Grigorios A. Pavliotis

Department of Mathematics Imperial College London

Mean field limits for weakly interacting diffusions : phase transitions, multiscale analysis and inference

St Entienne, 19 May, 2022

Imperial College London

- Study systems of interacting particles/agents, in the presence of noise.
- Consider systems for which collective behaviour emerges as the result of the interaction between agents.
- Interpret the emergence of collective behaviour as a disorder/order phase transition.
- Conditions for the existence of phase transitions, calculate transition points.
- Study the effect of colored noise, memory, inertia....
- Importance of fluctuations, hydrodynamic/macroscopic limits.
- Learn order parameters, predict phase transitions from data.

- Consider systems of interacting diffusions that exhibit phase transitions in their mean field limit.
- Interacting diffusions of this form are used as models for
  - Synchronization (Kuramoto).
  - Opinion formation (bounded confidence models, Hagselmann-Krause).
  - Systemic risk and cooperation.

The Kuramoto model: $\dot{x}_i = -\frac{1}{N}\sum_{j=1}^{N}\sin(x_i - x_j) + \sqrt{2\beta^{-1}}\dot{W}_i$.

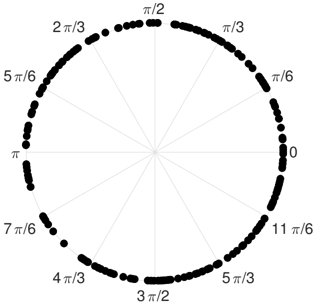

$\kappa < \kappa_c$, no phase locking

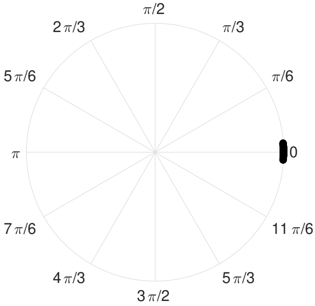

$\kappa > \kappa_c$, phase locking

Opinion dynamics[1]

- Cooperative dynamics:

$$\mathrm{d}x_t^{i,N} = -\frac{1}{N}\sum_{j=1}^{N} \varphi_\theta(||x_t^{i,N} - x_t^{j,N}||)(x_t^{i,N} - x_t^{j,N})\mathrm{d}t + \sigma \mathrm{d}w_t^{i,N}$$

- Interaction function

$$\varphi_\theta(r) = \theta_1 \exp\left[-\frac{0.01}{1-(r-\theta_2)^2}\right],$$

smooth approximation to $\widetilde{\varphi}_\theta(r) = \theta_1 \mathbf{1}_{r\in[0,\theta_2+1]}$

[1] Wang et al. J. Stat. Phys. 2017, Garnier et al Vietn. J. Math. 2017

Opinion dynamics

Imperial College London

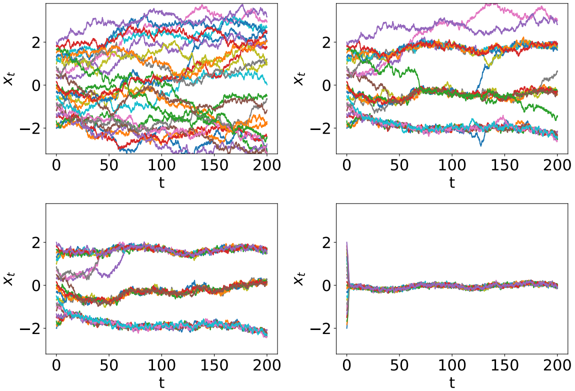

Figure: Sample trajectories of the system of interacting particles for $\theta_2 = \{0.0, 0.3, 0.5, 1.0\}$.

- Clustering.
- Definition of (not physics-informed) order parameter.
- Choice of correct boundary conditions, effect of radicals/extreme groups.[2]

[2] Goddard, P. et al IMA J. Appl. Math. 2022

Cooperative Behaviour/Systemic Risk[3]

- Consider a system of interacting diffusions in a bistable potential:

$$dX_t^i = \left(-V'(X_t^i) - \theta \left(X_t^i - \frac{1}{N} \sum_{j=1}^{N} X_t^j \right) \right) dt + \sqrt{2\beta^{-1}}\, dB_t^i.$$

- The total energy (Hamiltonian) is

$$W_N(\mathbf{X}) = \sum_{\ell=1}^{N} V(X^\ell) + \frac{\theta}{4N} \sum_{n=1}^{N} \sum_{\ell=1}^{N} (X^n - X^\ell)^2.$$

- We can pass rigorously to the mean field limit as $N \to \infty$ using, for example, martingale (Oelschlager 1984) or variational/$\Gamma-$convergence techniques (Carrillo, Delgadino, P. J. Func. Analysis (2020)).

[3] Dawson J. Stat. Phys. 1983, Garnier et al SIAM Math Finance 2013.

- Formally, using the law of large numbers we obtain the McKean SDE

$$dX_t = -V'(X_t)\,dt - \theta(X_t - \mathbb{E}X_t)\,dt + \sqrt{2\beta^{-1}}\,dB_t.$$

- The Fokker-Planck equation corresponding to this SDE is the McKean-Vlasov equation

$$\frac{\partial p}{\partial t} = \frac{\partial}{\partial x}\left(V'(x)p + \theta\left(x - \int_{\mathbb{R}} xp(x,t)\,dx\right)p + \beta^{-1}\frac{\partial p}{\partial x}\right).$$

- The McKean-Vlasov equation is a gradient flow, with respect to the Wasserstein metric, for the free energy functional

$$\mathcal{F}[\rho] = \beta^{-1}\int \rho \ln \rho\,dx + \int V\rho\,dx + \frac{\theta}{2}\int\int F(x-y)\rho(x)\rho(y)\,dxdy,$$

with $F(x) = \frac{1}{2}x^2$.

- The finite dimensional dynamics is reversible with respect to the Gibbs measure

$$\mu_N(dx) = \frac{1}{Z_N} e^{-\beta W_N(x^1,\ldots x^N)}\, dx^1 \ldots dx^N, \quad Z_N = \int_{\mathbb{R}^N} e^{-\beta W_N(x^1,\ldots x^N)}$$

- This can be written in the standard form of the Gibbs measure for an unbounded spin system:

$$\mu_N(dx) = \frac{1}{Z_N} e^{-\theta\beta \sum_{i \neq j} x_i x_j} \prod_{j=1}^{N} \pi(dx_i),$$

- the McKean dynamics can have more than one invariant measures, for nonconvex confining potentials and at sufficiently low temperatures (Dawson 1983, Tamura 1984, Shiino 1987).
- The density of the invariant measure(s) for the McKean dynamics satisfies the stationary nonlinear Fokker-Planck equation

$$\frac{\partial}{\partial x}\left(V'(x)p_\infty + \theta\left(x - \int_{\mathbb{R}} x p_\infty(x)\, dx\right) p_\infty + \beta^{-1}\frac{\partial p_\infty}{\partial x}\right) = 0.$$

- For the quadratic interaction potential a one-parameter family of solutions to the stationary McKean-Vlasov equation can be obtained: $p_\infty(x;\theta,\beta,m) = \frac{1}{Z_{\theta,\beta;m}} e^{-\beta\left(V(x)+\theta\left(\frac{1}{2}x^2 - xm\right)\right)}, \quad Z_{\theta,\beta;m} = \int_{\mathbb{R}} e^{-\beta\left(V(x)+\theta\left(\frac{1}{2}x^2 - xm\right)\right)}\, dx.$

- These solutions are subject to the constraint that they provide us with the correct formula for the first moment:

$$m = \int_{\mathbb{R}} x p_\infty(x;\theta,\beta,m)\,dx =: R(m;\theta,\beta).$$

- This is the **selfconsistency** equation:

$$m = R(m;\theta,\beta).$$

- The critical temperature can be calculated from

$$\mathrm{Var}_{p_\infty}(x)\Big|_{m=0} = \frac{1}{\beta\theta}.$$

Phase transitions

The number of solutions, and their stability, depends on the temperature.

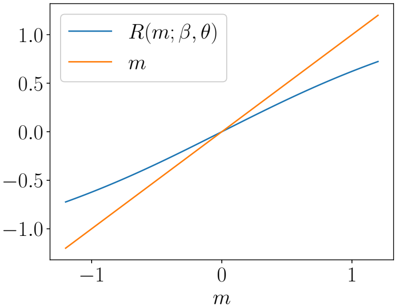

(a) $\beta = 1$: $m = 0$ is stable.

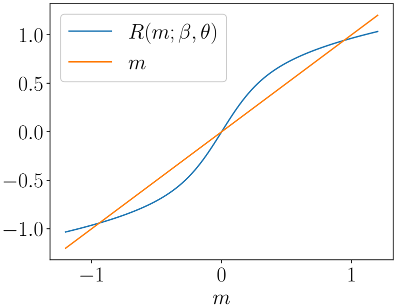

(b) $\beta = 9$: $m = 0$ is unstable.

The slope of $m - R(m;\beta,\theta)$ determines the *stability* of the steady states.

Stability of steady states

The stability of the solution can be seen from the associated free energy:

$$\mathcal{F}[\rho] = \beta^{-1} \int_{\mathbf{R}} \rho(x) \, \ln \rho(x) \, \mathrm{d}x + \int_{\mathbf{R}} V(x) \, \rho(x) \, \mathrm{d}x + \frac{\theta}{2} \int_{\mathbf{R}} \int_{\mathbf{R}} F(x-y) \, \rho(x) \, \rho$$

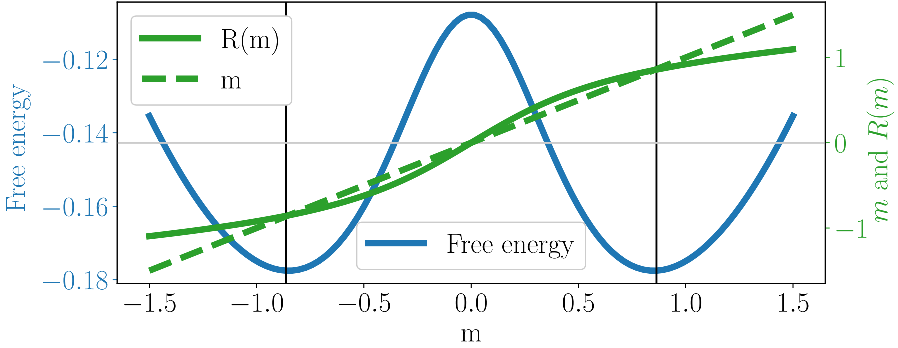

Bifurcation diagram

By solving the self-consistency equation for many values of β (or θ), we can construct the bifurcation diagram:

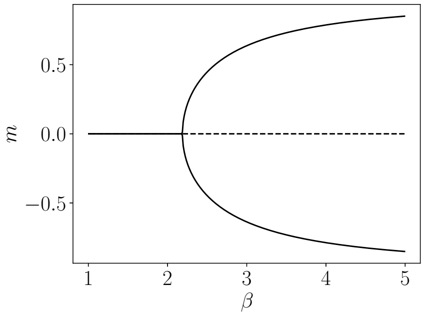

(a) Bifurcation diagram.

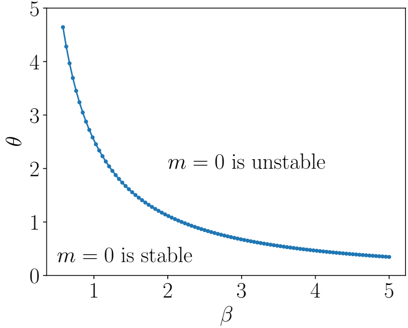

(b) Stability of the "$m = 0$" solution.

Equations for the moments[4]

- We can obtain an infinite system of equations for the moments of the McKean-Vlasov PDE:

$$\dot{M}_k(t) = k\big[(1-\theta)M_k(t) - M_{k+2}(t) + \beta^{-1}(k-1)M_{k-2}(t) \\ + \theta M_1(t)M_{k-1}(t)\big].$$

- We truncate the above system, with appropriate boundary conditions (cumulative moment expansion), and use it to obtain a low dimensional description of the dynamics **close to the phase transition**.

[4] Lucarini, Zagli, P. 2022

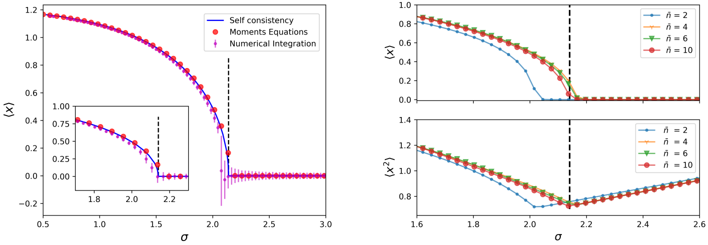

Figure: Approximation of the phase transition using the truncated moments representation. Left panel: $n = 10$. Comparison with calculation of $\langle X_t \rangle$ using the time average of the empirical mean over a long trajectory that includes several transitions between the two metastable states.

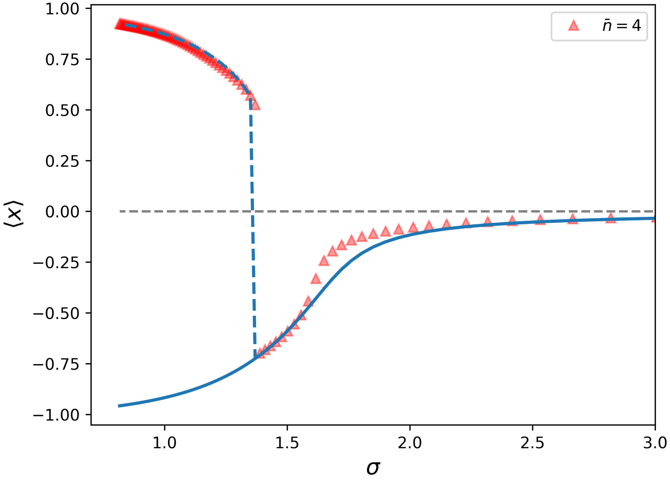

Figure: Approximation of the phase transition using the truncated moments representation for discontinuous phase transitions

Phase transitions for colored noise[5]

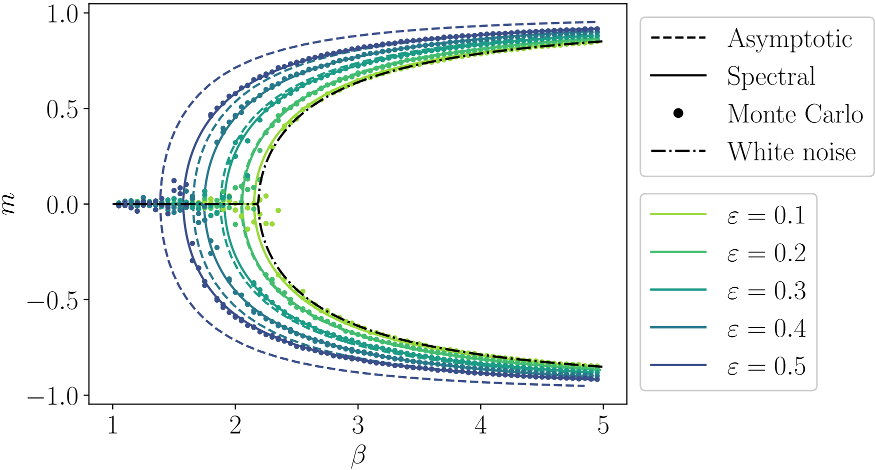

Figure: Bifurcation diagram with scalar OU noise

[5] S. Gomes, G.P., U. Vaes, SIAM MMS 2020.

The McKean-Vlasov equation on the torus

Arch. Rational Mech. Anal. 235 (2020) 635–690
Digital Object Identifier (DOI) https://doi.org/10.1007/s00205-019-01430-4

Check for updates

Long-Time Behaviour and Phase Transitions for the Mckean–Vlasov Equation on the Torus

J. A. Carrillo, R. S. Gvalani, G. A. Pavliotis & A. Schlichting

Communicated by A. Garroni

Abstract

We study the McKean–Vlasov equation

Nonlocal parabolic PDE

$$\frac{\partial \varrho}{\partial t} = \beta^{-1}\Delta\varrho + \kappa\nabla\cdot(\varrho\nabla W \star \varrho) \qquad \text{in } \mathbb{T}_L^d \times (0,T]$$

with periodic boundary conditions, $\varrho(\cdot,0) = \varrho_0 \in \mathcal{P}(\mathbb{T}_L^d)$,
$\mathbb{T}_L^d \hat{=} \left(-\frac{L}{2}, \frac{L}{2}\right)^d$

- $\varrho(\cdot,t) \in \mathcal{P}(\mathbb{T}_L^d)$ probability density of particles
- W coordinate-wise even interaction potential
- $\beta > 0$ inverse temperature (fixed)
- $\kappa > 0$ interaction strength (parameter)

Example: The noisy Kuramoto model

Imperial College London

The Kuramoto model: $W(x) = -\sqrt{\frac{2}{L}}\cos\left(2\pi k \frac{x}{L}\right), k \in \mathbb{Z}$

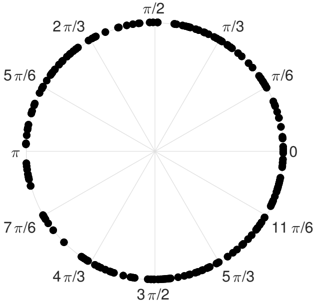

$\kappa < \kappa_c$, no phase locking

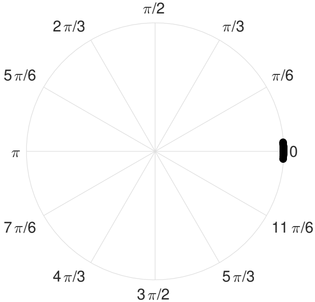

$\kappa > \kappa_c$, phase locking

Fourier representation $\widetilde{f}(k) = \langle f, w_k \rangle_{L^2(\mathbb{T}_L)}$ with $k \in \mathbb{Z}^d$

- A function $W \in L^2(\mathbb{T}_L^d)$ is ***H*-stable**, $W \in \mathbb{H}_s$, if

$$\widetilde{W}(k) = \langle W, w_k \rangle \geq 0, \quad \forall k \in \mathbb{Z}^d \ ,$$

- Decomposition of potential W into H-stable and H-unstable part

$$W_{\mathrm{s}}(x) = \sum_{k \in \mathbb{N}^d} (\langle W, w_k \rangle)_+ w_k(x) \qquad \text{and} \qquad W_{\mathrm{u}}(x) = W(x) - W_s(x) \ .$$

- Free energy functional $\mathscr{F}_\kappa$: Driving the W_2-gradient flow

$$\mathscr{F}_\kappa(\varrho) = \beta^{-1}\int_{\mathbb{T}_L^d} \varrho \log \varrho \,\mathrm{d}x + \frac{\kappa}{2}\iint\limits_{\mathbb{T}_L^d\times\mathbb{T}_L^d} W(x-y)\varrho(x)\varrho(y)\,\mathrm{d}x\,\mathrm{d}y \ .$$

- Dissipation: $\mathscr{F}_\kappa$ is Lyapunov-function

$$\mathcal{J}_\kappa(\varrho) = -\frac{\mathrm{d}}{\mathrm{d}t}\mathscr{F}_\kappa(\varrho) = \int_{\mathbb{T}_L^d} \left|\nabla \log \frac{\varrho}{e^{-\beta\kappa W\star\varrho}}\right|^2 \varrho\,\mathrm{d}x \ ,$$

- Kirkwood-Monroe fixed point mapping

$$F_\kappa(\varrho) = \varrho - \mathcal{T}\varrho = \varrho - \frac{1}{Z(\varrho,\kappa)} e^{-\beta\kappa W\star\varrho}, \quad \text{with} \quad Z(\varrho,\kappa) = \int_{\mathbb{T}_L^d} e^{-\beta\kappa W\star\varrho}$$

Characterization of stationary states: The following are equivalent

- ϱ is a stationary state: $\beta^{-1}\Delta\varrho + \kappa\nabla\cdot(\varrho\nabla W \star \varrho) = 0$.
- ϱ is a root of $F_\kappa(\varrho)$.
- ϱ is a global minimizer of $\mathcal{J}_\kappa(\varrho)$.
- ϱ is a critical point of $\mathscr{F}_\kappa(\varrho)$.

$\Rightarrow \varrho_\infty \equiv L^{-d}$ is a stationary state for all $\kappa > 0$.

Theorem

Under appropriate assumptions on the potential, for $\varrho_0 \in H^{3+d}(U) \cap \mathcal{P}_{\mathrm{ac}}(U)$, there exists a unique classical solution ϱ of the McKean-Vlasov equation such that $\varrho(\cdot, t) \in \mathcal{P}_{\mathrm{ac}}(U) \cap C^2(\overline{U})$ for all $t > 0$. Additionally, $\varrho(\cdot, t)$ is strictly positive and has finite entropy, i.e, $\varrho(\cdot, t) > 0$ and $S(\varrho(\cdot, t)) < \infty$, for all $t > 0$.

Theorem

(Convergence to equilibrium) Let $\varrho(x,t)$ be a classical solution of the Mckean–Vlasov equation with smooth initial data and smooth, even, interaction potential W. Then we have:

1. *If $0 < \kappa < \frac{2\pi}{3\beta L \|\nabla W\|_\infty}$, then $\|\varrho - \frac{1}{L}\|_2 \to 0$, exponentially, as $t \to \infty$,*
2. *If $\hat{W}(k) \geq 0$ for all $k \in \mathbb{Z}$ or $0 < \kappa < \frac{2\pi^2}{\beta L^2 \|\Delta W\|_\infty}$, then $\mathcal{H}\left(\varrho | \frac{1}{L}\right) \to 0$, exponentially, as $t \to \infty$,*

where $\hat{W}(k)$ represents the Fourier transform and $\mathcal{H}\left(\varrho | \frac{1}{L}\right)$ represents the relative entropy.

Nontrivial solutions to the stationary McKean–Vlasov equation?

- $W \notin \mathbb{H}_s$ is a necessary condition for the existence of nontrivial steady states.
- Numerical experiments indicate one, multiple, or possibly infinite solutions
- What determines the number of nontrivial solutions?
- Birfurcation analysis of $\varrho \mapsto F_\kappa(\varrho)$.

Example: Kuramoto model: $W(x) = -\sqrt{\frac{2}{L}}\cos(2\pi x/L)$

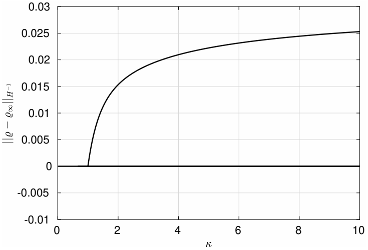

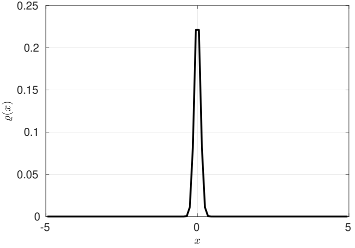

$\Rightarrow$ 1-cluster solution and uniform state ϱ_∞.

Theorem

London

(Local bifurcations) Let W *be smooth and even and let* $(1/L, \kappa)$ *represent the trivial branch of solutions. Then every* $k^* \in \mathbb{Z}, k > 0$ *such that*

1. $\operatorname{card}\Big\{k \in \mathbb{Z}, k > 0 : \hat{W}(k) = \hat{W}(k^*)\Big\} = 1$,
2. $\hat{W}(k) < 0,$

corresponds to a bifurcation point of the stationary McKean–Vlasov equation through the formula

$$\kappa_* = -\frac{\sqrt{L}}{\beta \hat{W}(k^*)} ,$$

with $(1/L, \kappa_*)$ *the bifurcation point.*

Definition (Transition point [Chayes & Panferov '10])

A parameter value $\kappa_c > 0$ is said to be a transition point of $\mathscr{F}_\kappa$ if it satisfies the following conditions,

1. For $0 < \kappa < \kappa_c$: ϱ_∞ is the unique minimiser of $\mathscr{F}_\kappa(\varrho)$
2. For $\kappa = \kappa_c$: ϱ_∞ is a minimiser of $\mathscr{F}_\kappa(\varrho)$.
3. For $\kappa > \kappa_c$: $\exists \varrho_\kappa \neq \varrho_\infty$, such that ϱ_κ is a minimiser of $\mathscr{F}_\kappa(\varrho)$.

Definition (Continuous and discontinuous transition point)

A transition point $\kappa_c > 0$ is a continuous transition point of $\mathscr{F}_\kappa$ if

1. For $\kappa = \kappa_c$: ϱ_∞ is the unique minimiser of $\mathscr{F}_\kappa(\varrho)$.
2. For any family of minimizers $\{\varrho_\kappa \neq \varrho_\infty\}_{\kappa > \kappa_c}$ it holds

$$\limsup_{\kappa \downarrow \kappa_c} \|\varrho_\kappa - \varrho_\infty\|_1 = 0.$$

A transition point $\kappa_c > 0$ which is not continuous is discontinuous.

Summary of critical points:

- κ_c transition point.
- κ_* bifurcation point.
- $\kappa_\sharp$ point of linear stability, i.e., $\kappa_\sharp = -\frac{L^{\frac{d}{2}}}{\beta \min_k \widetilde{W}(k)/\Theta(k)}$ with $k_\sharp = \arg\min \widetilde{W}(k)$.

  If there is exactly one $k_\sharp$, then $\kappa_\sharp = \kappa_*$ is a bifurcation point.

Theorem

(Discontinuous and continuous phase transitions) Let W be smooth and even and assume the free energy $\mathscr{F}_{\kappa,\beta}$ exhibits a transition point, $\kappa_c < \infty$. Then we have the following two scenarios:

1. *If there exist strictly positive $k^a, k^b, k^c \in \mathbb{Z}$ with $\hat{W}(k^a) = \hat{W}(k^b) = \hat{W}(k^c) = \min_k \hat{W}(k) < 0$ such that $k^a = k^b + k^c$ or $k^a = 2k^b$, then κ_c is a discontinuous transition point.*
2. *Let $k^\sharp = \arg\min_k \hat{W}(k)$ be well-defined with $\hat{W}(k^\sharp) < 0$. Let W_α denote the potential obtained by multiplying all the negative $\hat{W}(k)$ except $\hat{W}(k^\sharp)$ by some $\alpha \in (0, 1]$. Then if α is made small enough, the transition point κ_c is continuous.*

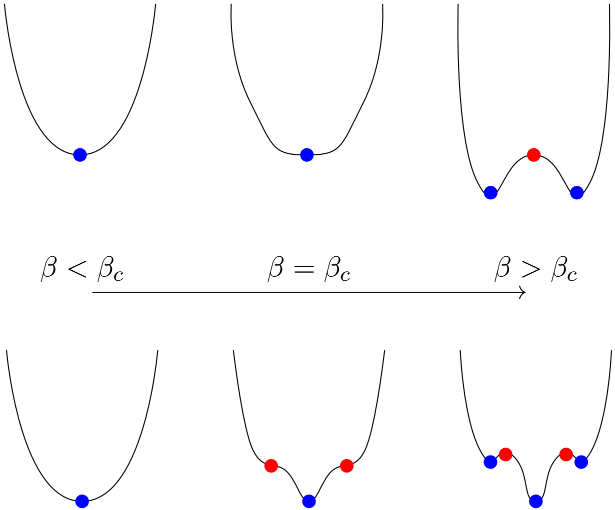

Figure: Continuous vs discontinuous phase transitions. Continuous phase transition (upper diagram): the unique critical point (shown in blue) loses its local stability through a local (pitchfork) bifurcation which gives rise to new locally stable critical points. Discontinuous phase transition (lower diagram): the unique critical point retains its local stability but new critical points arise in the free energy landscape through a saddle node bifurcation.

Proposition

The generalised Kuramoto model $W(x) = -w_k(x)$*, for some* $k \in \mathbb{N}, k \neq 0$ *exhibits a continuous transition point at* $\kappa_c = \kappa_\sharp$*. Additionally, for* $\kappa > \kappa_c$*, the equation* $F(\varrho, \kappa) = 0$ *has only two solutions in* $L^2(U)$ *(up to translations). The nontrivial one,* ϱ_κ *minimises* $\mathscr{F}_\kappa$ *for* $\kappa > \kappa_c$ *and converges in the narrow topology as* $\kappa \to \infty$ *to a normalised linear sum of equally weighted Dirac measures centred at the minima of* $W(x)$*.*

The noisy Hegselmann–Krause model for opinion dynamics

- The noisy Hegselmann–Krause system models the opinions of N interacting agents such that each agent is only influenced by the opinions of its immediate neighbours. The interaction potential is

$$W_{\mathrm{hk}}(x) = -\frac{1}{2}\left(\left(|x| - \frac{R}{2}\right)_{-}\right)^{2}$$

- for some $R > 0$. The ratio R/L measures the range of influence of an individual agent with $R/L = 1$ representing full influence.

The noisy Hegselmann–Krause model for opinion dynamics

- The Fourier transform of $W_{\mathrm{hk}}(x)$ is

$$\widetilde{W}_{\mathrm{hk}}(k) = \frac{(-\pi^2 k^2 R^2 + 2L^2)\sin\left(\frac{\pi k R}{L}\right) - 2\pi k L R \cos\left(\frac{\pi k R}{L}\right)}{4\sqrt{2}\pi^3 k^3 \sqrt{\frac{1}{L}}},$$

- with $k \in \mathbb{N}, k \neq 0$.
- the model has infinitely many bifurcation points for $R/L = 1$.

- We define a rescaled version of the potential

$$W_{\mathrm{hk}}^{R}(x) = -\frac{1}{2R^3}\left(\left(|x| - \frac{R}{2}\right)_-\right)^2,$$

which does not lose mass as $R \to 0$.

Proposition

For R small enough, the rescaled noisy Hegselmann–Krause model possesses a discontinuous transition point.

Imperial College London

Arch. Rational Mech. Anal.
Digital Object Identifier (DOI) https://doi.org/10.1007/s00205-021-01648-1

Check for updates

On the Diffusive-Mean Field Limit for Weakly Interacting Diffusions Exhibiting Phase Transitions

Matias G. Delgadino, Rishabh S. Gvalani & Grigorios A. Pavliotis

Communicated by P.-L. Lions

Abstract

The objective of this article is to analyse the statistical behaviour of a large number of weakly interacting diffusion processes evolving under the influence of a

The joint diffusive-mean field limit

We consider the following system of weakly interacting diffusions with a periodic interaction potential:

$$\mathrm{d}X_t^{i,\varepsilon} = -\nabla V(\varepsilon^{-1}X_t^{i,\varepsilon})\,dt - \frac{1}{N}\sum_{i\neq j}^{N}\nabla W(\varepsilon^{-1}(X_t^{i,\varepsilon} - X_t^{j,\varepsilon}))\,dt + \sqrt{2\beta^{-1}}\,dB_t^i$$

with W, V chosen to be 1-periodic.

Let $\rho^{\varepsilon,N} = \mathrm{Law}(X_t^{1,\varepsilon}, \dots, X_t^{N,\varepsilon})$ and consider the diffusive rescaling

$$\rho^{\varepsilon,N}(x,t) := \varepsilon^{-Nd}\nu^N(\varepsilon^{-1}x, \varepsilon^{-2}t) \in \mathcal{P}((\mathbb{R}^d)^N)\,.$$

Interpretation: zooming out in space and looking at sufficiently long (diffusive) times.

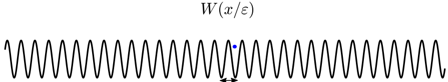

The 1d problem (Freidlin, Lifson-Jackson, 1962)

- Consider a single particle moving in a periodic potential

$$dX_t = -\sin(X_t)\, dt + \sqrt{2\beta^{-1}}\, dW_t$$

- From the martingale central limit theorem it follows that the rescaled process converges weakly to a Brownian motion

$$\varepsilon X_{t/\varepsilon^2} \to \sqrt{2D_\beta}\, W_t$$

- where $D_\beta = \frac{\beta^{-1}}{Z_+ \, Z_-}, \quad Z_\pm = \frac{1}{2\pi} \int_0^{2\pi} e^{\mp \cos(x)}\, dx.$
- This formula was obtained by Lifson-Jackson (J. Chem. Phys., 1962) by doing a mean exit time calculation.
- Similar result in the multidimensional case. Upper and lower bounds on the covariance matrix of the effective Brownian motion $\frac{\beta^{-1}}{Z_+ \, Z_-} \|\xi\|^2 \leq \langle D_\beta \xi, \xi \rangle \leq \beta^{-1} \|\xi\|^2, \quad \forall\, \xi \in \mathbb{R}^d.$

The quotiented N-particle system

$$\mathrm{d}\dot{X}_t^i = -\nabla V(\dot{X}_t^i)\,dt - \frac{1}{N}\sum_{i\neq j}^{N} \nabla W(\dot{X}_t^i - \dot{X}_t^j)\,dt + \sqrt{2\beta^{-1}}\,d\dot{B}_t^i\,,$$

$\dot{X}_t^i \in \mathbb{T}^d$ and $\dot{B}_t^i$ are $\mathbb{T}^d$-valued Wiener processes.

This is a reversible process with respect to the N-particle Gibbs measure

$$M_N(x) = \frac{e^{-H^N(x)}}{\int\limits_{\mathbb{T}^{dN}} e^{-H^N(y)}\ \mathrm{d}y}\,,$$

and the law $\widetilde{\nu}^N$ evolves according to

$$\begin{cases} \partial_t \widetilde{\nu}^N & = \beta^{-1}\Delta\widetilde{\nu}^N + \nabla\cdot(\nabla H^N \widetilde{\nu}^N), \quad (t,x)\in(0,\infty)\times(\mathbb{T}^d)^N \\ \widetilde{\nu}^N(0) & = \widetilde{\nu}_0^N := \sum_{k\in\mathbb{Z}^d} \nu_0^N(k+x) \in \mathcal{P}((\mathbb{T}^d)^N) \end{cases}$$

Periodic rearrangement of ν^N.

Theorem (The diffusive limit)

$$\rho^{N,*} = \lim_{\varepsilon \to 0} \rho^{\varepsilon,N}$$

exists (with convergence in weak-⋆) and satisfies

$$\partial_t \rho^{N,*} = \nabla \cdot (A^{\mathrm{eff},N} \nabla \rho^{N,*}),$$

where the diffusion matrix is given by

$$A^{\mathrm{eff},N} = \int_{\mathbb{T}^N} (I + \nabla \Psi^N(y))\, M_N(y)\, dy$$

with M_N *the invariant measure of the quotiented process on the torus, and* Ψ^N *the solution of the corrector problem*

$$\nabla \cdot (\nabla \Psi^N M^N) = -\nabla M^N, \quad x \in \mathbb{T}^N.$$

The diffusive limit is affected by the problem on the torus, that exhibits phase transitions in the mean field limit.

- Question: do the mean field and homogenization limits commute?

$$
\begin{tikzcd}
\rho^{\varepsilon,N} \arrow[rr, bend left, dash, "\varepsilon \to 0"] \arrow[d, bend right, dash, "N \to \infty"'] & & \rho^{N,*} \arrow[d, bend left, dash, "N \to \infty"] \\
\rho^{\varepsilon,\otimes N} \arrow[rr, bend right, dash, "\varepsilon \to 0"'] & & ?
\end{tikzcd}
$$

Theorem (Delgadino–Gvalani–P. ’21)

Let $\mathcal{F}_\beta$ be the free energy on the torus and assume it exhibits a phase transition at some $\beta = \beta_c$. Then for $\beta < \beta_c$

$$\lim_{N\to\infty}\lim_{\varepsilon\to 0}\rho^{\varepsilon,N} = \lim_{\varepsilon\to 0}\lim_{N\to\infty}\rho^{\varepsilon,N}\,.$$

On the other hand if $\beta > \beta_c$, there exists initial data $\rho_0^{\otimes N}$ such that

$$\lim_{N\to\infty}\lim_{\varepsilon\to 0}\rho^{\varepsilon,N} \neq \lim_{\varepsilon\to 0}\lim_{N\to\infty}\rho^{\varepsilon,N}\,.$$

Consider the quotiented Kuramoto model with a confining potential $V(x) = \cos(2\pi x)$. Then there exists a $\beta = \beta_c$ such that:

- For $\beta < \beta_c$, there exists a unique steady state given by

$$\widetilde{\nu}^{\min}(x) = Z_{\min}^{-1} e^{a^{\min}\cos(2\pi x)}\,, \qquad Z_{\min} = \int_{\mathbb{T}} e^{a^{\min}\cos(2\pi x)}\, \mathrm{d}x\,,$$

  for some $a^{\min} = a^{\min}(\beta), a^{\min} > 0$, which is the unique minimiser of the periodic mean field energy $\widetilde{E}_{MF}$.

- For $\beta > \beta_c$, there exist at least 2 steady states given by

$$\widetilde{\nu}^{\min}(x) = Z_{\min}^{-1} e^{a^{\min}\cos(2\pi x)}\,, \qquad Z_{\min} = \int_{\mathbb{T}} e^{a^{\min}\cos(2\pi x)}\, \mathrm{d}x\,,$$
$$\widetilde{\nu}^{*}(x) = Z_{*}^{-1} e^{a^{*}\cos(2\pi x)}\,, \qquad Z_{*}^{-1} = \int_{\mathbb{T}} e^{a^{*}\cos(2\pi x)}\, \mathrm{d}x\,,$$

  where $a^* < 0 < a^{\min}$ and both constants depend on β. Here $\widetilde{\nu}^{\min}$ is the unique minimiser and $\widetilde{\nu}^*$ is a non-minimising critical points of the periodic mean field energy $\widetilde{E}_{MF}$. Moreover, $a^* \neq -a^{\min}$.

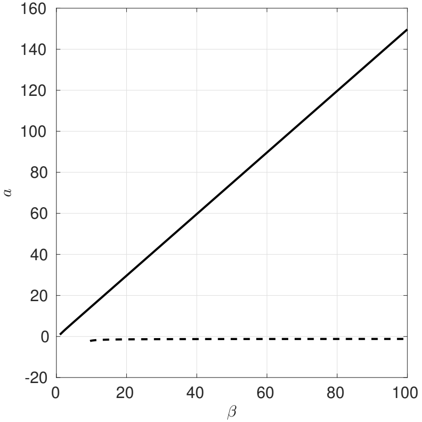

Figure: $a^{\min}$ (solid line) and a^* (dotted line) for $\eta = 0.5$. The two effective diffusion coefficients are $A_*^{\text{eff}} = \frac{\beta^{-1}}{I_0(-a^*)^2}$ and $A_{\min}^{\text{eff}} = \frac{\beta^{-1}}{I_0(a^{\min})^2}$.

Sketch of proof for $N \to \infty$ followed by $\varepsilon \to 0$

- Pass to the mean field limit to obtain $\boldsymbol{X}^\varepsilon(t)$.
- For the associated mean field SDE on the torus consider a moving corrector problem:

$$\nabla \cdot (\widetilde{\mu}^\varepsilon(t) \nabla \chi) = -\nabla(\widetilde{\mu}^\varepsilon)\,, \quad \widetilde{\mu}^\varepsilon(t) \sim \exp(-\beta(W \star \widetilde{\nu}(t) - V))$$

  and obtain time-dependent estimates:

$$\|\chi_i\|_{C^m(\mathbb{T}^d)} \lesssim 1$$
$$\|\partial_t \chi_i\|_{C^m(\mathbb{T}^d)} \lesssim \sum_{m=1}^{k} d_2^m(\widetilde{\nu}(t), \widetilde{\nu}^*)\,.$$

- Using coupling techniques (a' la Eberle et al.) prove an initial data dependent version of the martingale CLT.
- Pass to the limit as $\varepsilon \to 0$.

Sketch of proof for $\varepsilon \to 0$ followed by $N \to \infty$

- Need to pass to the limit in the diffusion matrix $A^{\text{eff,N}}$:

$$A^{\text{eff},N} = \beta^{-1} \int_{(\mathbb{T}^d)^N} (I + \nabla \Psi^N(y))\ M_N(y)\ \mathrm{d}y\,.$$

- Key idea $M_N \approx M_{N-1}(M_N)_1$ as $N \to \infty$ + natural uniform in N estimate on Ψ^N:

$$\begin{aligned}
&\int_{(\mathbb{T}^d)^N} (I + \nabla \Psi^N(y))\ \ (M_N - M_{N-1}(M_N)_1)\ \mathrm{d}y \\
&= \int_{(\mathbb{T}^d)^N} (I + \nabla \Psi^N(y))\ \left(\frac{M_N}{M_{N-1}} - (M_N)_1 \right) M_{N-1}\ \mathrm{d}y \\
&\leq \left\| I + \nabla \Psi^N \right\|_{L^2(M_{N-1})} \left\| \left(\frac{M_N}{M_{N-1}} - (M_N)_1 \right) \right\|_{L^2(M_{N-1})} .
\end{aligned}$$

- The function $M_N/(M_{N-1})$ is symmetric in all but one of its variables. Use techniques due to Lions pass to $N \to \infty$ on $C(\mathcal{P}(\mathbb{T}^d))$. Similarly pass to $N \to \infty$ to obtain $M_{N-1} \to \delta_{\tilde{\nu}^{\min}} \in \mathcal{P}(\mathcal{P}(\mathbb{T})^d)$.
- Enough information to pass to the limit in the PDE.

Consider the fluctuations around the mean field for the Kuramoto model $\mathcal{G}^N := \sqrt{N}\left(\frac{1}{N}\sum_{i=1}^{N}\delta_{x_i} - \widetilde{\nu}^{\min}\right)$. We have that $\mathcal{G}^N$ converges in law to $\mathcal{G}^\infty$ whose law is the unique invariant measure of the following linear stochastic PDE

$$d\mathcal{G}^\infty = \left(\beta^{-1}\partial_{xx}\mathcal{G}^\infty + (2\pi)^2\cos(2\pi x) * \mathcal{G}^\infty\right)\mathrm{d}t + \sqrt{2\beta^{-1}}d\xi,$$

where $\xi(t,x) = \sum_{k\in\mathbb{Z}} 2\pi k e_k(x)\dot{B}_k(t)$. We can find the invariant measure explicitly for each mode

$$\mathrm{Law}(\langle\mathcal{G}^\infty, e_k\rangle) = \begin{cases} \mathcal{N}\left(0, \frac{2}{2-\beta}\right) & |\mathrm{k}|=1 \\ \mathcal{N}(0,1) & |\mathrm{k}|\neq 1, \end{cases}$$

We can clearly identify the phase transition $\beta_c = 2$ when the SPDE no longer supports an invariant measure.

PHASE TRANSITIONS, LOGARITHMIC SOBOLEV INEQUALITIES, AND UNIFORM-IN-TIME PROPAGATION OF CHAOS FOR WEAKLY INTERACTING DIFFUSIONS

BY MATÍAS G. DELGADINO[1], RISHABH S. GVALANI[2], GRIGORIOS A. PAVLIOTIS[3] AND SCOTT A. SMITH[4]

[1] *Department of Mathematics, The University of Texas at Austin* *matias.delgadino@math.utexas.edu*

[2] *Max-Planck-Institut für Mathematik in den Naturwissenschaften* *gvalani@mis.mpg.de*

[3] *Department of Mathematics, Imperial College London* *g.pavliotis@imperial.ac.uk*

[4] *Academy of Mathematics and Systems Sciences, Chinese Academy of Sciences* *ssmith@amss.ac.cn*

Abstract In this article, we study the mean field limit of weakly interacting diffusions for confining and interaction potentials that are not necessarily convex. We explore the relationship between the large N limit of the constant in the logarithmic Sobolev inequality (LSI) for the N-particle system and the presence or absence of phase transitions for the mean field limit. The non-degeneracy of the LSI constant is shown to have far reaching consequences,

We consider $\{X_t^i\}_{i=1,\ldots,N} \subset \mathbb{R}^d$, the positions of N indistinguishable interacting particles at time $t \geq 0$, satisfying the following system of SDEs:

$$\begin{cases} dX_t^i = -\nabla V(X_t^i)\,\mathrm{d}t - \dfrac{1}{N}\displaystyle\sum_{j=1}^{N} \nabla_1 W(X_t^i, X_t^j)\,\mathrm{d}t + \sqrt{2\beta^{-1}}dB_t^i \\ \mathrm{Law}(X_0^1, \ldots, X_0^N) = \rho_{\mathrm{in}}^{\otimes N} \in \mathcal{P}_{2,\mathrm{sym}}((\mathbb{R}^d)^N), \end{cases} \tag{0.1}$$

where $V : \mathbb{R}^d \to \mathbb{R}$, $W : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$, $\beta^{-1} > 0$ is the inverse temperature, $B_t^i, i = 1, \ldots, N$ are independent d-dimensional Brownian motions, and the initial position of the particles is i.i.d with law ρ_{in}.

Assumption

The confining potential V is lower semicontinuous, bounded below, K_V-convex for some $K_V \in \mathbb{R}$ and there exists $R_0 > 0$ and $\delta > 0$, such that $V(x) \geq |x|^\delta$ for $|x| > R_0$. The interaction potential W is lower semicontinuous, K_W-convex for some $K_W \in \mathbb{R}$, bounded below, symmetric $W(x,y) = W(y,x)$, vanishes along the diagonal $W(x,x) = 0$, and there exists C such that

$$|\nabla_1 W(x,y)| \leq C(1 + |W(x,y)| + V(x) + V(y))$$

- Given $\rho(t)$, the solution to the McKean-Vlasov PDE, we differentiate to obtain the dissipation

$$\frac{\mathrm{d}}{\mathrm{d}t}E^{MF}[\rho(t)]-\inf E^{MF} = -\int_\Omega |\beta^{-1}\nabla \log \rho(t)+\nabla W \star \rho(t)+\nabla V|^2 \rho(t)\ \mathrm{d}$$

- We have exponentially fast convergence provided that the **infinite volume log Sobolev constant**

$$0 < \lambda_{\mathrm{LSI}}^{\infty} := \inf_{\substack{\rho\in\mathcal{P}(\Omega)\\ \rho\notin\mathcal{K}}} \frac{D(\rho)}{E^{MF}[\rho(t)] - \inf E^{MF}},$$

is positive, where

$$\mathcal{K} = \{\rho \in \mathcal{P}(\Omega)\ :\ E^{MF}[\rho] = \inf E^{MF}\}.$$

Theorem

Under the above Assumption, we have

$$\limsup_{N\to\infty} \lambda_{\mathrm{LSI}}^N \leq \lambda_{\mathrm{LSI}}^\infty.$$

Moreover, if the mean field energy $E^{MF} : \mathcal{P}(\Omega) \to \mathbb{R} \cup \{+\infty\}$ *given by*

$$E^{MF}[\rho] := \beta^{-1} \int_\Omega \rho \log(\rho)\ \mathrm{d}x + \frac{1}{2} \int_{\Omega^2} W(x,y)\ \mathrm{d}\rho(x)\ \mathrm{d}\rho(y) + \int_\Omega V(x)\ \mathrm{d}\rho(x)$$

admits a critical point that is not a minimiser, then $\lambda_{\mathrm{LSI}}^\infty = 0$*, and there exists* $C > 0$ *such that*

$$\lambda_{\mathrm{LSI}}^N \leq \frac{C}{N}.$$

When $\lambda_{\mathrm{LSI}}^{\infty} > 0$, we can show that the **regularized log Sobolev constant**

$$\lambda_{\mathrm{LSI}}^{N,\varepsilon} := \inf_{\rho^N \colon \overline{\mathcal{E}}(\rho^N | M_N) > \varepsilon} \frac{\beta^{-1} \overline{\mathcal{I}}(\rho^N | M_N)}{\overline{\mathcal{E}}(\rho^N | M_N)}$$

does not degenerate:

$$\lim_{N \to \infty} \lambda_{\mathrm{LSI}}^{N,\varepsilon} \geq \lambda_{\mathrm{LSI}}^{\infty} > 0.$$

Theorem

Under the above assumption, assume that $\lambda_{\mathrm{LSI}}^{\infty} > 0$, and that ρ_{in} in (0.1) has finite energy and bounded higher order moments,

$$E^{MF}[\rho_{\mathrm{in}}] < \infty \qquad \text{and} \qquad \int_{\Omega} |x|^{2+\delta} \, \mathrm{d}\rho_{\mathrm{in}} < \infty, \qquad \text{for some } \delta > 0.$$

Then, for every $\varepsilon > 0$, there exists $N_0 \in \mathbb{N}$, such that for every $N > N_0$ we have

- We are not able to fully characterise the limit of $\lambda_{\mathrm{LSI}}^{N}$ in terms of the mean field limit. We formulate the following conjecture.
  - Under the above assumptions, we have the equality

$$\lim_{N\to\infty} \lambda_{\mathrm{LSI}}^{N} = \lambda_{\mathrm{LSI}}^{\infty}.$$

- These results suggest that:
  a. the absence of phase transitions
  b. the non-degeneracy of the infinite volume log Sobolev constant
  c. the validity uniform-in-time propagation of chaos
- are all equivalent

Theorem

Under the above Assumption, assume that $\limsup_{N\to\infty} \lambda_{\mathrm{LSI}}^N > 0$*. Then, there exists a unique steady state* ρ_β *to the McKean-Vlasov PDE. Moreover, there exists* $C > 0$*, such that*

$$\overline{d}_2^2(\rho_\beta^{\otimes N}, M_N) \leq \frac{2}{\lambda_{\mathrm{LSI}}^N}\overline{\mathcal{E}}(\rho_\beta^{\otimes N}|M_N) \leq \frac{2}{(\lambda_{\mathrm{LSI}}^N)^2}\overline{\mathcal{I}}(\rho_\beta^{\otimes N}|M_N) \leq \frac{C}{N}.$$

Theorem

Under the above assumption, let ρ^N *and* ρ *denote the unique solutions to the particle and mean field dynamics. Assume that* ρ_{in} *has finite energy* $E^{MF}[\rho_{\mathrm{in}}] < \infty$, *that the gradient of the square of the interaction potential is uniformly integrable*
$\sup_{t\in[0,\infty]} \int_\Omega |\nabla_1 W|^2 \star \rho(t)\rho(t)\ \mathrm{d}x < \infty$, *and that*
$\liminf_{N\to\infty} \lambda_{\mathrm{LSI}}^N =: \lambda^\infty > 0$*. Then,*
$\overline{d}_2(\rho^N(t), \rho^{\otimes N}(t)) \leq \frac{C}{N^\theta}$ *for all* $t > 0$*, where*

$$\theta = \begin{cases} 1/2 & \text{if } K_V + K_W > 0 \\ \frac{1}{2}\frac{\lambda^\infty}{\lambda^\infty - 2(K_V + K_W)} & \text{if } K_V + K_W < 0 \end{cases}$$

with K_V *and* K_W *the convexity constants of* V *and* W *in A1,A2. In the case where* $K_V + K_W(1 - 1/N) = 0$*, we can pick any* $\theta < 1/2$*.*

Inference for the McKean-SDE

Parameter Estimation for the McKean-Vlasov Stochastic Differential Equation *

L. Sharrock†, N. Kantas†, P. Parpas‡, and G.A. Pavliotis†

Abstract. We consider the problem of parameter estimation for a stochastic McKean-Vlasov equation, and the associated system of weakly interacting particles. We first establish consistency and asymptotic normality of the offline maximum likelihood estimator for the interacting particle system in the limit as the number of particles $N \to \infty$. We then propose an online estimator for the parameters of the McKean-Vlasov SDE, which evolves according to a continuous-time stochastic gradient descent algorithm on the asymptotic log-likelihood of the interacting particle system. We prove that this estimator converges in $\mathbb{L}^1$ to the stationary points of the asymptotic log-likelihood of the McKean-Vlasov SDE in the joint limit as $N \to \infty$ and $t \to \infty$, under suitable conditions which guarantee ergodicity and uniform-in-time propagation of chaos. We then demonstrate, under the additional condition of global strong concavity, that our estimator converges in $\mathbb{L}^2$ to the unique maximiser of this asymptotic log-likelihood function, and establish an $\mathbb{L}^2$ convergence rate. We also obtain analogous results under the condition that, rather than observing multiple trajectories of the interacting particle system, we instead observe multiple independent replicates of the McKean-Vlasov SDE itself or, less realistically, a single sample path of the McKean-Vlasov SDE and its law. Our theoretical results are demonstrated via two numerical examples, a linear mean field model and a stochastic opinion dynamics model.

Key words. McKean-Vlasov equation, nonlinear diffusion, maximum likelihood, parameter estimation, consistency, asymptotic normality, stochastic gradient descent

AMS subject classifications. 60F05, 60F25, 60H10, 62F12

1. Introduction. In this paper, we consider a family of McKean-Vlasov stochastic differential equations (SDEs) on $\mathbb{R}^d$, parametrised by $\theta \in \mathbb{R}^p$, of the form

$$\mathrm{d}x_t^\theta = B(\theta, x_t^\theta, \mu_t^\theta)\mathrm{d}t + \sigma(x_t^\theta)\mathrm{d}w_t, \quad t \geq 0 \tag{1.1}$$

$$\mu_t^\theta = \mathcal{L}(x_t^\theta), \tag{1.2}$$

where $B : \mathbb{R}^p \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}^d$, $\sigma : \mathbb{R}^d \to \mathbb{R}^{d\times d}$ are Borel measurable functions, $(w_t)_{t\geq 0}$ is a $\mathbb{R}^d$-valued standard Brownian motion, and $\mathcal{L}(x_t^\theta)$ denotes the law of of x_t^θ. We assume that $x_0 \in \mathbb{R}^d$, or that x_0 is a $\mathbb{R}^d$-valued random variable with law μ_0, independent of $(w_t)_{t\geq 0}$. This equation is non-linear in the sense of McKean [60, 61, 79]; in particular, the coefficients depend on the law of the solution, in addition to the solution itself. We will restrict our attention to the case in which the dependence on the law only enters linearly in the drift, namely, that

$$B(\theta, x, \mu) = b(\theta, x) + \int_{\mathbb{R}^d} \phi(\theta, x, y)\mu(\mathrm{d}y), \tag{1.3}$$

*Funding: The first author was funded by the EPRSC CDT in the Mathematics of Planet Earth (grant number EP/L016613/1) and the National Physical Laboratory. The second, third, and fourth authors were partially funded under a J.P. Morgan A.I. Research Award (2022). The fourth author was partially supported by the EPSRC (grant number EP/P031587/1).

†Department of Mathematics, Imperial College London, South Kensington, London, SW7 2AZ, UK (louis.sharrock16@imperial.ac.uk, n.kantas@imperial.ac.uk, g.pavliotis@imperial.ac.uk)

‡Department of Computing, Imperial College London, South Kensington, London, SW7 2AZ, UK (panos.parpas@imperial.ac.uk)

1

Eigenfunction martingale estimators for interacting particle systems and their mean field limit

Grigorios A. Pavliotis * Andrea Zanoni †

Abstract

We study the problem of parameter estimation for large exchangeable interacting particle systems when a sample of discrete observations from a single particle is known. We propose a novel method based on martingale estimating functions constructed by employing the eigenvalues and eigenfunctions of the generator of the mean field limit, linearized around the (unique) invariant measure of the mean field dynamics. We then prove that our estimator is asymptotically unbiased and asymptotically normal when the number of observations and the number of particles tend to infinity, and we provide a rate of convergence towards the exact value of the parameters. Finally, we present several numerical experiments which show the accuracy of our estimator and corroborate our theoretical findings, even in the case the mean field dynamics exhibit more than one steady states.

AMS subject classifications. 35Q70, 35Q83, 60J60, 62M15, 65C30.

Key words. Interacting particle systems, exchangeability, mean field limit, inference, Fokker–Planck operator, eigenvalue problem, martingale estimators.

1 Introduction

Interacting particle systems and, more generally interacting multiagent models, appear frequently in the natural and social sciences. In addition to the well known applications, e.g., plasma physics [22] and stellar dynamics [7], new applications include, e.g., the modeling of chemotaxis [40], pedestrian dynamics [24,30], crowd dynamics [32], urban modeling [14], models for opinion formation [18,21], collective behavior [11], and models for systemic risk [20]. In many of these applications, the phenomenological models involve unknown parameters that need to be estimated from data. This is particularly the case for multiagent models used in the social sciences and in economics, where no physics-informed choices of parameters are available. Learning parameters or even models, in a nonparametric setting, from data is becoming an increasingly important aspect of the overall mathematical modeling strategy. This is particularly the case in view of the huge quantity of available data in different areas, which allows the development of accurate data-driven techniques for learning parameters from data.

*Department of Mathematics, Imperial College London, London SW7 2AZ, UK, g.pavliotis@imperial.ac.uk

†Institute of Mathematics, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland, andrea.zanoni@epfl.ch

1

Data generation process

Three cases for model hypothesis. We observe data as either:

1. N dependent paths generated from IPS, $\left\{(x_s^{\theta^*,i})_{s\in[0,t]}\right\}_{i=1}^{N}$
2. N independent paths generated from McKean SDE, $\left\{(x_s^{\theta^*,[i]})_{s\in[0,t]}\right\}_{i=1}^{N}$
3. A single path from McKean SDE and the full evolution $u_t^{\theta^*}(x)$ of the density for all x

Case 3 not practically relevant - almost standard SDE estimation case

In each case we wish to estimate true parameter θ^* used to generate the data

Data generation process

θ^* is true parameter. Data can be generated by different models.

Cases for $s \in [0, t]$:

1. Paths of particles of IPS:

$$\mathrm{d}x_s^{\theta^*,i,N} = B(\theta^*, x_s^{\theta^*,i,N}, \mu_s^{\theta,N})ds + \sigma \mathrm{d}w_s^{i,N}$$

2. N independent paths of McKean SDE:

$$\mathrm{d}x_s^{\theta^*,[i]} = B(\theta^*, x_s^{\theta^*,[i]}, \mu_s^{\theta})\mathrm{d}s + \sigma \mathrm{d}w_s^{[i]}$$

3. One path of McKean SDE and the full evolution of the density for all x:

$$\mathrm{d}x_s^{\theta^*} = B(\theta^*, x_s^{\theta^*}, u_s^{\theta})\mathrm{d}s + \sigma \mathrm{d}w_s$$
$$\partial_s u_s^{\theta^*}(x) = \nabla \left[\frac{1}{2}\sigma\sigma^T \nabla u_s^{\theta^*}(x) + u_s^{\theta^*}(x) B(\theta^*, x, u_s^{\theta})\right]$$

- Notation:

$$B(\theta, x, \mu) = b(\theta, x) + \int_{\mathbb{R}^d} \phi(\theta, x, y)\mu(dy)$$

Parameter estimation using likelihood inference

- Ideal Girsanov likelihood for McKean SDE

$$\begin{aligned}\mathcal{L}_t(\theta; X) &= \log \frac{\mathrm{d}\mathbb{P}_t^{\theta}}{\mathrm{d}\mathbb{P}_t^{\theta*}}(X) \\ &= \int_0^t L(\theta, x_s, \mu_s)\mathrm{d}s + \int_0^t \langle G(\theta, x_s, \mu_s), \mathrm{d}w_s \rangle.\end{aligned}$$

with

$$\begin{aligned}B(\theta, x, \mu) &= b(\theta, x) + \int_{\mathbb{R}^d} \phi(\theta, x, y)\mu(\mathrm{d}y) \\ G(\theta, x, \mu) &= B(\theta, x, \mu) - B(\theta^*, x, \mu) \\ L(\theta, x, \mu) &= -\frac{1}{2}||G(\theta, x, \mu)||^2\end{aligned}$$

Likelihood approximations and ML estimation

- Cases 1-2 require particle approximations:

$$\mathcal{L}_t^N(\theta) = \frac{1}{N}\sum_{i=1}^{N} \underbrace{\left[\int_0^t L(\theta, x_s^{i,N}, \mu_s^N)\mathrm{d}s + \int_0^t \left\langle G(\theta, x_s^{i,N}, \mu_s^N), \mathrm{d}w_s^{i,N}\right\rangle\right]}_{=\mathcal{L}_t^{i,N}(\theta)},$$

  with

$$\mu_t^N = \frac{1}{N}\sum_{j=1}^{N} \delta_{x_t^{j,N}}$$

  Resulting approximate MLE:

$$\hat{\theta}_t^N = \arg\sup_{\theta\in\mathbb{R}^p} \mathcal{L}_t^N(\theta).$$

- Case 3 no need for approximation

$$\hat{\theta}_t = \arg\sup_{\theta\in\mathbb{R}^p} \mathcal{L}_t(\theta).$$

MLE convergence: the offline case

- Offline MLE:
  - observations after a fixed time interval $[0,t]$: $\left\{(x_s^i)_{s\in[0,t]}\right\}_{i=1}^N$ (dropping θ^* superscripts above)
  - estimate θ^* using $\hat{\theta}_t$ or $\hat{\theta}_t^N$
- There are different asymptotic regimes
  - $t \to \infty$: $\hat{\theta}_t$ well understood [Levanony et al 94, Wen et al 16] resp.
  - $N \to \infty$: (fixed t) previously established with linear parametrisation in θ [Kasonga 90]
- Weak consistency and asymptotic normality holds

Offline parameter estimation - convergence

Theorem

Under the above and some more identifiability assumptions, let $\Theta \subseteq \mathbb{R}^p$ *be a compact set and* $\theta^* \in \Theta$*. Then for any fixed* t *as* $N \to \infty$*, we have*

$$\hat{\theta}_t^N \xrightarrow{\mathbb{P}} \theta^*$$

and

$$N^{\frac{1}{2}}(\hat{\theta}_t^N - \theta^*) \xrightarrow{\mathcal{D}} \mathcal{N}(0, I_t^{-1}(\theta^*))$$

with

$$[I_t(\theta^*)]_{kl} = \int_0^t \int_{\mathbb{R}^d} [\nabla_\theta B(\theta^*, x, \mu_s)]_k [\nabla_\theta B(\theta^*, x, \mu_s)]_l \mu_s(\mathrm{d}x)\mathrm{d}s.$$

Bistable Potential: Offline MLE

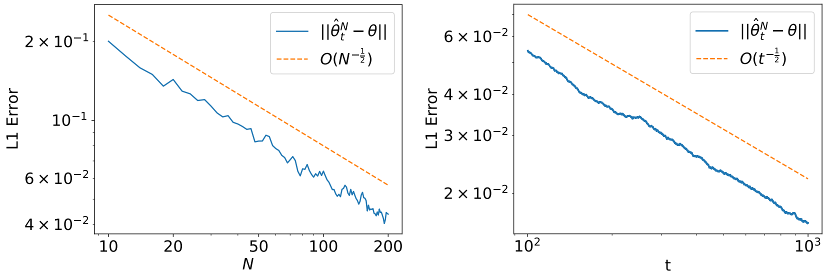

(a) $\hat{\theta}^N_{1,t=5}$.

(b) $\hat{\theta}^{N=2}_{2,t}$.

Figure: Log-log plot of the $\mathbb{L}^1$ error of the offline MLE for $t = 0.5$ and $N \in \{20, \ldots, 200\}$ (left hand panel), and for $t \in [100, 1000]$ and $N = 2$ (right hand panel).

Online parameter estimation

Ideal approach (case 3):

- seek to maximise ergodic likelihood

$$\begin{aligned}\widetilde{\mathcal{L}}(\theta) &= \lim_{t\to\infty} \frac{1}{t}\mathcal{L}_t(\theta) \\ &= -\frac{1}{2}\int_{\mathbb{R}^d} ||B(\theta,x,\mu_\infty) - B(\theta^*,x,\mu_\infty)||^2 \mu_\infty(\mathrm{d}x).\end{aligned}$$

- use stochastic gradient descent

$$\mathrm{d}\theta_t = \gamma_t\left(\underbrace{\nabla_\theta L(\theta_t,x_t,\mu_t)\mathrm{d}t}_{\text{(noisy) ascent term}} + \underbrace{\nabla_\theta B(\theta_t,x_t,\mu_t)\mathrm{d}w_t}_{\text{noise term}}\right)$$

with γ_t decreasing step size

- A more clear view

$$\mathrm{d}\theta_t = \gamma_t(\underbrace{\nabla_\theta \widetilde{\mathcal{L}}(\theta_t)\mathrm{d}t}_{\text{(true) ascent term}} + \underbrace{(\nabla_\theta L(\theta_t,x_t,\mu_t) - \nabla_\theta \widetilde{\mathcal{L}}(\theta_t))\mathrm{d}t}_{\text{fluctuations term}} + \underbrace{\nabla_\theta B(\theta_t,x_t,\mu_t)\mathrm{d}w_t}_{\text{noise term}})$$

Online parameter estimation

- Implementable approach uses particle approximations for likelihood terms:

$$\widetilde{\mathcal{L}}^{i,N}(\theta) = \lim_{t\to\infty} \frac{1}{t}\mathcal{L}_t^{i,N}(\theta) \text{ or } \widetilde{\mathcal{L}}^{N}(\theta) = \lim_{t\to\infty} \frac{1}{t}\mathcal{L}_t^{N}(\theta).$$

- At each particle use:

$$\mathrm{d}\theta_t^{i,N} = \gamma_t \left[\nabla_\theta L(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\mathrm{d}t + \nabla_\theta B(\theta_t^{i,N}, x_t^{i,N}, \mu_t^N)\mathrm{d}w_t^i\right]$$

  and average

$$\theta_t^N = \frac{1}{N}\sum_{i=1}^{N}\theta_t^{i,N}.$$

- Cases 1 and 2 are identical up to specification of the data.

Online parameter estimation convergence

Theorem

Under assumptions above and some more, we have in $\mathbb{L}^1$ (for both Cases 1,2)

$$\lim_{N\to\infty}\lim_{t\to\infty} ||\nabla_\theta \widetilde{\mathcal{L}}(\theta_t^N)|| = 0,$$

$$\lim_{N\to\infty}\lim_{t\to\infty} ||\nabla_\theta \widetilde{\mathcal{L}}(\theta_t^{i,N})|| = 0,$$

and for Case 3

$$\lim_{t\to\infty} ||\nabla_\theta \widetilde{\mathcal{L}}(\theta_t)|| = 0,$$

For Case 1 if N is finite:

$$\lim_{t\to\infty} ||\nabla_\theta \widetilde{\mathcal{L}}^{i,N}(\theta_t^{i,N})|| = \lim_{t\to\infty} ||\nabla_\theta \widetilde{\mathcal{L}}^{N}(\theta_t^{N})|| = 0$$

Online ML numerical results - opinion dynamics

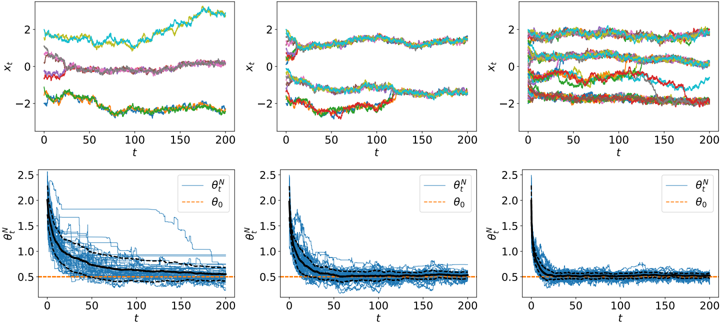

Figure: Independent runs of parameter estimates for the range parameter θ_2, $\theta_2(0) \sim \mathcal{U}([1.5, 2.5])$, and $N = \{10, 20, 50\}$ (left to right). Top: typical saple data trajectories with $\theta_2^* = 0.5$.

Online ML numerical results - opinion dynamics

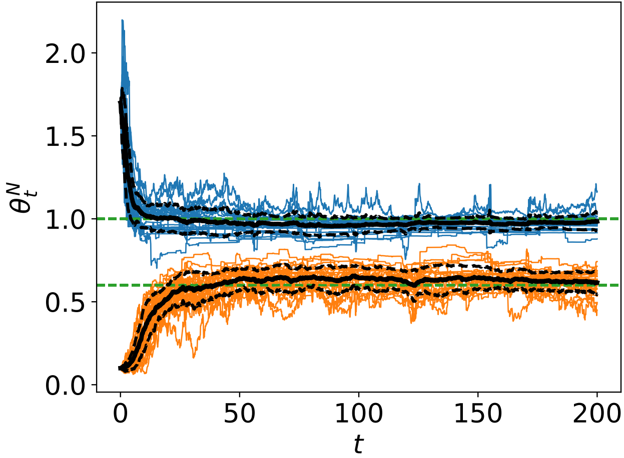

Figure: Sequence on online parameter estimates for two range parameters $\theta_{2,1}$ (blue) and $\theta_{2,2}$ (orange); interaction is a smooth approximation to $\widetilde{\varphi}_\theta(r) = \sum_{i=1}^{p} \theta_{1,i} \mathbf{1}_{r \in [0, \theta_{2,i}]}$.

Martingale Eigenfunction Estimator

Consider a large system of N interacting particles $(n = 1, \ldots, N)$

$$\mathrm{d}X_t^{(n)} = -V'(X_t^{(n)}; \alpha)\,\mathrm{d}t - \kappa\Big(X_t^{(n)} - \underbrace{\bar{X}_t^N}_{\frac{1}{N}\sum_{n=1}^N X_t^{(n)}}\Big)\,\mathrm{d}t + \sqrt{2\sigma}\,\mathrm{d}W_t^{(n)}$$

If $N \to \infty \implies$ *mean field limit* (nonlinear SDE)

$$\mathrm{d}X_t = -V'(X_t; \alpha)\,\mathrm{d}t - \kappa\Big(X_t - \underbrace{m_t}_{\mathbb{E}[X_t]}\Big)\,\mathrm{d}t + \sqrt{2\sigma}\,\mathrm{d}W_t$$

which can have **multiple** invariant measures

Goal: estimate unknown parameter $\theta = (\alpha, \kappa, \sigma)$ given $M + 1$ discrete observations of one single particle $\{X_{m\Delta}^{(n^*)}\}_{m=0}^M$

Methodology

Idea: employ martingale estimating functions based on eigenvalues and eigenfunctions of the generator of the mean field dynamics

$$\mathcal{L}_t u(x) = -\Big(V'(x;\alpha) + \kappa(x - m_t)\Big)u'(x) + \sigma u''(x)$$

Issue: the generator is time-dependent $\implies$ replace m_t with the expectation m w.r.t. the "right" invariant measure ρ

$$\mathcal{L}u(x) = -\Big(V'(x;\alpha) + \kappa(x - \underbrace{m}_{\mathbb{E}^\rho[X]})\Big)u'(x) + \sigma u''(x)$$

Procedure:

- Compute the first J eigenpairs $-\mathcal{L}\phi_j(x;\theta) = \lambda_j(\theta)\phi_j(x;\theta)$
- Construct the martingale estimating function $G^J_{M,N}(\theta)$
- Solve the nonlinear system $G^J_{M,N}(\theta) = 0 \implies \widehat{\theta}^J_{M,N}$

Algorithm

Input: Observations $\{X_{m\Delta}^{(n^*)}\}_{m=0}^{M}$
Distance between two consecutive observations Δ
Number of eigenvalues and eigenfunctions J
Smooth functions $\{\psi_j(x;\theta)\}_{j=1}^{J}$
Confining potential V

Output: Estimation $\widehat{\theta}_{M,N}^{J}$ of θ

1: Find the invariant measure ρ and compute m
2: Consider the equation $\sigma\phi''(x;\theta) - (V'(x;\alpha) + \kappa(x-m))\,\phi'(x;\theta) + \lambda(\theta)\phi(x;\theta) = 0$
3: Compute the first J eigenvalues $\{\lambda_j(\theta)\}_{j=1}^{J}$ and eigenfunctions $\{\phi_j(\cdot;\theta)\}_{j=1}^{J}$
4: Construct the function $g_j(x,y;\theta) = \psi_j(x;\theta)\left(\phi_j(y;\theta) - e^{-\lambda_j(\theta)\Delta}\phi_j(x;\theta)\right)$
5: Construct the score function $G_{M,N}^{J}(\theta) = \frac{1}{M}\sum_{m=0}^{M-1}\sum_{j=1}^{J} g_j(\widetilde{X}_m^{(n)}, \widetilde{X}_{m+1}^{(n)};\theta)$
6: Let $\widehat{\theta}_{M,N}^{J}$ be the solution of the nonlinear system $G_{M,N}^{J}(\theta) = 0$

Convergence analysis

- Asymptotic **unbiasedness**:

$$\text{if } M = o(N) \text{ then } \widehat{\theta}^J_{M,N} \to \theta$$

$$\text{in probability as } M, N \to \infty$$

- **Rate** of convergence: $\dfrac{1}{\sqrt{M}} + \dfrac{1}{\sqrt{N}}$

- Asymptotic **normality**:

$$\text{if } M = o(\sqrt{N}) \text{ then } \sqrt{M}(\widehat{\theta}^J_{M,N} - \theta) \to \mathcal{N}(0, \Gamma^J)$$

$$\text{in distribution as } M, N \to \infty$$

Numerical experiments

Consider the double well potential $V(x;\alpha) = \alpha \cdot \left(\frac{x^4}{4} \quad -\frac{x^2}{2} \right)^\top$

Issue: phase transition occurs

Plots below the phase transition:

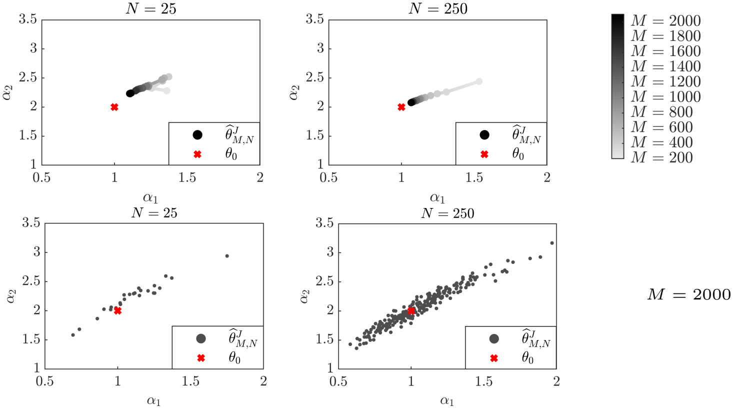

Numerical experiments

Consider the double well potential $V(x;\alpha) = \alpha \cdot \left(\frac{x^4}{4} \quad -\frac{x^2}{2} \right)^\top$

Issue: phase transition occurs

Plots above the phase transition:

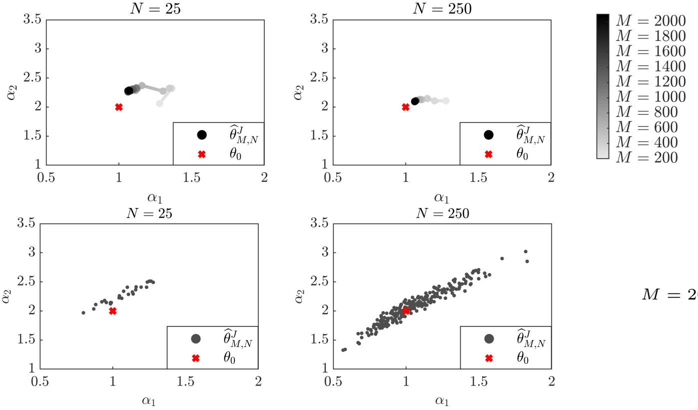

$M = 2000$

Numerical experiments

Consider the nonsymmetric potential $V(x;\alpha) = \alpha \cdot \begin{pmatrix} \frac{x^4}{4} & \frac{x^2}{2} & x \end{pmatrix}^\top$ Issue: invariant measures around each critical point of the potential