IMPERIAL

SciCADE 2026 · New ideas in molecular dynamics

Controlling interacting particle systems

A Feedback-Control Perspective on Sampling for Langevin and McKean–Vlasov Dynamics

Dante Kalise, **Lucas M. Moschen**, and Grigorios A. Pavliotis

Department of Mathematics, Imperial College London

Why Control Probability Laws?
Since we're in Scotland...

In a Scottish ceilidh, a caller guides the group's choreography in real time. Photo © Dave Conner.

- Sampling and **many-particle systems**, modelled via SDEs, lead to mean field PDEs for the law $\mu_t \in \mathcal{P}_2(\Omega)$.
- The dynamics can be **slow** or converge to the **undesired** equilibrium.
- **Goal:** modify the dynamics so that, for $\mathsf{dist}(\mu_0, \bar{\mu})$ small,

$$\mathsf{dist}(\mu_t, \bar{\mu}) \leq Ce^{-\delta t}\,\mathsf{dist}(\mu_0, \bar{\mu}).$$

Control viewpoint in feedback form

$$\partial_t \mu_t = \mathcal{G}(\mu_t) + \mathcal{B}(\mu_t)u(t), \qquad u(t) = \mathcal{K}(\mu_t - \bar{\mu}).$$

The Sampling Problem

McKean–Vlasov dynamics: slow convergence and controlled burn-in

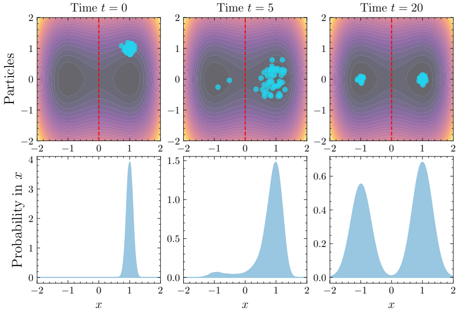

Sampling target $\bar{\mu} \propto e^{-(x^4 - 2x^2)/4}$

$$\mu_t^N = \frac{1}{N} \sum_{i=1}^{N} \delta_{X_t^i}$$

* McKean (1969), Sznitman (1991); Carrillo, McCann and Villani (2003).

- **Overdamped Langevin:**

$$dX_t^i = -\nabla\big(V + W * \mu_t^N\big)(X_t^i)\, dt + \sqrt{2\sigma}\, dB_t^i.$$

- λ_1 = spectral gap of the linearized Fokker–Planck operator around $\bar{\mu}$ controls the rate

$$\|\mu_t - \bar{\mu}\|_{L^2(\bar{\mu}^{-1})} \le e^{-\lambda_1 t} \|\mu_0 - \bar{\mu}\|_{L^2(\bar{\mu}^{-1})}.$$

Nonconvex V or strong $W \Rightarrow \lambda_1$ exponentially small.

- **Controlled burn-in:**

$$dX_t^i = -\nabla\big(V + W * \mu_t^N + U[\mu_t^N]\big)(X_t^i)\, dt + \sqrt{2\sigma}\, dB_t^i.$$

The McKean–Vlasov Equation as a Gradient Flow

Free energy $\mathcal{F} = \mathcal{E} + \sigma\,\mathrm{Ent}$ and descent dynamics

- On $\Omega = \mathbb{T}^d$, define the free energy

$$\mathcal{F}(\mu) := \mathcal{E}(\mu) + \sigma \mathrm{Ent}(\mu) = \int_\Omega V\,d\mu + \frac{1}{2}\iint_{\Omega\times\Omega} W(x-y)\,d\mu(x)\,d\mu(y) + \sigma\int_\Omega \mu \log \mu\,dx.$$

- The mean-field PDE is the W_2-gradient flow of $\mathcal{F}$:

$$\partial_t \mu_t = \nabla\cdot\left(\mu_t \nabla \frac{\delta\mathcal{F}}{\delta\mu}(\mu_t)\right) = \sigma\Delta\mu_t + \nabla\cdot(\mu_t\nabla V) + \nabla\cdot(\mu_t\nabla(W*\mu_t))\,.$$

  Here $\delta\mathcal{F}/\delta\mu$ is the first variation.

- Stationary states solve $\dfrac{\delta\mathcal{F}}{\delta\mu}(\bar\mu) = \mathrm{const} \quad\Longleftrightarrow\quad \bar\mu \propto \exp\left\{-\dfrac{V + W*\bar\mu}{\sigma}\right\}$.

  Convergence depends on the curvature of $\mathcal{F}$ around $\bar\mu$.

* Jordan, Kinderlehrer, Otto (1998), Villani (2009); Carrillo, Gvalani, Pavliotis and Schlichting (2020); Monmarché and Reygner (2023)

Where This Fits

Three neighbouring literatures

Control of SDEs

A social planner controls agents at the particle level minimising

$$\int_0^T \int \ell(x, \mu_t, u)\, d\mu_t\, dt.$$

Bensoussan, Frehse & Yam (2013); Carmona & Delarue (2018); Andersson & Djehiche (2011); Pham & Wei (2017); Fornasier & Solombrino (2014); Albi, Choi, Fornasier & Kalise (2017); Djetté, Possamai & Tan (2022); Cardaliaguet, Daudin, Jackson & Souganidis (2023); Barbu (2023).

Control of PDEs

Direct control at the distribution level using FP equations.

$$\partial_t \mu = \mathcal{G}(\mu) + \mathcal{B}(\mu)\, u$$

Blaquère (1992); Fleig & Guglielmi (2017); Aronna & Tröltzsch (2021); Breiten, Kunisch & Pfeiffer (2018); Breiten & Kunisch (2023); Duprez, Morancey & Rossi (2019); Pogodaev & Rossi (2024); Bicego, Kalise & Pavliotis (2025); Albi, Bicego & Kalise (2025).

Sampling & long-time

Exponential mixing, propagation of chaos, metastability.

$$\mathrm{W}_2(\mu_t, \bar{\mu}) \le Ce^{-\lambda t}$$

Bolley, Gentil & Guillin (2012); Carrillo, Gvalani, Pavliotis & Schlichting (2020); Raginsky, Rakhlin & Telgarsky (2017); Chizat (2022); Lelièvre, Nier & Pavliotis (2013); Duncan, Nüsken & Pavliotis (2017); Wang & Chizat (2026); Monmarché (2025); Monmarché & Reygner (2025).

This talk: feedback control $u(t) = \mathcal{K}(\mu_t - \bar{\mu})$ at the PDE level, with prescribed exponential rate δ and particle implementation $u_j^N(t) = \frac{\pi_j}{N} \sum_\ell \alpha_j(X_t^\ell)$.

The Local Stabilisation Problem
Relating the Hessian with the linearisation

Given a target equilibrium $\bar{\mu}$ and a rate $\delta > 0$, design **feedback control** so that, locally,

$$\|\mu_t - \bar{\mu}\|_{H^{-1}_{\bar{\mu}}} \leq Ce^{-\delta t}\|\mu_0 - \bar{\mu}\|_{H^{-1}_{\bar{\mu}}}, \quad \text{where } \|\rho\|_{H^{-1}_{\bar{\mu}}} := \|\mathcal{I}_{\bar{\mu}}^{-1}\rho\|_X, \quad \langle \xi, \eta \rangle_X = \int_\Omega \nabla\xi \cdot \nabla\eta \, d\bar{\mu}.$$

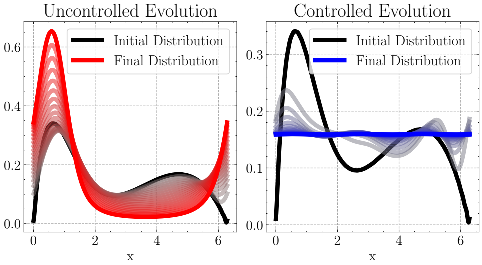

Noisy Kuramoto test case: $V = 0$, $W = -K\cos x$ and the equilibrium $\bar{\mu} = 1/(2\pi)$ is unstable when $K > 2\sigma$.

- Let $\rho = \mu - \bar{\mu} = -\nabla \cdot (\bar{\mu}\nabla\xi) =: \mathcal{I}_{\bar{\mu}}\xi$.
- The Wasserstein Hessian is represented by a self-adjoint operator with **compact resolvent**

$$A = \mathrm{Hess}_{W_2}\mathcal{F}(\bar{\mu}) \text{ on } X = \dot{H}^1_{\bar{\mu}},$$

and measures the local curvature of $\mathcal{F}$.

Linearising $\partial_t\mu_t = \mathcal{G}(\mu_t)$ around $\bar{\mu}$

$$L = -\mathcal{I}_{\bar{\mu}} A \mathcal{I}_{\bar{\mu}}^{-1}, \qquad \partial_t\rho = L\rho, \qquad \partial_t\xi = -A\xi.$$

Spectral Analysis of the Wasserstein Hessian
What prevents fast convergence?

Since A has compact resolvent, choose an X-orthonormal eigenbasis

$$Ae_k = \lambda_k e_k, \qquad \xi(t) = e^{-At}\xi_0 = \sum_{k\geq 1} e^{-\lambda_k t}\langle \xi_0, e_k\rangle_X e_k, \qquad \lambda_1 \leq \cdots \leq \lambda_k \to +\infty.$$

For a desired rate $\delta > 0$, set the **slow spectral subspace**

$$\mathcal{X}_\delta = \operatorname{span}\{e_k : \lambda_k \leq \delta\}, \qquad P_\delta \xi = \sum_{\lambda_k \leq \delta} \langle \xi, e_k\rangle_X e_k, \qquad m := \dim \mathcal{X}_\delta < \infty.$$

Only $P_\delta\xi$ can obstruct the prescribed decay rate.

Idea: add m profiles α_j for the slow modes, $\mathcal{X}_\delta \subseteq \operatorname{span}\{\alpha_1, \ldots, \alpha_m\}$,

$$\partial_t \mu_t = \nabla \cdot \left[\mu_t \nabla\left(\frac{\delta\mathcal{F}}{\delta\mu}(\mu_t) + \sum_{j=1}^{m} u_j(t)\alpha_j\right)\right].$$

* Breiten, Kunisch and Pfeiffer (2018), Kalise, Moschen, Pavliotis (2025).

Control on the Slow Spectral Window

Actuation on the slow modes

The controlled linearised dynamics takes the form

$$\partial_t \rho = L\rho + \sum_{j=1}^{m} u_j(t)\, \nabla \cdot (\bar{\mu} \nabla \alpha_j), \quad \partial_t \xi = -A\xi - Bu,$$

where $Bu := \sum_{j=1}^{m} u_j \alpha_j$. We set $\alpha_j = e_j$ so that $BB^* = P_\delta$. Under this setting, the system satisfies the infinite-dimensional δ-**Hautus** condition

$$\ker(\lambda I + A) \cap \ker(B^*) = \{0\} \quad \text{for all } \Re\lambda \geq -\delta,$$

guaranteeing the existence of a **stabilising Riccati feedback** under an additional detectability assumption.

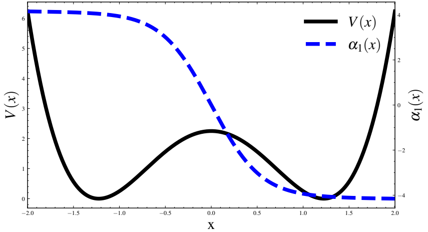

One spatial control profile.

Using the Riccati Equation to Define the Control
Linear-quadratic optimal control formulation

- **Optimal Control Problem:** Penalise the weighted linear dynamics by
$$J(\xi_0, u) := \frac{1}{2}\int_0^\infty e^{2\delta t}\left(\langle M\xi(t), \xi(t)\rangle_X + |u(t)|^2\right) dt.$$
- **Feedback Law:** The optimal control is $u^*(t) = B^*\Pi\xi(t)$, where Π is the stabilising non-negative solution to the Riccati equation
$$-(A-\delta I)\Pi - \Pi(A-\delta I) - \Pi P_\delta \Pi + M = 0.$$
- **Closed-loop Eigenvalues:** If $Me_k = m_k e_k$, then $\Pi e_k = \pi_k e_k$. The eigenvalues of the closed-loop operator $A_\Pi := A + P_\delta \Pi$ are explicitly shifted on the slow modes:
$$\lambda_k^\Pi = \begin{cases} \delta + \sqrt{(\lambda_k - \delta)^2 + m_k} & k \leq m, \\ \lambda_k & k > m. \end{cases}$$

Outcome: Linear Exponential Stabilisation

$$\partial_t \xi = -A_\Pi \xi \quad \implies \quad \|\xi(t)\|_X \leq Ce^{-\delta t}\|\xi(0)\|_X.$$

Convexification Interpretation

The feedback changes the local curvature at the target

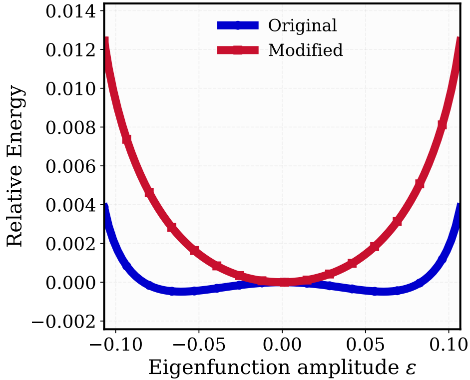

- **Closed-loop free energy:**

$$\mathcal{F}_{\mathsf{cl}}(\mu) = \mathcal{F}(\mu) + \frac{1}{2}\sum_{j=1}^{m} \pi_j \left(\int_{\mathbb{T}^d} e_j \, d(\mu - \bar{\mu}) \right)^2 .$$

- **Modified curvature:** The Wasserstein Hessian of $\mathcal{F}_{\mathsf{cl}}$ at $\bar{\mu}$ coincides with the closed-loop operator:

$$\mathrm{Hess}_{\mathrm{W}_2} \mathcal{F}_{\mathsf{cl}}(\bar{\mu}) = A + P_\delta \Pi = A_\Pi .$$

There exists $\lambda > \delta/4$ s.t. $\forall \mu_0, \mu_1 \in \mathcal{V}_\star$ (appropriate class of distributions around $\bar{\mu}$) and $s \in [0,1]$,

$$\mathcal{F}_{\mathsf{cl}}(\mu_s) \le (1-s)\mathcal{F}_{\mathsf{cl}}(\mu_0) + s\mathcal{F}_{\mathsf{cl}}(\mu_1) - \frac{\lambda}{2} s(1-s)\mathrm{W}_2^2(\mu_0, \mu_1).$$

Local Exponential Stabilisation

Guarantee for the fully nonlinear closed-loop PDE

Theorem (Local exponential stabilisation)

For every rate $\gamma \in (0, \delta)$*, there exist* $\varepsilon, C_\gamma > 0$ *such that if the initial perturbation* $\nu_0 = \mu_0 - \bar{\mu}$ *satisfies* $\mathcal{I}_{\bar{\mu}}^{-1}\nu_0 \in \dot{H}^2(\mathbb{T}^d)$ *and* $\|\nu_0\|_{L^2(\mathbb{T}^d)} \leq \varepsilon$*, the closed-loop solution is global in an appropriate sense and*

$$\|\mu_t - \bar{\mu}\|_{H_{\bar{\mu}}^{-1}} \leq C_\gamma e^{-\gamma t}\|\mu_0 - \bar{\mu}\|_{L^2(\mathbb{T}^d)}.$$

Local Exponential Stabilisation

Guarantee for the fully nonlinear closed-loop PDE

Theorem (Local exponential stabilisation)

For every rate $\gamma \in (0, \delta)$*, there exist* $\varepsilon, C_\gamma > 0$ *such that if the initial perturbation* $\nu_0 = \mu_0 - \bar{\mu}$ *satisfies* $\mathcal{I}_{\bar{\mu}}^{-1}\nu_0 \in \dot{H}^2(\mathbb{T}^d)$ *and* $\|\nu_0\|_{L^2(\mathbb{T}^d)} \le \varepsilon$*, the closed-loop solution is global in an appropriate sense and*

$$\|\mu_t - \bar{\mu}\|_{H_{\bar{\mu}}^{-1}} \le C_\gamma e^{-\gamma t}\|\mu_0 - \bar{\mu}\|_{L^2(\mathbb{T}^d)}.$$

How the proof works

- **Weighted equation:** Setting $w_t := e^{\gamma t}\nu_t$ gives $\partial_t w_t - (L_\Pi + \gamma I)w_t = \mathfrak{R}_\gamma(w)$.
- **Linear estimate:** Since $\gamma < \delta$, maximal regularity yields $\|w\|_{\mathcal{Y}_T} \le C_{\text{lin}}\big(\|\nu_0\|_{L^2} + \|\mathfrak{R}_\gamma(w)\|\big)$, uniformly in T.
- **Quadratic remainder:** Near $\bar{\mu}$, $\|\mathfrak{R}_\gamma(w)\| \le C_{\text{rem}}\|w\|_{\mathcal{Y}_T}^2$.
- **Bootstrap:** Thus $\|w\|_{\mathcal{Y}_T} \le C_{\text{lin}}\|\nu_0\|_{L^2} + C\|w\|_{\mathcal{Y}_T}^2$. For small ν_0, the solution remains below the continuation radius, so $T_{\max} = \infty$ and $\nu_t = O(e^{-\gamma t})$.

Local Exponential Stabilisation

Guarantee for the fully nonlinear closed-loop PDE

Theorem (Local exponential stabilisation)

For every rate $\gamma \in (0, \delta)$*, there exist* $\varepsilon, C_\gamma > 0$ *such that if the initial perturbation* $\nu_0 = \mu_0 - \bar{\mu}$ *satisfies* $\mathcal{I}_{\bar{\mu}}^{-1}\nu_0 \in \dot{H}^2(\mathbb{T}^d)$ *and* $\|\nu_0\|_{L^2(\mathbb{T}^d)} \leq \varepsilon$*, the closed-loop solution is global in an appropriate sense and*

$$\|\mu_t - \bar{\mu}\|_{H^{-1}_{\bar{\mu}}} \leq C_\gamma e^{-\gamma t}\|\mu_0 - \bar{\mu}\|_{L^2(\mathbb{T}^d)}.$$

Corollary (Wasserstein stabilisation)

By the comparison inequality $\mathrm{W}_2(\mu_t, \bar{\mu}) \leq 2\|\mu_t - \bar{\mu}\|_{H^{-1}_{\bar{\mu}}}$*, we have*

$$\mathrm{W}_2(\mu_t, \bar{\mu}) \leq 2C_\gamma e^{-\gamma t}\|\nu_0\|_{L^2(\mathbb{T}^d)}.$$

Kuramoto under Feedback

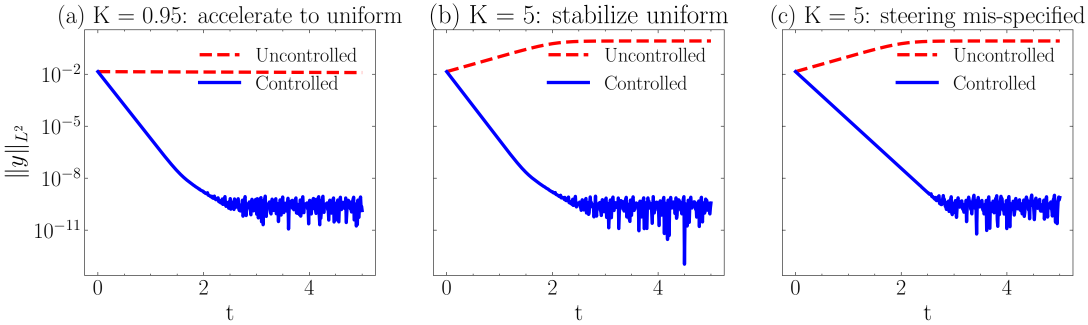

Figure: Control of the noisy Kuramoto model ($\sigma = 0.5$, $W(x) = -K\cos x$). Panels (a) and (b) target the uniform distribution $\bar{\mu} = (2\pi)^{-1}$, while panel (c) targets a selected translated synchronised equilibrium.

From the PDE to the Particle System

How do we apply the PDE control in the SDEs?

The PDE control is computed with $\alpha_j = e_j$. By plugging in the empirical measure, we have

$$u_j(t) = \pi_j \int_\Omega e_j \, d\mu_t \quad \Rightarrow \quad u_j^N(t) = \frac{\pi_j}{N} \sum_{\ell=1}^{N} e_j(X_t^\ell).$$

Advantage: The control is evaluated directly from particle positions. **No density reconstruction.**

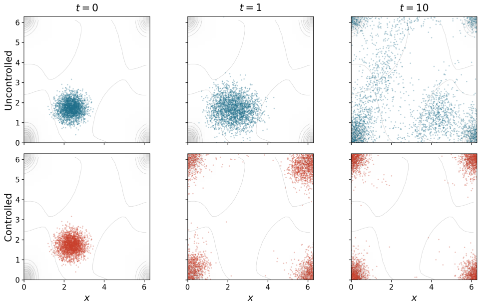

Uncontrolled (top) vs. feedback-controlled (bottom).

Controlled N-particle system

$$dX_t^i = -\left(\nabla V(X_t^i) + \frac{1}{N} \sum_{k=1}^{N} \nabla W(X_t^i - X_t^k) + \sum_{j=1}^{m} u_j^N(t) \nabla e_j(X_t^i) \right) dt + \sqrt{2\sigma}\, dB_t^i.$$

The First Estimate

Finite-time propagation of chaos

Theorem (Standard finite-time consistency)

For every $T > 0$ and $t \in [0, T]$, the mean squared Wasserstein distance satisfies

$$\mathbb{E}\mathrm{W}_2^2(\mu_t^N, \bar{\mu}) \leq \frac{C_0}{N}(e^{C_1 t} - 1) + C_1 \beta_d(N) + C_2 e^{-2\delta t}\|\nu_0\|_{L^2(\mathbb{T}^d)}^2.$$

- $\beta_d(N)$ is the empirical measure error and comes from Fournier–Guillin (2015).
- **Issue:** The Gronwall error $e^{C_1 t}/N$, due to the coupling error $\mathrm{W}_2(\mu_t^N, \bar{\mu}_t^N)$, blows up for large t.
- The estimate ceases to be useful beyond the timescale $t \approx (\log N)/C_1$.
- **Idea:** The PDE provides strict **local contraction** $\mathrm{W}_2(S_T^{\Pi}\mu, \bar{\mu}) \leq q\mathrm{W}_2(\mu, \bar{\mu})$ for $q < 1$. We can avoid global Gronwall by restarting the estimate over time blocks of length T.

The Block Recurrence Strategy
Restarting the estimate on short intervals

- **PDE Contraction:** Over a fixed block length T and $\mathrm{W}_2(\mu, \bar{\mu}) \leq R$,

$$\mathrm{W}_2(S_T^{\Pi}\mu, \bar{\mu}) \leq q\mathrm{W}_2(\mu, \bar{\mu}) \quad (q < 1).$$

- **Conditional Consistency:** Over this same short block T,

$$\mathbb{P}\left(\mathrm{W}_2\left(\mu^N_{(k+1)T}, S_T^{\Pi}\mu^N_{kT}\right) \geq \sqrt{C_T\beta_d(N)} + r\right) \leq e^{-c_T N r^2}.$$

The Recurrence Relation

Setting $D_k = \mathrm{W}_2(\mu^N_{kT}, \bar{\mu})$, on the event that $D_k \leq R$,

$$D_{k+1} \leq qD_k + \mathrm{W}_2\left(\mu^N_{(k+1)T}, S_T^{\Pi}\mu^N_{kT}\right).$$

Long-Time Confinement

PDE contraction is stable under empirical approximation

Theorem

For large N and a given initial configuration $\mathbf{X}_0 = \mathbf{x}$, if $\mathrm{W}_2(\mu_0^N, \bar{\mu})$ is small, the empirical measure satisfies:

$$\mathbb{E}_{\mathbf{x}} \mathrm{W}_2^2(\mu_t^N, \bar{\mu}) \le C e^{-2\gamma t} \mathrm{W}_2^2(\mu_0^N, \bar{\mu}) + C\beta_d(N) + C \min\{1, t e^{-cN}\}.$$

- **Burn-in phase:** The distance decays exponentially until it hits the statistical floor $\beta_d(N)$.
- **Metastability:** The error is bounded by $C\beta_d(N)$ for exponentially long times $t \le e^{c' N}$.
- **Conclusion:** The feedback provides a particle-level burn-in accelerator.

Opinion Dynamics under Feedback

Restoring the unsynchronised state in the Hegselmann–Krause model

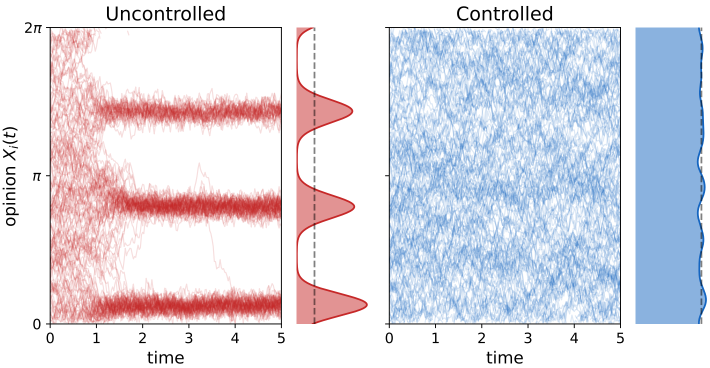

Figure: The empirical dynamics of the bounded confidence model ($N = 5000$, $\sigma = 0.225$, and wrapped Gaussian interaction W). Uncontrolled particles fragment into multiple clusters, while the spectral feedback restores the unsynchronised state at $\bar{\mu} = 1/(2\pi)$.

What Was Hard

Technical obstacles and open questions

1. Propagation of chaos with feedback

- Standard finite-time estimates degrade exponentially: $C_T \propto e^T$.
- Solved via the **block recurrence** argument.
- Given the local nature, uniform-in-time propagation of chaos remains open, while permanent pathwise confinement **cannot hold**: there is a positive probability that the empirical measure leaves the attraction region.

2. Computing the slow eigenfunctions

- The feedback requires solving $Ae_k = \lambda_k e_k$ offline.
- Fine in low dimensions; very expensive in high dimensions or complex potentials.
- **Open:** Can we *learn* the slow modes online from particle data?

Broader Question: Is this competitive with enhanced sampling methods at equal computational cost?

What Was Hard

Technical obstacles and open questions

1. Propagation of chaos with feedback

- Standard finite-time estimates degrade exponentially: $C_T \propto e^T$.
- Solved via the **block recurrence** argument.
- Given the local nature, uniform-in-time propagation of chaos remains open, while permanent pathwise confinement **cannot hold**: there is a positive probability that the empirical measure leaves the attraction region.

2. Computing the slow eigenfunctions

- The feedback requires solving $Ae_k = \lambda_k e_k$ offline.
- Fine in low dimensions; very expensive in high dimensions or complex potentials.
- **Open:** Can we *learn* the slow modes online from particle data?

Broader Question: Is this competitive with enhanced sampling methods at equal computational cost?

Honest answer: we don't know yet.

Summary and Open Questions

What I showed you

- A **feedback design** for Wasserstein gradient flows: Riccati gains on the slow Hessian modes, from particle observables at cost $O(mN)$.
- **PDE guarantee:** local exponential stabilisation at prescribed rate δ.
- **Particle guarantee:** ensemble stays near $\bar{\mu}$ for time $e^{c'N}$, error $\lesssim \beta_d(N)$.

What I don't know yet

- Can feedback help **globally**, not only near $\bar{\mu}$?
- Can we **learn** the slow eigenfunctions online from particle data?
- What happens if the normalising constant Z is unknown? $\rightarrow$ work in progress
- Extension to **underdamped Langevin**?

Happy to discuss any of these!

Papers:

2025: Linearization-Based Feedback Stabilization of McKean–Vlasov PDEs

2026: Feedback Control and Local Convexification of Wasserstein Gradient Flows

lucas.moschen22@imperial.ac.uk

Multi-Species Coupled Dynamics
What matters is the gradient-flow structure with entropy

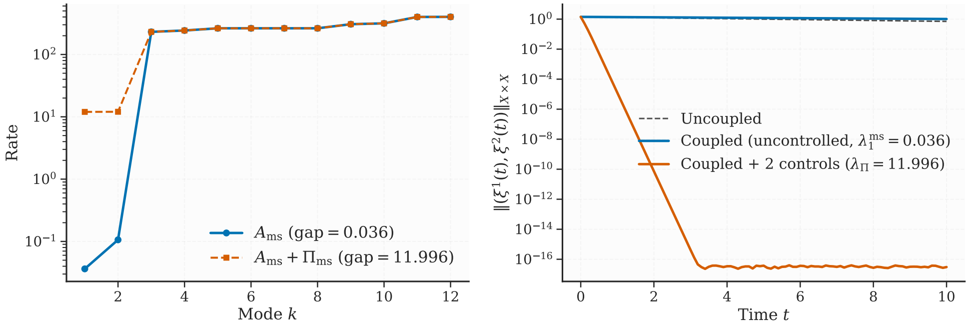

Figure: Feedback stabilisation of a two-species system on $\mathbb{T}$ with $V(x) = -2\cos(4\pi x)$, $\sigma = 0.5$, and cross-interaction $W_{ij}(x) = -0.5\cos(2\pi x)$. A two-control Riccati feedback computed on the product space $X \times X$ shifts the two slow joint modes above the target rate $\delta = 5$ (left, red).

Local Displacement Convexity

From a pointwise Hessian bound to a regional property

- We now pass from the strict Hessian lower bound at $\bar{\mu}$ to a local property.
- **Admissible neighbourhood:** We fix a class $\mathcal{V} \subseteq \mathrm{dom}(\mathcal{F})$ where the Hessian is well-defined, and define an L^∞-neighbourhood of the target:

$$\mathcal{V}_r := \left\{ \mu \in \mathcal{V} : \left\| \frac{\mu}{\bar{\mu}} - 1 \right\|_{L^\infty(\mathbb{T}^d)} \leq r \right\}.$$

- **Geodesic subset:** Let $\mathcal{V}_\star \subseteq \mathcal{V}_{r_*}$ be a maximal subset containing $\bar{\mu}$ connected by constant-speed W_2-geodesics.

Theorem (Local λ-displacement convexity)

For r_* sufficiently small, $\mathcal{F}_{\mathrm{cl}}$ is λ-displacement convex on $\mathcal{V}_\star$. Equivalently, for every $\mu_0, \mu_1 \in \mathcal{V}_\star$ and $s \in [0,1]$,

$$\mathcal{F}_{\mathrm{cl}}(\mu_s) \leq (1-s)\mathcal{F}_{\mathrm{cl}}(\mu_0) + s\mathcal{F}_{\mathrm{cl}}(\mu_1) - \frac{\lambda}{2} s(1-s)\mathrm{W}_2^2(\mu_0, \mu_1).$$