, then there exists a constant $\mathbf{C} > 0$ depending only on T, the interaction kernels and on R such that, if u^* is an optimal solution for the mean-field problem, then

$$J^N(u^*) \leq \inf_{u \in \mathcal{U}} J^\infty(u) + \mathbf{C} W_2(\mu_0^N, \mu_0^\infty)$$

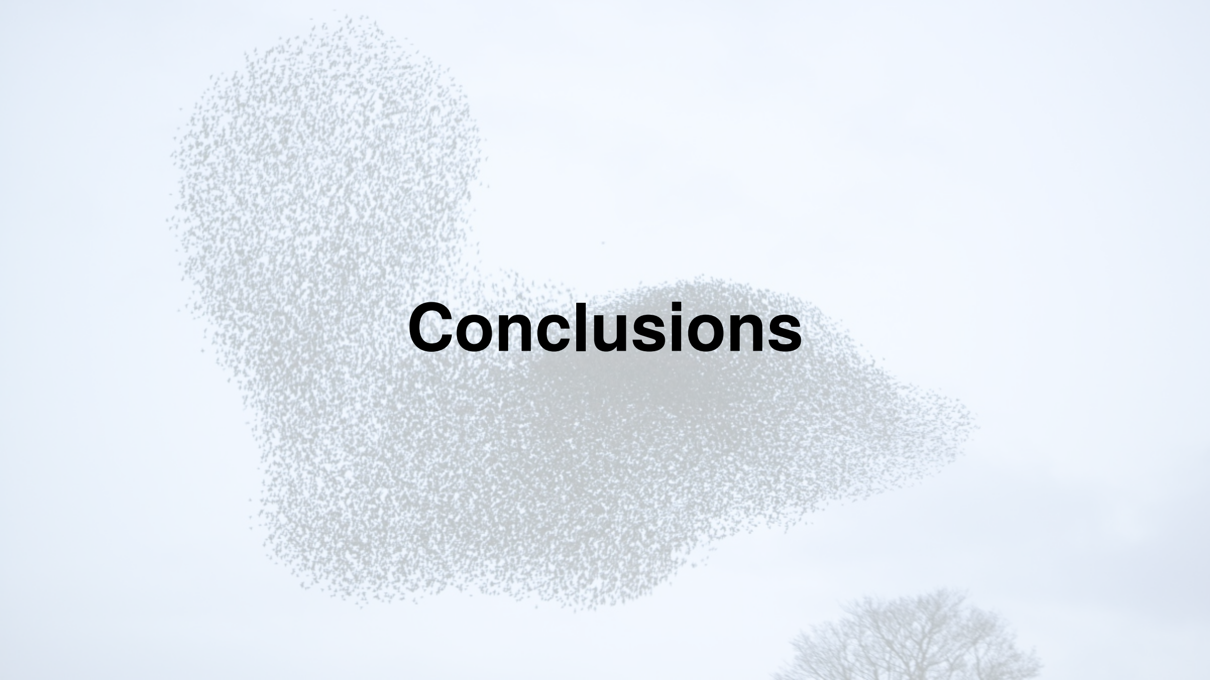

Conclusions

A brief recap

A brief recap

A brief recap

- **THE PROBLEM**: Optimal control of a system made of a large number of interacting agents by controlling only a few of them.

A brief recap

- **THE PROBLEM**: Optimal control of a system made of a large number of interacting agents by controlling only a few of them.
- **THE MODEL**: A coupled PDE-ODE system which arises from the mean-field limit of the non-controllable agents.

A brief recap

- **THE PROBLEM**: Optimal control of a system made of a large number of interacting agents by controlling only a few of them.
- **THE MODEL**: A coupled PDE-ODE system which arises from the mean-field limit of the non-controllable agents.
- **THE RESULT**: First order optimality conditions à la Pontryagin which holds for general non-local continuity equations coupled with a controlled ODE.

Advantages of our result

Advantages of our result

Our result allow us to use terminal costs with low regularity, one above all being the **Wasserstein distance** from a reference probability distribution, which is one of the most natural terminal cost choices.

Advantages of our result

Our result allow us to use terminal costs with low regularity, one above all being the **Wasserstein distance** from a reference probability distribution, which is one of the most natural terminal cost choices.

We can work with rather **general** coupled PDE-ODE systems, not only those arising as mean-field limits of finite dimensional ones.

Advantages of our result

Our result allow us to use terminal costs with low regularity, one above all being the **Wasserstein distance** from a reference probability distribution, which is one of the most natural terminal cost choices.

We can work with rather **general** coupled PDE-ODE systems, not only those arising as mean-field limits of finite dimensional ones.

We can keep track of the time evolution of the **probability distributions** rather than having to follow **each single agent** state and costate.

Advantages of our result

Our result allow us to use terminal costs with low regularity, one above all being the **Wasserstein distance** from a reference probability distribution, which is one of the most natural terminal cost choices.

We can work with rather **general** coupled PDE-ODE systems, not only those arising as mean-field limits of finite dimensional ones.

We can keep track of the time evolution of the **probability distributions** rather than having to follow **each single agent** state and costate.

Optimality conditions are **finite dimensional** and don't scale with the discretization as they depend only on the number of the controllable agents which is fixed.

Thank you for the attention!