Recent Developments in Schrödinger Bridge and Optimal Density Steering

Abhishek Halder

Department of Aerospace Engineering
Iowa State University
Ames, IA 50011

Joint work with I. Nodozi, S. Haddad, K.F. Caluya (UC Santa Cruz), B. Singh (Ford Greenfield Labs), A. Mesbah (UC Berkeley)

University of Arizona, September 18, 2023

Density Steering via State Feedback

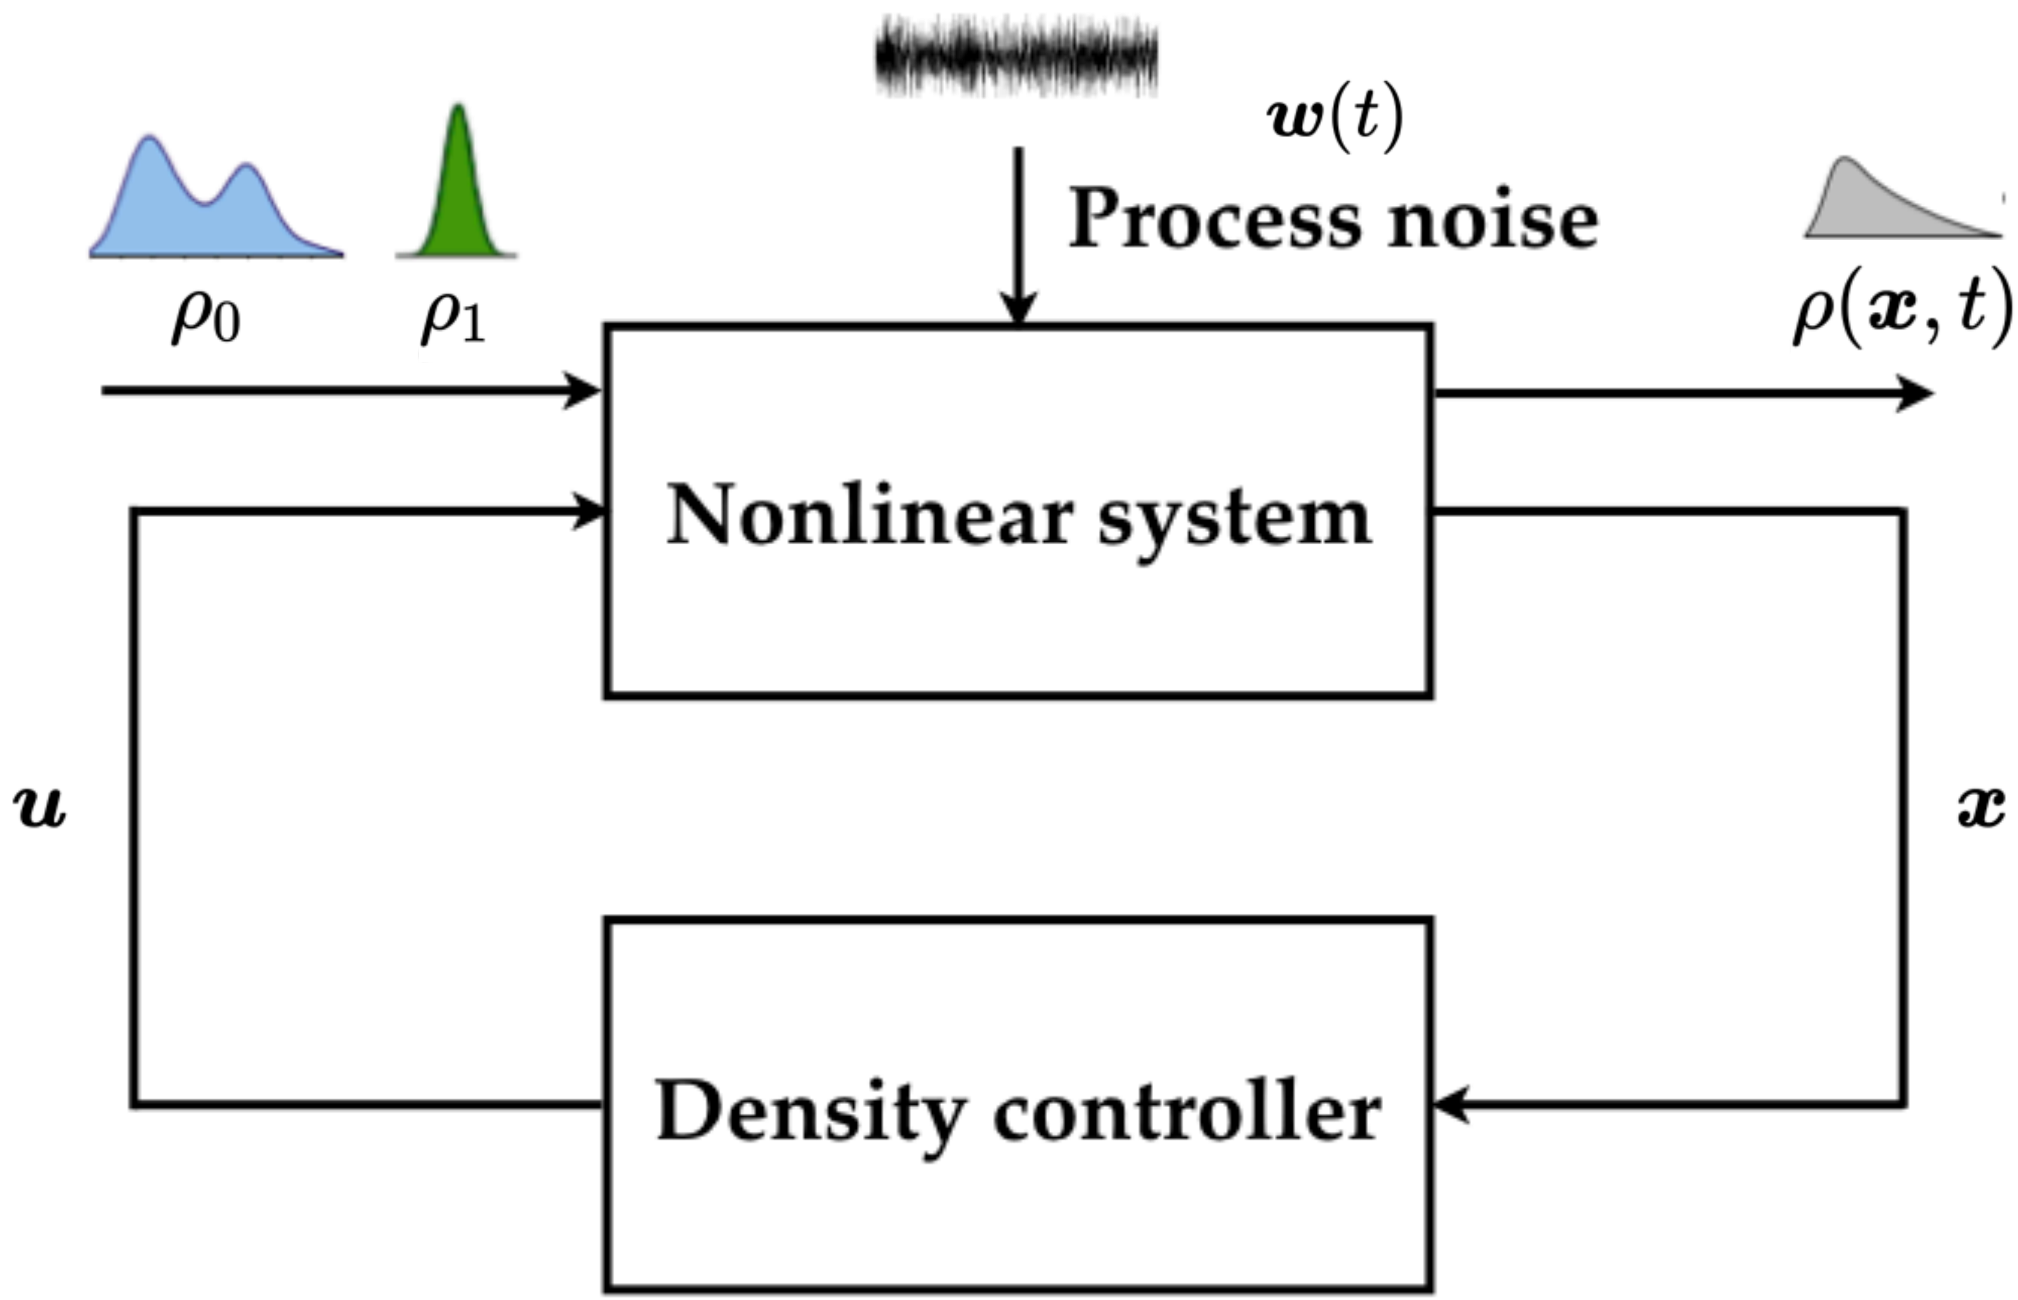

Motivating Applications

Distribution $\sim$ Probability

Distribution $\sim$ Population

Motivating Applications

Distribution ~ Probability

Spacecraft landing with desired statistical accuracy

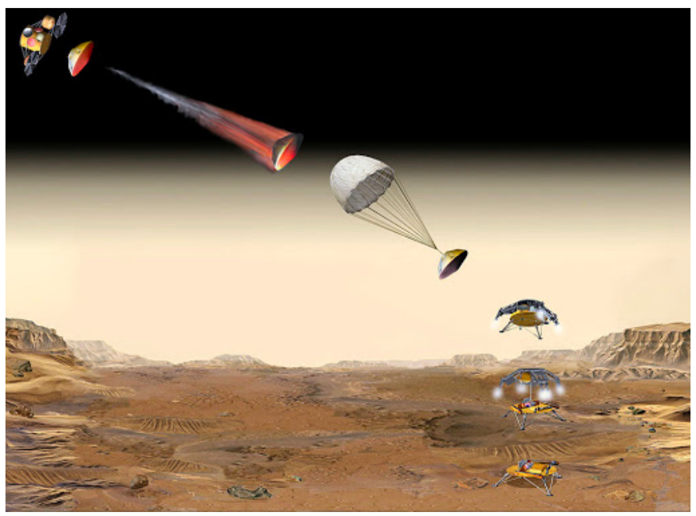

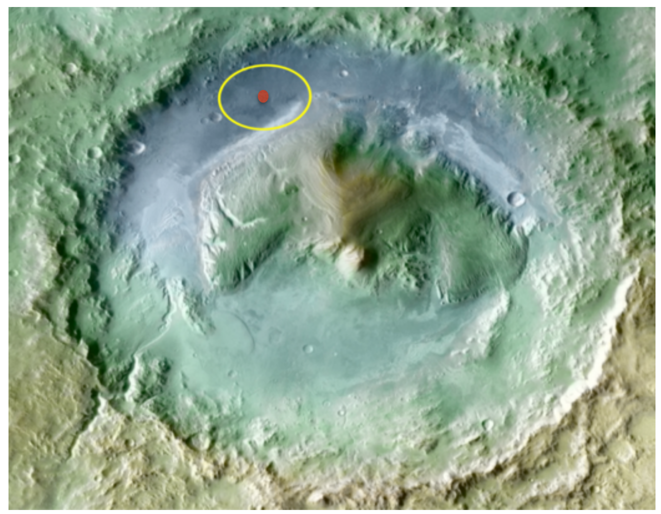

Gale Crater (4.49S, 137.42E)

Risk management for automated driving in multi-lane highways

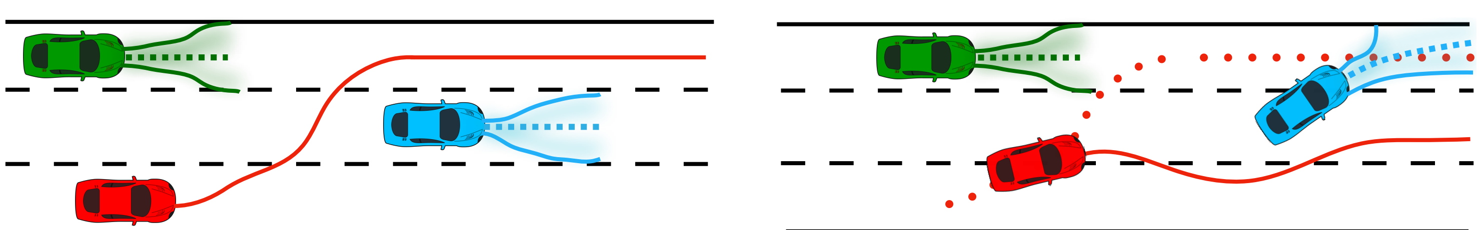

Control of uncertainties

Distribution ~ Population

Motivating Applications

Distribution ~ Probability

Spacecraft landing with desired statistical accuracy

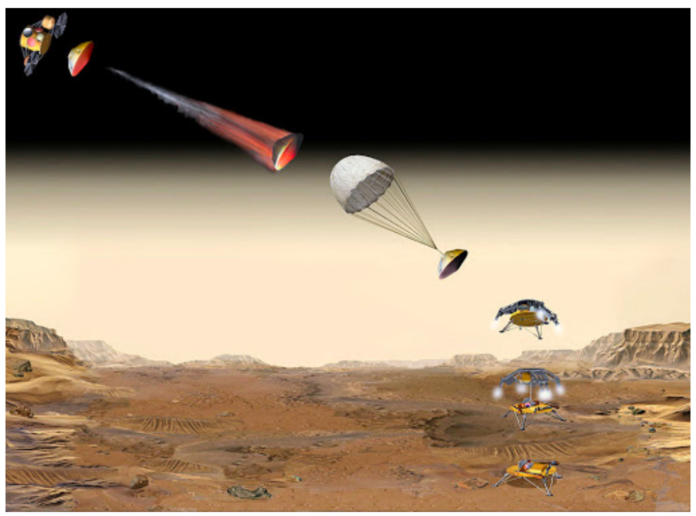

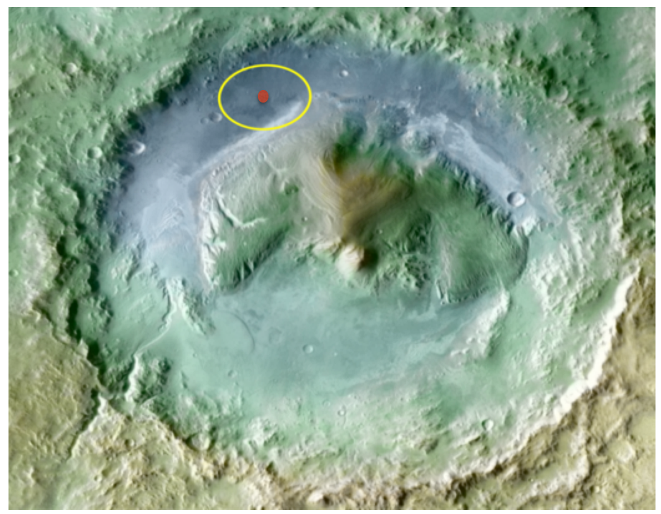

Gale Crater (4.49S, 137.42E)

Risk management for automated driving in multi-lane highways

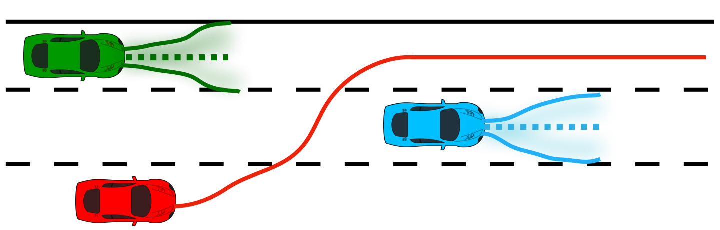

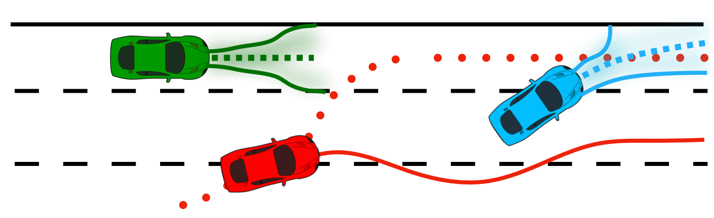

Control of uncertainties

Distribution ~ Population

Dynamic shaping of swarms

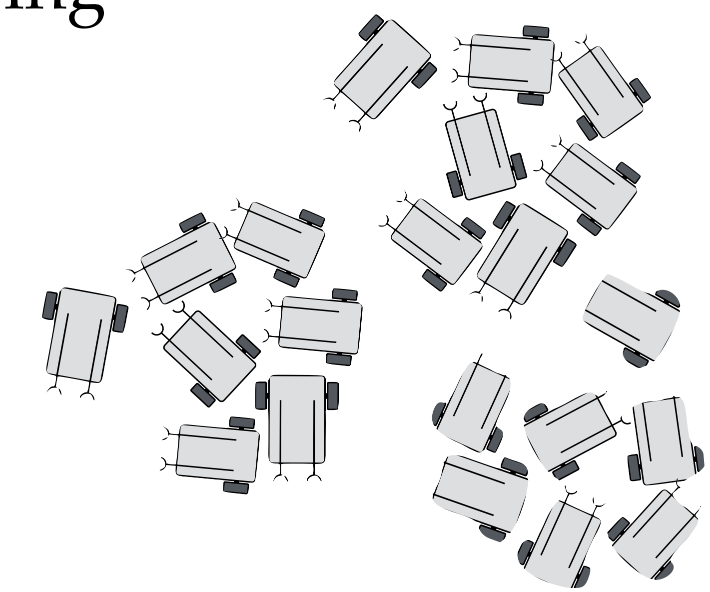

Feedback sync. and desync. of neuronal population

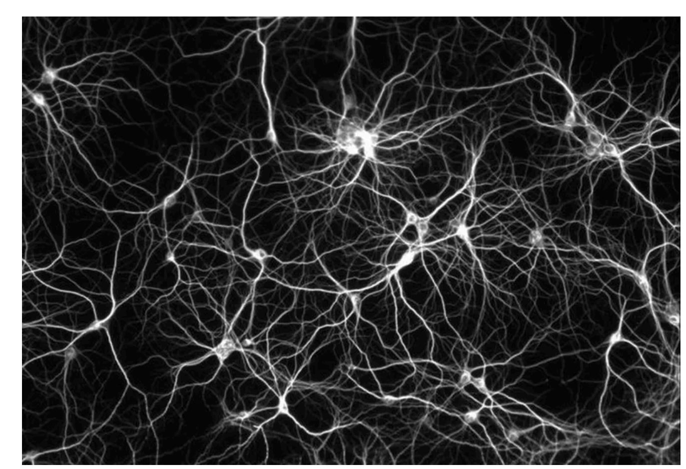

Control of ensemble

Growing Interest in Systems-Control

- **Early literature in covariance control:** Skelton et. al. [late '80s and early '90s]

- **Optimal control of Liouville PDE:** Brockett et. al. [2007-12]

- **Finite horizon linear quadratic Gaussian steering:** Chen-Georgiou-Pavon [2014-18]

- **Input-constrained covariance steering:** Bakolas [2018], Okamoto-Tsiotras [2019]

- **Linear quadratic Gaussian steering with terminal cost:** Wendel-Halder [2016], Balci-Bakolas [2020], Balci-Halder-Bakolas [2021]

- **State constraints: soft probabilistic:** Tsiotras et. al. [2018-20], **hard deterministic**: Caluya-Halder [2021]

-

Growing Interest in Systems-Control (contd.)

- **Density steering with nonlinear drift:** Chen-Georgiou-Pavon [2015], Caluya-Halder [2020-22], Nodozi-Halder [2022]

- **Application to automated driving:** Haddad-Caluya-Halder-Singh [2021]

- **Application to self-assembly:** Nodozi-Halder-Mesbah [2022]

State Feedback Density Steering

Steer joint state PDF via feedback control over finite time horizon

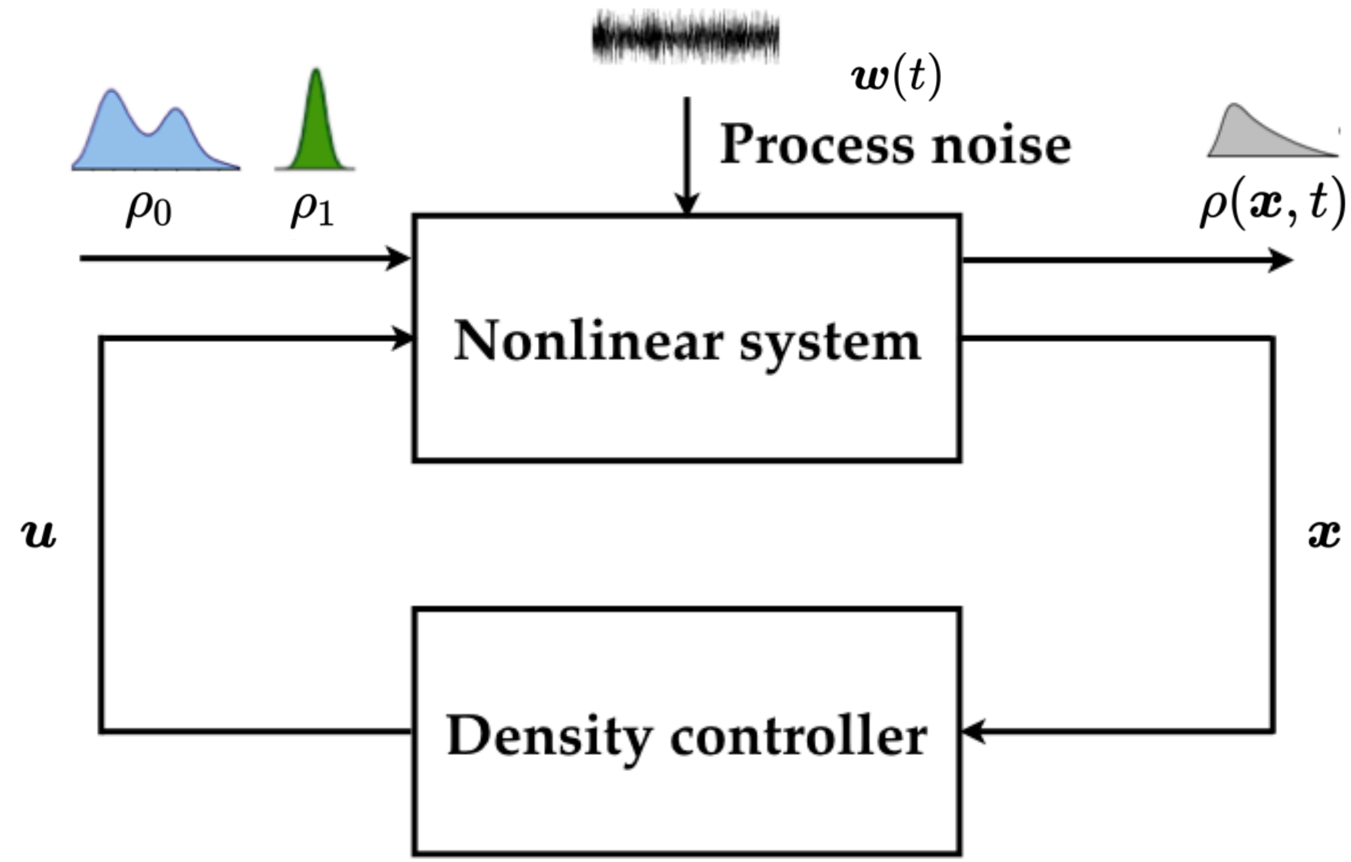

Common scenario: $\boldsymbol{G} \equiv \boldsymbol{B}$

$$\underset{u \in \mathcal{U}}{\text{minimize}} \quad \mathbb{E}\left[\int_0^1 \left(\frac{1}{2}\|\boldsymbol{u}(t, \boldsymbol{x}_t^u)\|_2^2 + q(t, \boldsymbol{x}_t^u)\right) \mathrm{d}t\right]$$

subject to

$$\mathrm{d}\boldsymbol{x}_t^u = \{\boldsymbol{f}(t, \boldsymbol{x}_t^u) + \boldsymbol{B}(t, \boldsymbol{x}_t^u)\boldsymbol{u}\}\mathrm{d}t + \sqrt{2}\boldsymbol{G}(t, \boldsymbol{x}_t^u)\mathrm{d}\boldsymbol{w}_t$$

$$\boldsymbol{x}_0^u := \boldsymbol{x}_t^u(t=0) \sim \rho_0, \quad \boldsymbol{x}_1^u := \boldsymbol{x}_t^u(t=1) \sim \rho_1$$

Optimal Control Problem over PDFs

$\mathcal{P}_{01} :=$

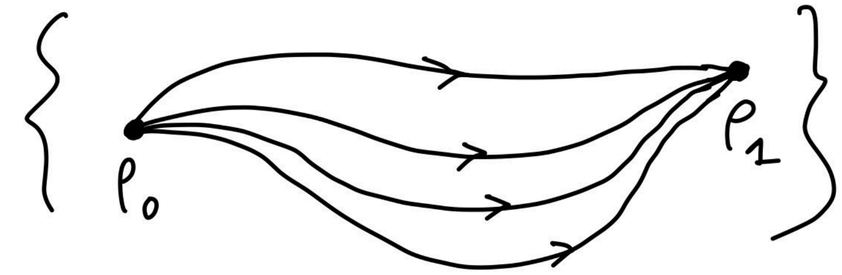

Diffusion tensor: $\boldsymbol{D} := \boldsymbol{G}\boldsymbol{G}^\top$

Hessian operator w.r.t. state: Hess

$$\inf_{(\rho,\boldsymbol{u})\in\mathcal{P}_{01}\times\mathcal{U}} \quad \int_{\mathbb{R}^n}\int_0^1 \left(\frac{1}{2}\|\boldsymbol{u}(t,\boldsymbol{x}_t^{\boldsymbol{u}})\|_2^2 + q(t,\boldsymbol{x}_t^{\boldsymbol{u}})\right)\rho(t,\boldsymbol{x}_t^{\boldsymbol{u}})\,\mathrm{d}t\,\mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}}$$

subject to

$$\frac{\partial\rho}{\partial t} + \nabla\cdot((\boldsymbol{f}+\boldsymbol{B}\boldsymbol{u})\,\rho) = \langle \mathrm{Hess}, \boldsymbol{D}\rho\rangle$$

$$\rho(t=0,\boldsymbol{x}_0^{u}) = \rho_0, \quad \rho(t=1,\boldsymbol{x}_1^{u}) = \rho_1$$

Controlled Fokker-Planck or Kolmogorov's forward PDE

Zero process noise $\rightsquigarrow$ Optimal mass transport

Dynamic optimal mass transport
with prior dynamics $\boldsymbol{f}$

$\mathcal{P}_{01} :=$ 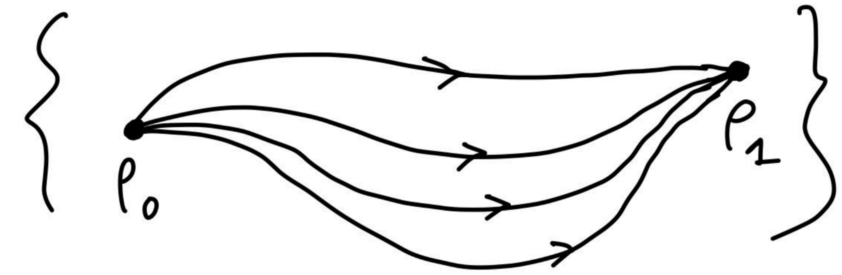

$$\inf_{(\rho,\boldsymbol{u})\in\mathcal{P}_{01}\times\mathcal{U}} \quad \int_{\mathbb{R}^n}\int_0^1 \left(\frac{1}{2}\|\boldsymbol{u}(t,\boldsymbol{x}_t^{\boldsymbol{u}})\|_2^2 + q(t,\boldsymbol{x}_t^{\boldsymbol{u}})\right)\rho(t,\boldsymbol{x}_t^{\boldsymbol{u}})\,\mathrm{d}t\,\mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}}$$

subject to

$$\frac{\partial\rho}{\partial t} + \nabla\cdot((\boldsymbol{f}+\boldsymbol{B}\boldsymbol{u})\,\rho) = \cancelto{\mathbf{0}}{\langle\mathrm{Hess},\boldsymbol{D}\rho\rangle}$$

$$\rho(t=0,\boldsymbol{x}_0^{u}) = \rho_0, \quad \rho(t=1,\boldsymbol{x}_1^{u}) = \rho_1$$

Controlled Liouville PDE

Optimal Control Problem over PDFs

Existence-uniqueness needs regularity assumptions + endpoint PDFs ρ_0, ρ_1 having finite second moments

Are known to hold for many practical classes of nonlinearities

This talk: will focus on a few important classes

Necessary Conditions of Optimality (Assuming $G \equiv B$)

Coupled nonlinear PDEs + linear boundary conditions

Controlled Fokker-Planck or Kolmogorov's forward PDE

$$\frac{\partial \rho^{\text{opt}}}{\partial t} + \nabla \cdot \left(\left(\boldsymbol{f} + \boldsymbol{D} \nabla \psi \right) \rho^{\text{opt}} \right) = \langle \text{Hess}, \boldsymbol{D} \rho \rangle$$

Hamilton-Jacobi-Bellman-like PDE

$$\frac{\partial \psi}{\partial t} + \langle \nabla \psi, \boldsymbol{f} \rangle + \langle \boldsymbol{D}, \text{Hess}(\psi) \rangle + \frac{1}{2} \langle \nabla \psi, \boldsymbol{D} \nabla \psi \rangle = q$$

Boundary conditions:

$$\rho^{\text{opt}}(\cdot, t = 0) = \rho_0, \quad \rho^{\text{opt}}(\cdot, t = 1) = \rho_1$$

Optimal control: $\boldsymbol{u}^{\text{opt}} = \boldsymbol{B}^\top \nabla \psi$

Feedback Synthesis via the Schrödinger System

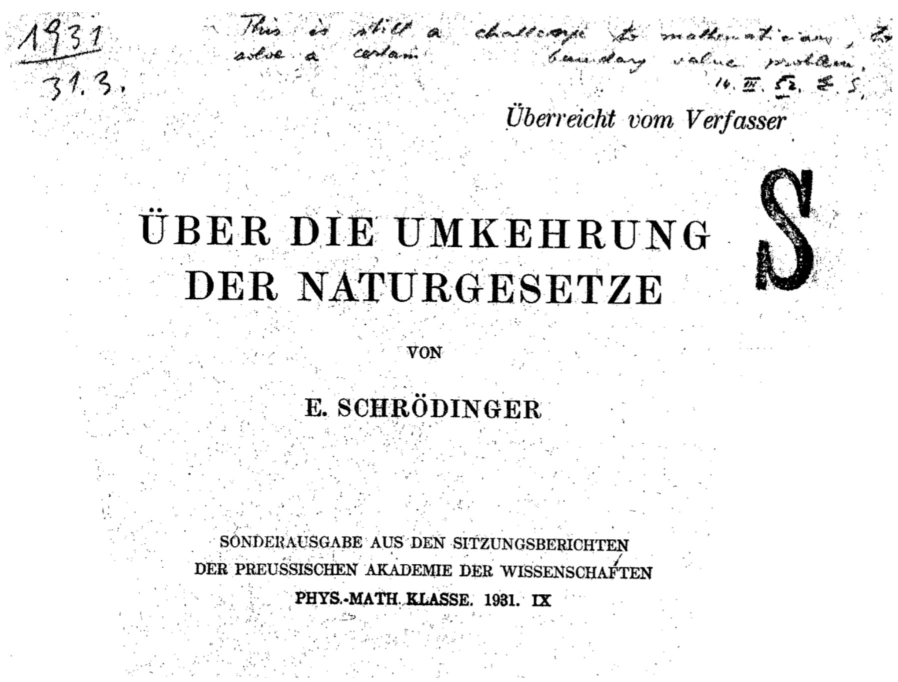

1931 31.3.

This is still a challenge to mathematicians, to solve a certain boundary value problem. 14.III.52. E.S.

Überreicht vom Verfasser

S

ÜBER DIE UMKEHRUNG DER NATURGESETZE

VON

E. SCHRÖDINGER

SONDERAUSGABE AUS DEN SITZUNGSBERICHTEN DER PREUSSISCHEN AKADEMIE DER WISSENSCHAFTEN PHYS.-MATH. KLASSE. 1931. IX

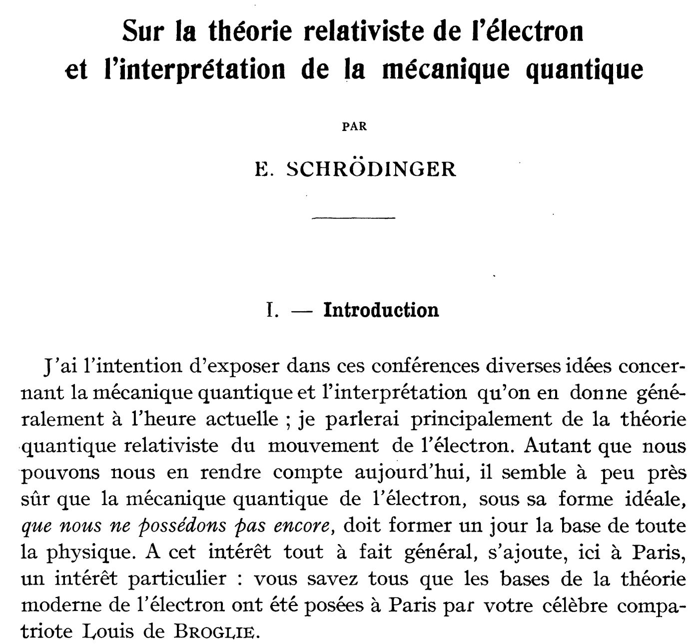

Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique

PAR

E. SCHRÖDINGER

I. — **Introduction**

J'ai l'intention d'exposer dans ces conférences diverses idées concernant la mécanique quantique et l'interprétation qu'on en donne généralement à l'heure actuelle ; je parlerai principalement de la théorie quantique relativiste du mouvement de l'électron. Autant que nous pouvons nous en rendre compte aujourd'hui, il semble à peu près sûr que la mécanique quantique de l'électron, sous sa forme idéale, *que nous ne possédons pas encore*, doit former un jour la base de toute la physique. A cet intérêt tout à fait général, s'ajoute, ici à Paris, un intérêt particulier : vous savez tous que les bases de la théorie moderne de l'électron ont été posées à Paris par votre célèbre compatriote Louis de BROGLIE.

Hopf-Cole a.k.a. Fleming's logarithmic transform:

$$\left(\rho^{\mathrm{opt}}, \psi\right) \mapsto \left(\widehat{\varphi}, \varphi\right)$$ — **Schrödinger factors**

$$\widehat{\varphi}(x,t) = \rho^{\mathrm{opt}}(x,t)\exp\left(-\psi\left(x,t\right)\right)$$

$$\varphi(x,t) = \exp\left(\psi\left(x,t\right)\right) \qquad \text{for all} \qquad (x,t) \in \mathbb{R}^n \times [0,1]$$

Feedback Synthesis via the Schrödinger System

2 coupled nonlinear PDEs $\longrightarrow$ boundary-coupled linear PDEs!!

Uncontrolled forward-backward Kolmogorov PDEs:

$$\frac{\partial \widehat{\varphi}}{\partial t} = -\nabla \cdot (\widehat{\varphi} \boldsymbol{f}) + \langle \text{Hess}, \boldsymbol{D}\widehat{\varphi} \rangle - q\widehat{\varphi}, \qquad \widehat{\varphi}_0 \varphi_0 = \rho_0,$$

$$\frac{\partial \varphi}{\partial t} = -\langle \nabla \varphi, \boldsymbol{f} \rangle - \langle \text{Hess}(\varphi), \boldsymbol{D} \rangle + q\varphi, \qquad \widehat{\varphi}_1 \varphi_1 = \rho_1,$$

Optimal controlled joint state PDF: $\rho^{\text{opt}}(\boldsymbol{x}, t) = \widehat{\varphi}(\boldsymbol{x}, t)\varphi(\boldsymbol{x}, t)$

Optimal control: $\boldsymbol{u}^{\text{opt}}(\boldsymbol{x}, t) = 2\boldsymbol{B}^\top \nabla_{\boldsymbol{x}} \log \varphi(\boldsymbol{x}, t)$

Feedback Synthesis via the Schrödinger System

$$\boxed{\text{2 coupled nonlinear PDEs} \longrightarrow \text{boundary-coupled linear PDEs!!}}$$

Uncontrolled forward-backward Kolmogorov PDEs:

$$\frac{\partial \widehat{\varphi}}{\partial t} = \boxed{-\nabla \cdot (\widehat{\varphi}\boldsymbol{f}) + \langle \mathrm{Hess}, \boldsymbol{D}\widehat{\varphi}\rangle - q\widehat{\varphi},} \qquad \widehat{\varphi}_0\varphi_0 = \rho_0,$$

$$\frac{\partial \varphi}{\partial t} = \boxed{-\langle \nabla\varphi, \boldsymbol{f}\rangle - \langle \mathrm{Hess}(\varphi), \boldsymbol{D}\rangle + q\varphi,} \qquad \widehat{\varphi}_1\varphi_1 = \rho_1,$$

$\mathcal{L}_{\text{forward}}\widehat{\varphi}$ $\qquad$ $\mathcal{L}_{\text{backward}}\varphi$

Optimal controlled joint state PDF: $\rho^{\text{opt}}(\boldsymbol{x}, t) = \widehat{\varphi}(\boldsymbol{x}, t)\varphi(\boldsymbol{x}, t)$

Optimal control: $\boldsymbol{u}^{\text{opt}}(\boldsymbol{x}, t) = 2\boldsymbol{B}^{\top}\nabla_{\boldsymbol{x}} \log \varphi(\boldsymbol{x}, t)$

Fixed Point Recursion Over Pair $\left(\varphi_1, \hat{\varphi}_0\right)$

$$
\begin{tikzcd}
\hat{\varphi}_0(\boldsymbol{x}) \arrow[rr, "\int \mathcal{L}_{\text{forward}}"] & & \hat{\varphi}_1(\boldsymbol{x}) \arrow[d, "\rho_1(\boldsymbol{x}) \oslash \hat{\varphi}_1(\boldsymbol{x})"] \\
\varphi_0(\boldsymbol{x}) \arrow[u, "\rho_0(\boldsymbol{x}) \oslash \varphi_0(\boldsymbol{x})"] & & \varphi_1(\boldsymbol{x}) \arrow[ll, "\int \mathcal{L}_{\text{backward}}"]
\end{tikzcd}
$$

This recursion is contractive in the Hilbert metric!!

Feedback Density Control: $\boldsymbol{f} \equiv \boldsymbol{0}, \boldsymbol{B} = \boldsymbol{G} \equiv \boldsymbol{I}, q \equiv 0$

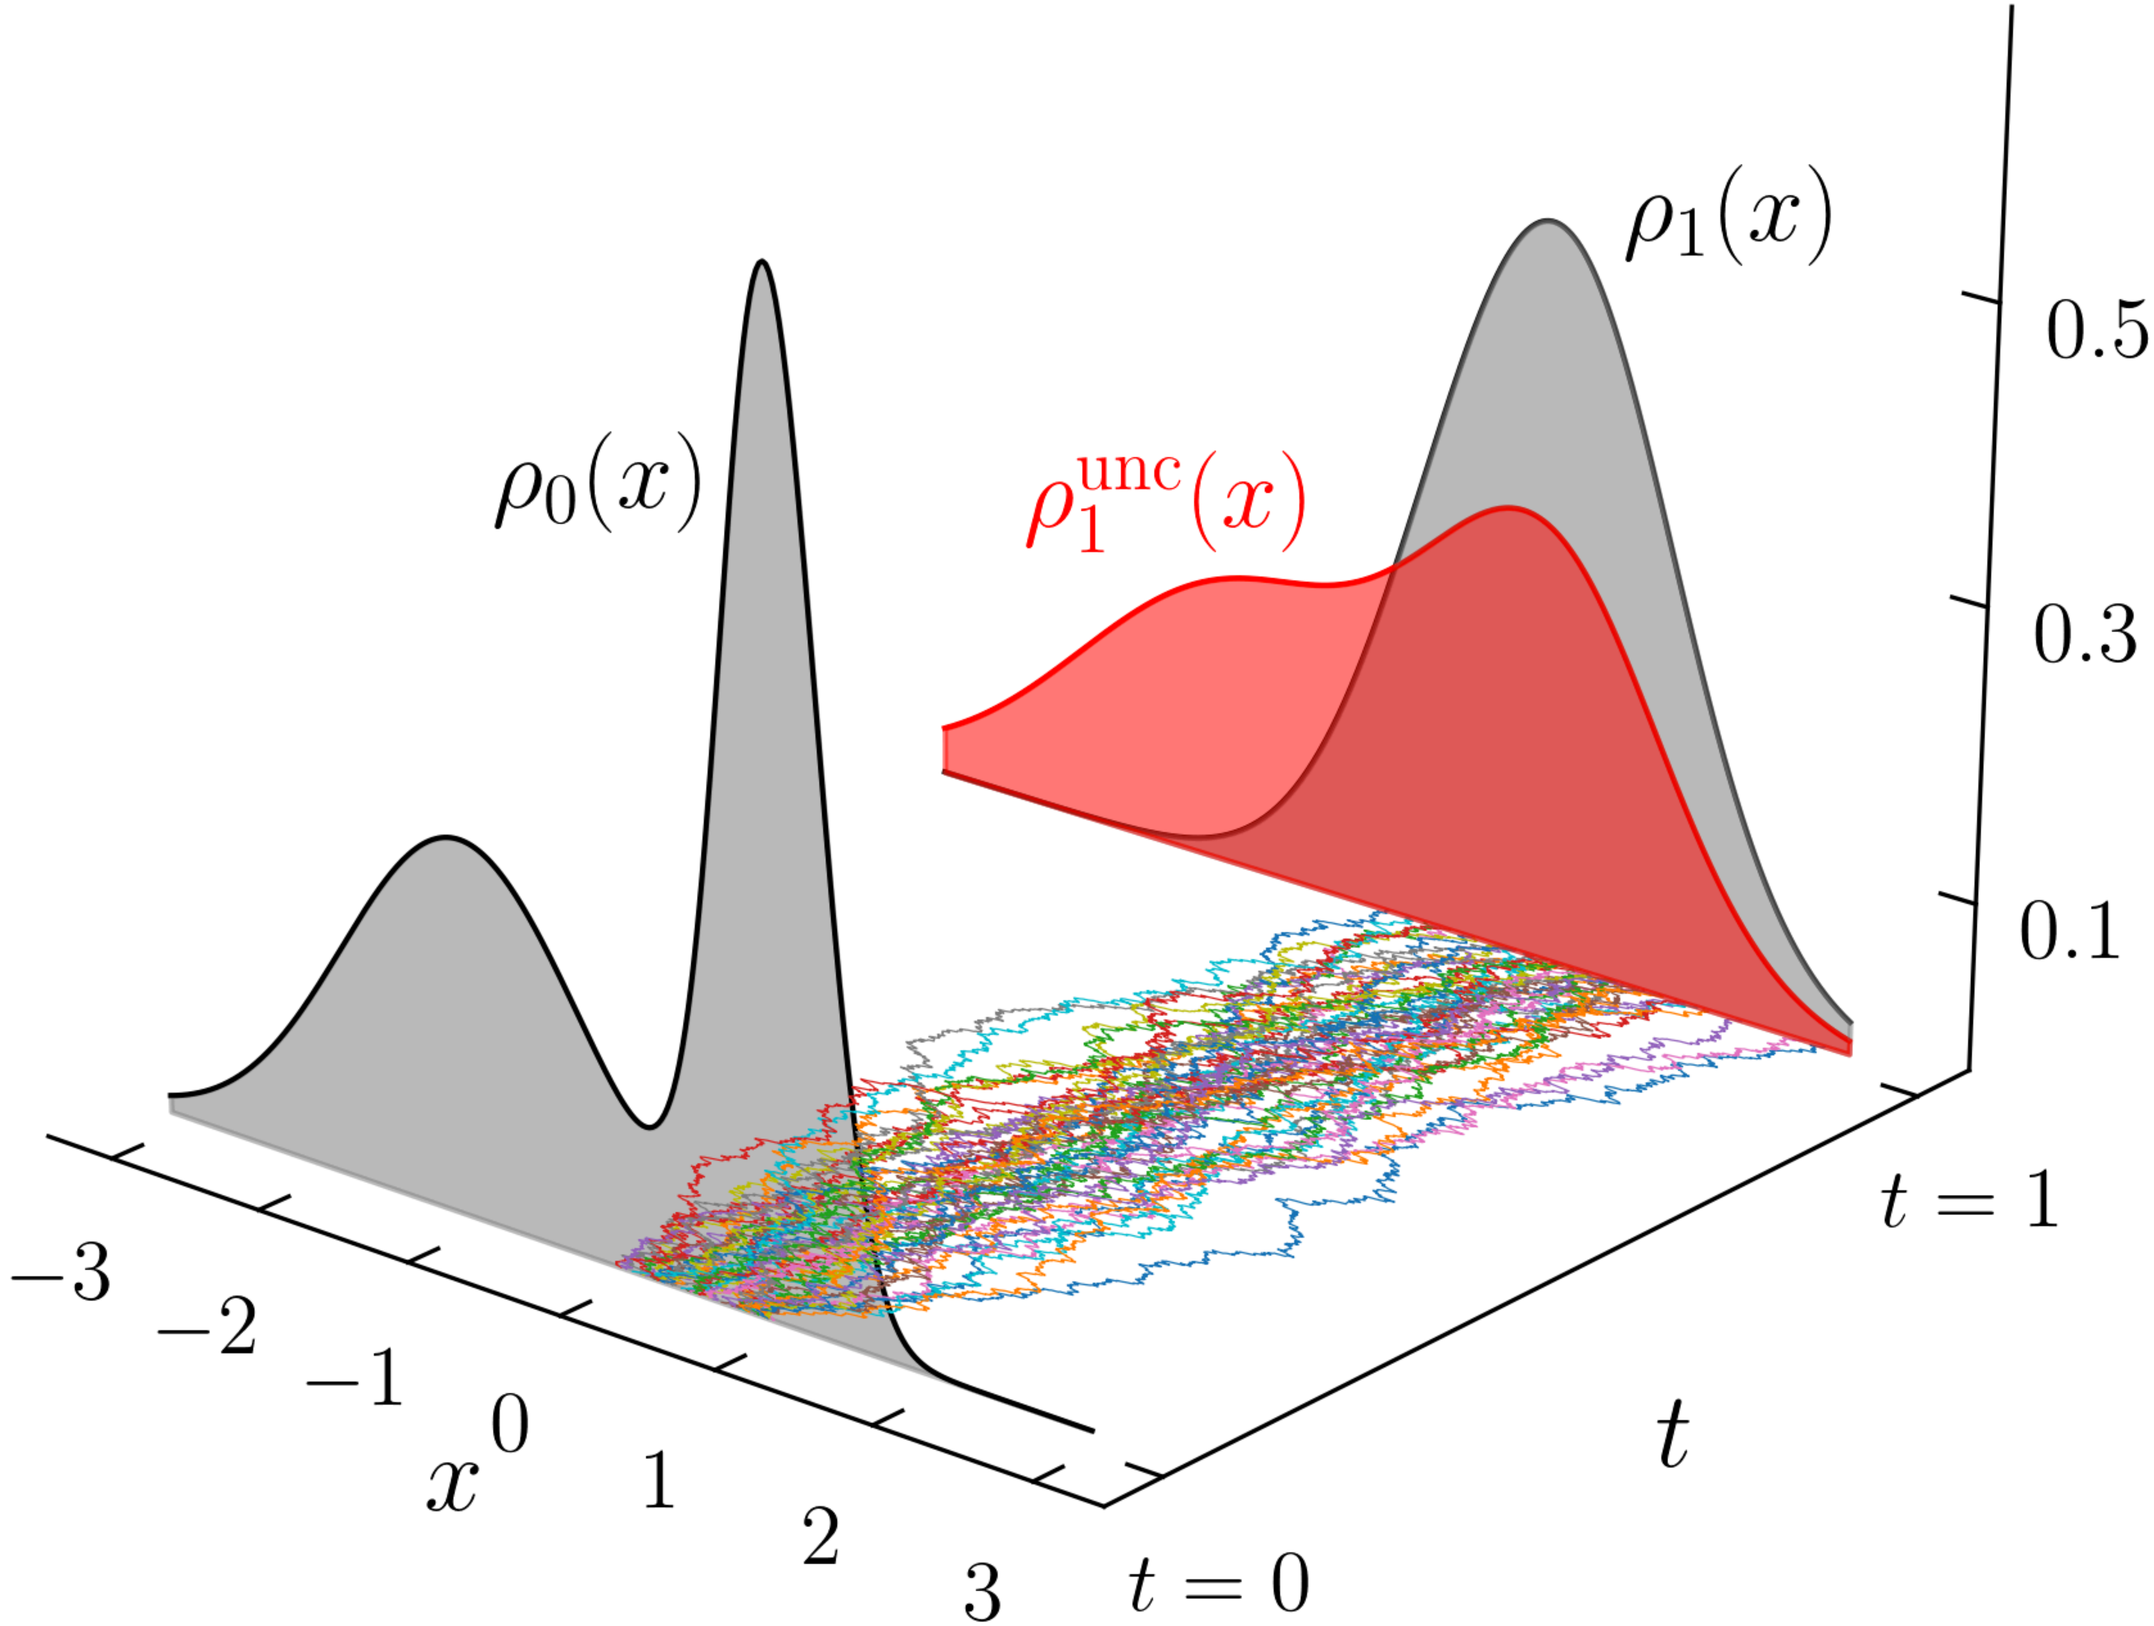

Zero prior dynamics

Feedback Density Control: $f \equiv Ax, B = G, q \equiv 0$

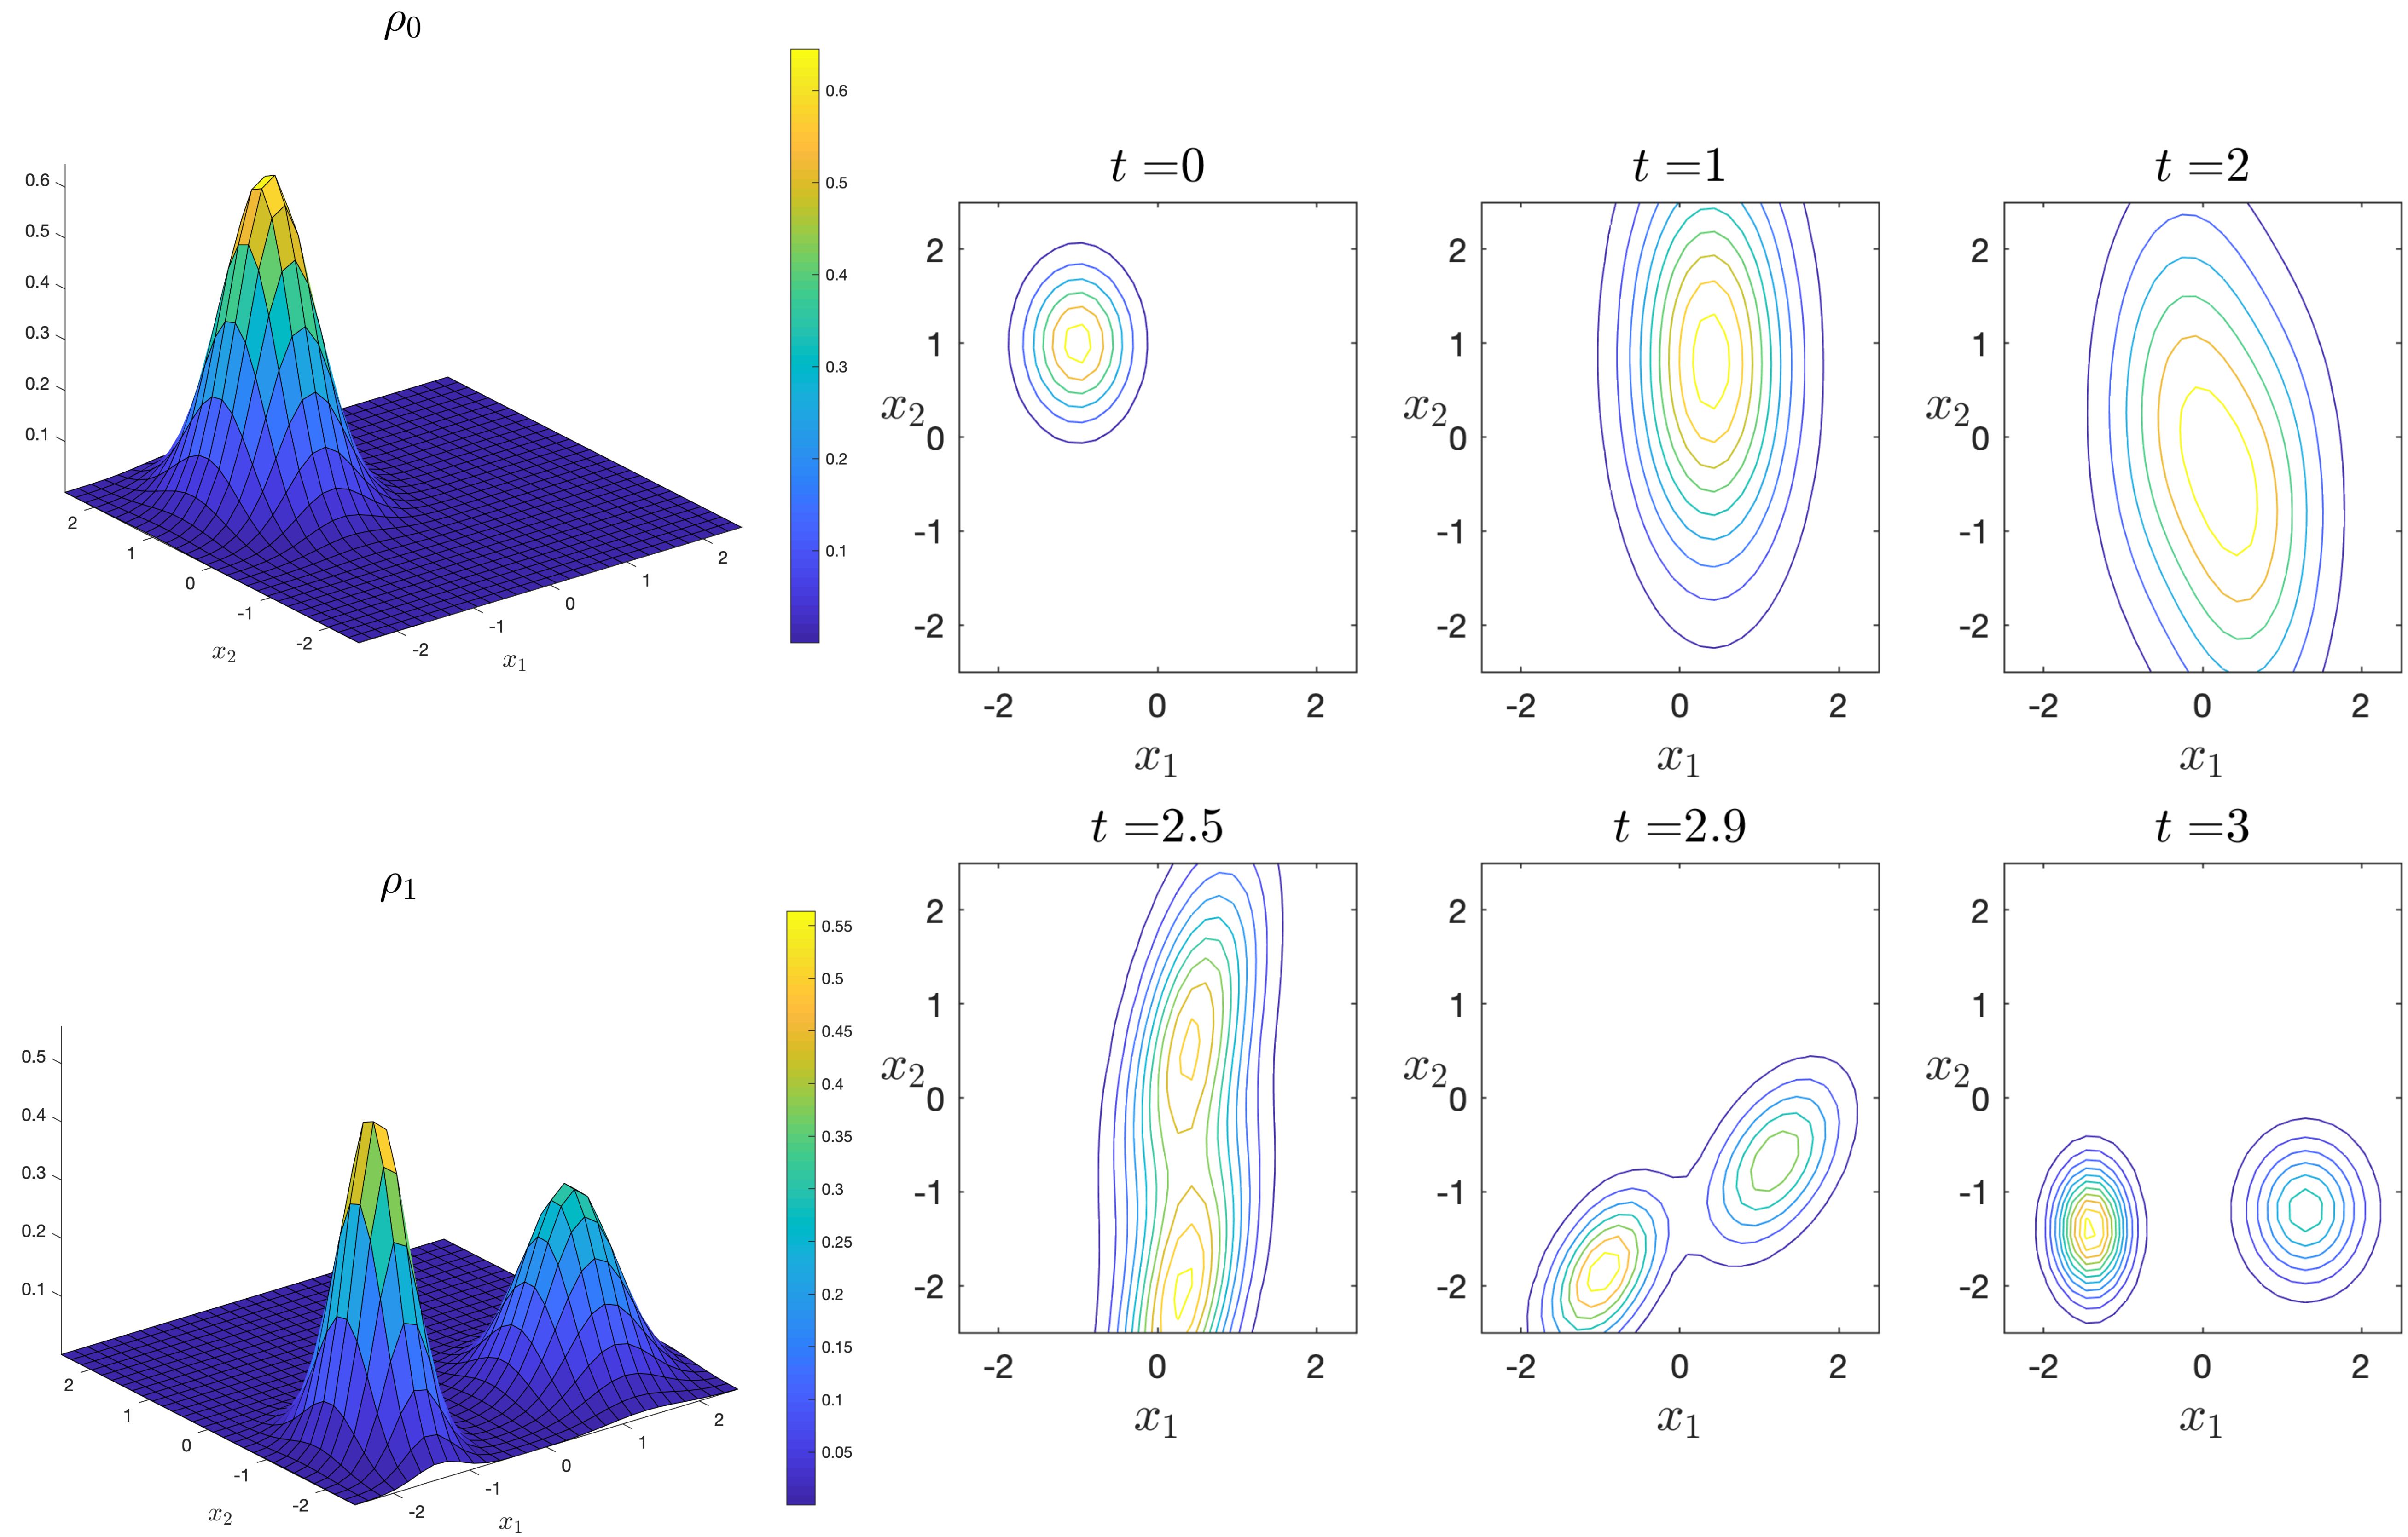

Linear prior dynamics

In general …

Need (uncontrolled) forward AND backward Kolmogorov solvers

Bad news: Need two different solvers

Good news: Sometimes one solver suffices!!!
If not, use Feynman-Kac path integral for backward

Even better: it is possible to design generalized gradient flow solvers based on point clouds!!

Brief Detour: Generalized Gradient Flow

PDE Initial Value Problem:

$$\frac{\partial \rho}{\partial t} = \mathcal{L}\rho, \quad \rho(\cdot, t=0) = \rho_0$$

$$\Updownarrow$$

Proximal Recursion:

$$\varrho_k = \operatorname{prox}^{d}_{\tau\Phi}\left(\varrho_{k-1}\right) := \underset{\varrho}{\arg\inf}\left\{\frac{1}{2}\operatorname{dist}^2\left(\varrho, \varrho_{k-1}\right) + \tau\Phi(\varrho)\right\}, \quad k \in \mathbb{N}$$

Given the operator $\mathcal{L}$, design the pair $(\operatorname{dist}, \Phi)$ such that

$$\varrho_k \xrightarrow{\tau\downarrow 0} \rho\left(\cdot, t = k\tau\right)$$

Generalized Gradient Flow

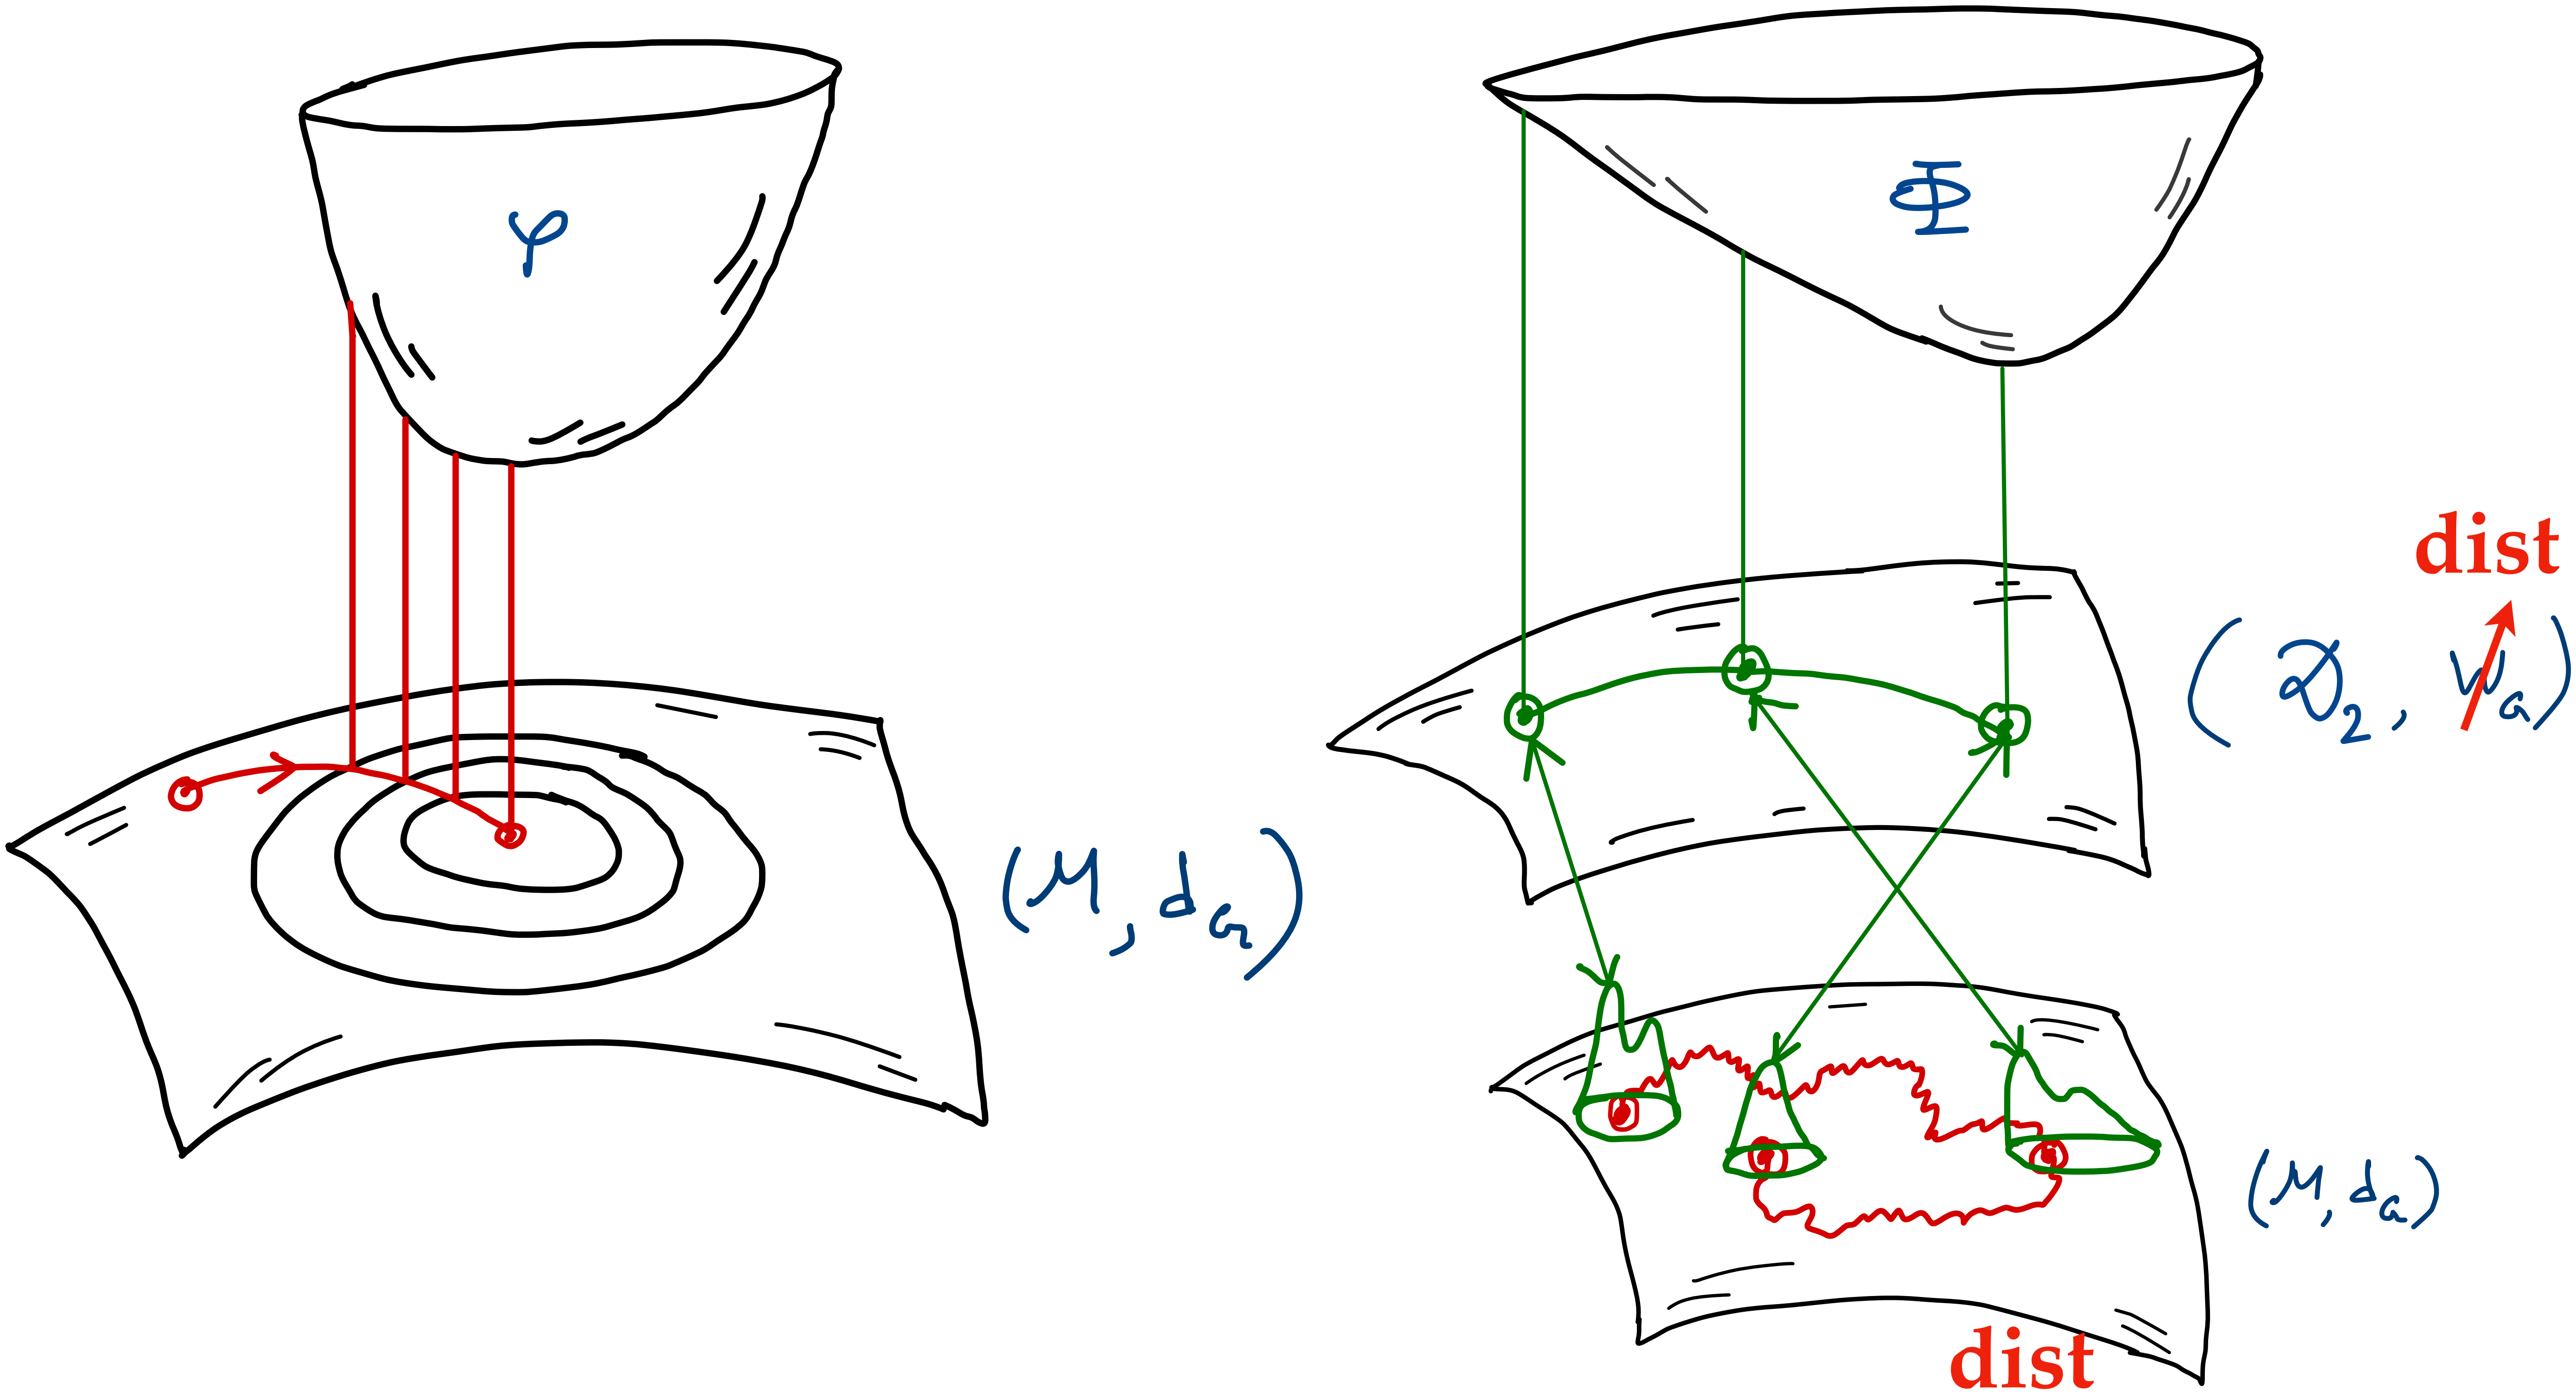

Design pair $(\overset{\text{dist}}{\cancel{W}}_G, \Phi)$

Brief Detour: Generalized Gradient Flow

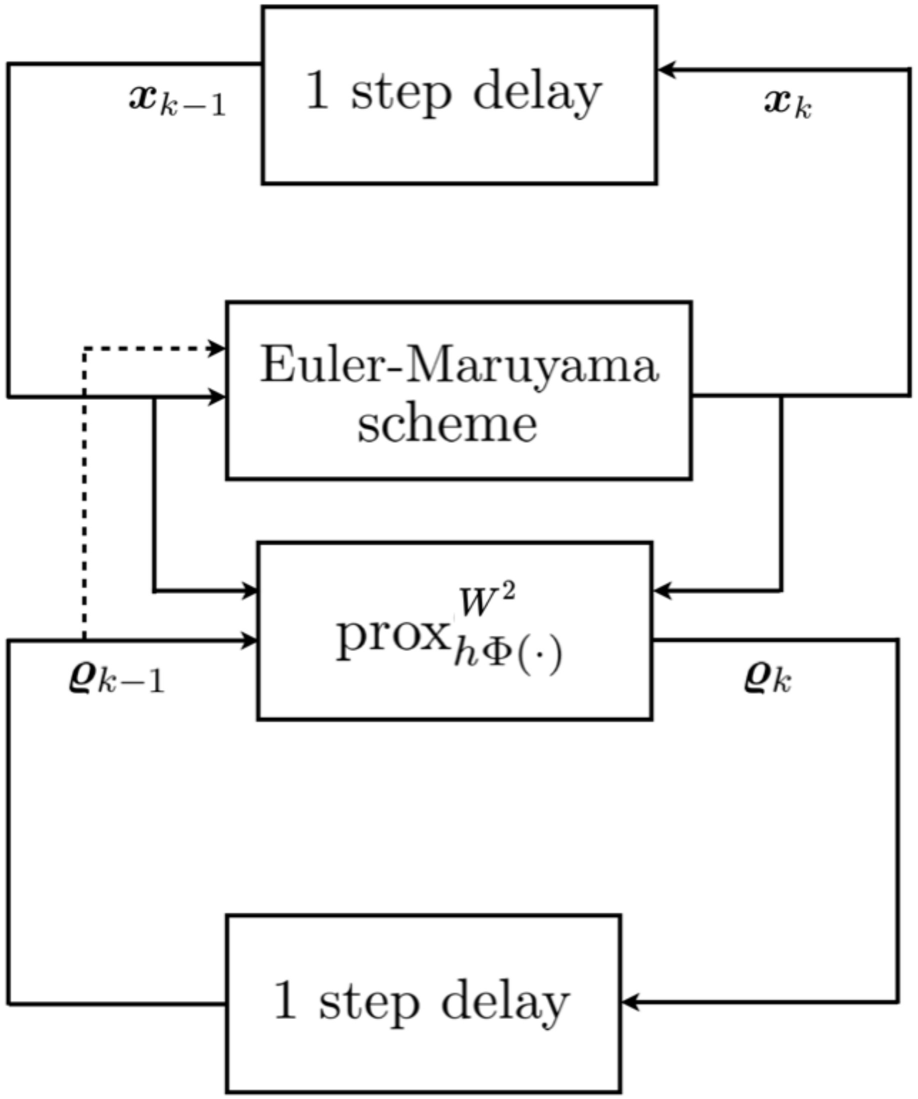

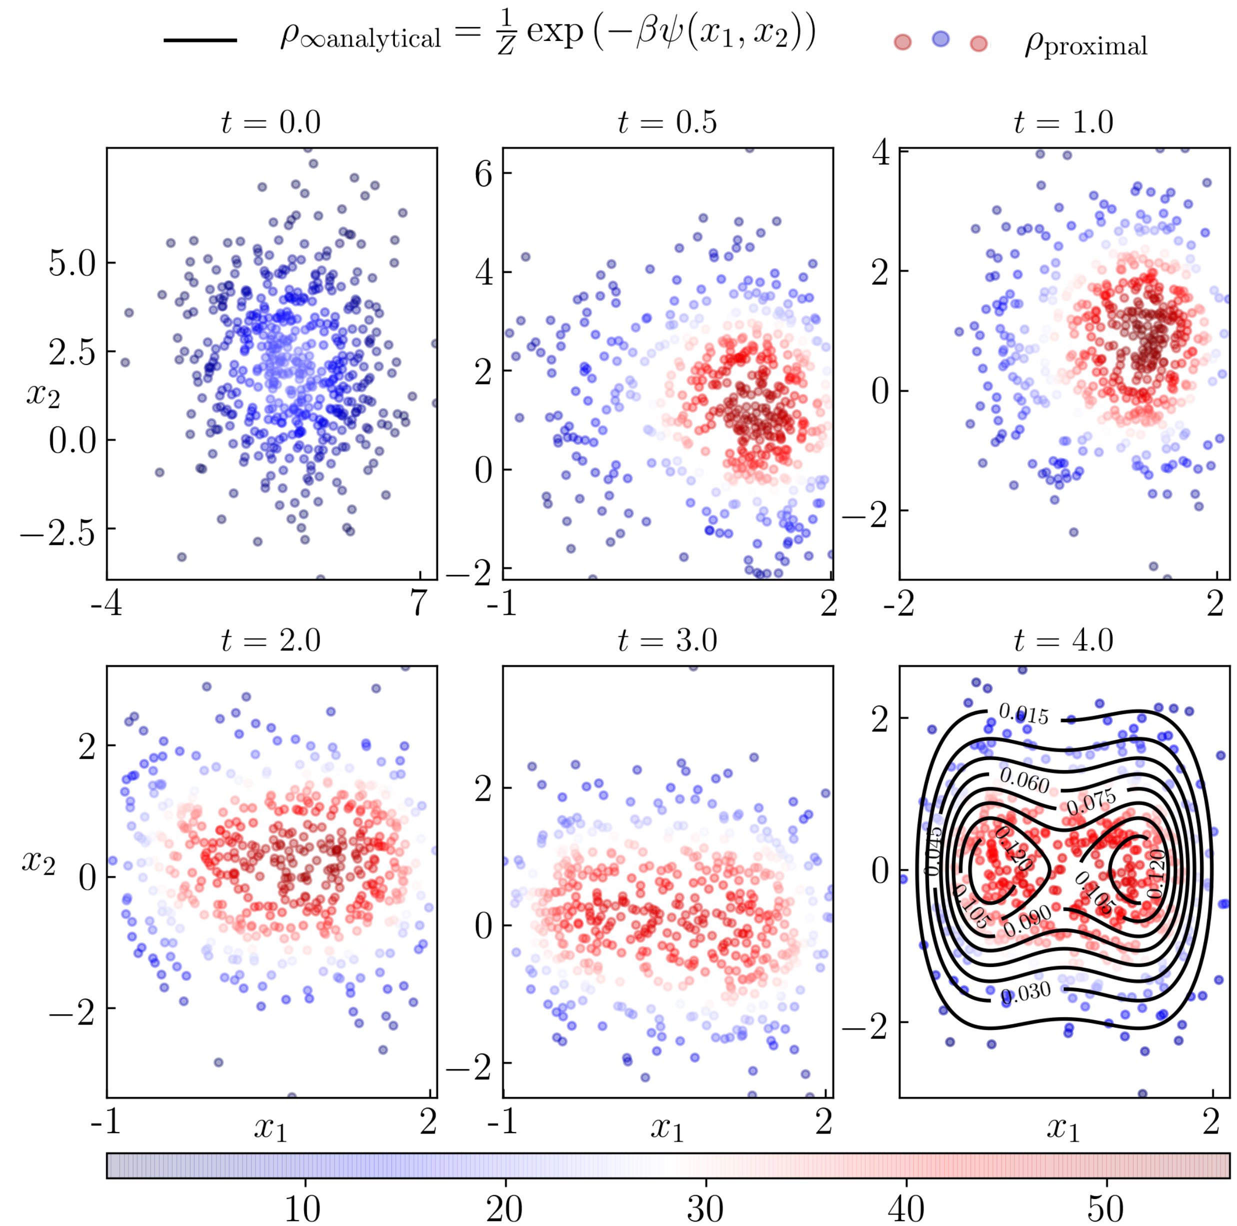

Details:

— K.F. Caluya, and A.H., Gradient flow algorithms for density propagation in stochastic systems, Vol. 65, No. 10, pp. 3991–4004, *IEEE Trans. Automatic Control, 2019*.

— K.F. Caluya, and A.H., Proximal recursion for solving the Fokker-Planck equation, *ACC 2019*.

— A.H., K.F. Caluya, B. Travacca, S.J. Moura, Hopfield neural network flow: a geometric viewpoint, Vol. 31, No. 11, pp. 4869–4880, *IEEE Trans. Neural Networks and Learning Systems*, 2020.

Single solver suffices for …

Example: gradient drift

$$\mathrm{d}\boldsymbol{x} = \{-\nabla V(\boldsymbol{x}) + \boldsymbol{u}(\boldsymbol{x},t)\}\,\mathrm{d}t + \sqrt{2\epsilon}\,\mathrm{d}\boldsymbol{w}$$

Assume: $\boldsymbol{x} \in \mathbb{R}^n$, $V \in C^2(\mathbb{R}^n)$

Example: mixed conservative-dissipative drift

$$\begin{pmatrix}\mathrm{d}\boldsymbol{\xi}\\ \mathrm{d}\boldsymbol{\eta}\end{pmatrix} = \begin{pmatrix}\boldsymbol{\eta}\\ -\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) - \kappa\boldsymbol{\eta} + \boldsymbol{u}(\boldsymbol{x},t)\end{pmatrix}\mathrm{d}t + \sqrt{2\epsilon\kappa}\begin{pmatrix}\boldsymbol{0}_{m\times m}\\ \boldsymbol{I}_{m\times m}\end{pmatrix}\mathrm{d}\boldsymbol{w}$$

Assume: $\boldsymbol{\xi},\boldsymbol{\eta} \in \mathbb{R}^m, \boldsymbol{x} := (\boldsymbol{\xi},\boldsymbol{\eta})^\top \in \mathbb{R}^n$, $n = 2m$, $V \in C^2(\mathbb{R}^m)$, $\inf V > -\infty$, $\mathrm{Hess}(V)$ unif. bounded

Feedback Density Control: Gradient Drift, $q \equiv 0$

Theorem

For $t \in [0,1]$, let $s := 1 - t$.

Define the change-of-variables $\varphi \mapsto q \mapsto p$ as

$$q(\boldsymbol{x}, s) := \varphi(\boldsymbol{x}, s) = \varphi(\boldsymbol{x}, 1-t),$$

$$p(\boldsymbol{x}, s) := q(\boldsymbol{x}, s)\exp\left(-V(\boldsymbol{x})/\epsilon\right).$$

Then the pair $(\hat{\varphi}, p)$ solves

$$\frac{\partial \hat{\varphi}}{\partial t} = \nabla \cdot (\hat{\varphi} \nabla V) + \epsilon \Delta \hat{\varphi}, \quad \hat{\varphi}\left(\boldsymbol{x}, 0\right) = \hat{\varphi}_0(\boldsymbol{x}),$$

$$\frac{\partial p}{\partial s} = \nabla \cdot (p \nabla V) + \epsilon \Delta p, \quad p\left(\boldsymbol{x}, 0\right) = \varphi_1(\boldsymbol{x})\exp\left(-V(\boldsymbol{x})/\epsilon\right).$$

Feedback Density Control: Mixed Conservative-Dissipative Drift, $q \equiv 0$

Theorem

For $t \in [0,1]$, let $s := 1 - t$. Also, let $\boldsymbol{\vartheta} := -\boldsymbol{\eta}$.

Define the change-of-variables $\varphi \mapsto q \mapsto \widetilde{p} \mapsto p$ as

$$q(\boldsymbol{\xi}, \boldsymbol{\eta}, s) := \varphi(\boldsymbol{\xi}, \boldsymbol{\eta}, s) = \varphi(\boldsymbol{\xi}, \boldsymbol{\eta}, 1-t),$$

$$\widetilde{p}(\boldsymbol{\xi}, -\boldsymbol{\eta}, s) := q(\boldsymbol{\xi}, \boldsymbol{\eta}, s) \exp\left(-\tfrac{1}{\epsilon}\left(\tfrac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right),$$

$$p\left(\boldsymbol{\xi}, \boldsymbol{\vartheta}, s\right) := \widetilde{p}(\boldsymbol{\xi}, -\boldsymbol{\eta}, s).$$

Then the pair $(\hat{\varphi}, p)$ solves

$$\frac{\partial \hat{\varphi}}{\partial t} = -\langle \boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}} \hat{\varphi} \rangle + \nabla_{\boldsymbol{\eta}} \cdot \left(\hat{\varphi}\left(\nabla_{\boldsymbol{\xi}} V\left(\boldsymbol{\xi}\right) + \kappa \boldsymbol{\eta}\right)\right) + \epsilon\kappa \Delta_{\boldsymbol{\eta}} \hat{\varphi},$$

$$\frac{\partial p}{\partial s} = -\langle \boldsymbol{\vartheta}, \nabla_{\boldsymbol{\xi}} p \rangle + \nabla_{\boldsymbol{\vartheta}} \cdot \left(p\left(\nabla_{\boldsymbol{\xi}} V\left(\boldsymbol{\xi}\right) + \kappa \boldsymbol{\vartheta}\right)\right) + \epsilon\kappa\, \Delta_{\boldsymbol{\vartheta}} p,$$

$$\hat{\varphi}\left(\boldsymbol{\xi}, \boldsymbol{\eta}, 0\right) = \hat{\varphi}_0(\boldsymbol{\xi}, \boldsymbol{\eta}),$$

$$p(\boldsymbol{\xi}, \boldsymbol{\vartheta}, 0) = \varphi_1(\boldsymbol{\xi}, -\boldsymbol{\vartheta}) \exp\left(-\tfrac{1}{\epsilon}\left(\tfrac{1}{2}\|\boldsymbol{\vartheta}\|_2^2 + V(\boldsymbol{\xi})\right)\right).$$

Feedback Density Control via Wasserstein prox.

Design proximal recursions over discrete time pair:

$(t_{k-1}, s_{k-1}) := ((k-1)\tau, (k-1)\sigma)$, $k \in \mathbb{N}$, and τ, σ are step-sizes.

The recursions are of the form:

$$\begin{pmatrix} \hat{\phi}_{t_{k-1}} \\ \omega_{s_{k-1}} \end{pmatrix} \mapsto \begin{pmatrix} \hat{\phi}_{t_k} \\ \omega_{s_k} \end{pmatrix} := \begin{pmatrix} \underset{\hat{\phi} \in \mathcal{P}_2(\mathbb{R}^n)}{\arg\inf} \ \frac{1}{2} d^2\left(\hat{\phi}_{t_{k-1}}, \hat{\phi}\right) + \tau F(\hat{\phi}) \\ \underset{\omega \in \mathcal{P}_2(\mathbb{R}^n)}{\arg\inf} \ \frac{1}{2} d^2\left(\omega_{s_{k-1}}, \omega\right) + \sigma F(\omega) \end{pmatrix}$$

Consistency guarantees:

$$\hat{\phi}_{t_{k-1}}(\boldsymbol{x}) \to \hat{\varphi}(\boldsymbol{x}, t = (k-1)\tau) \quad \text{in} \quad L^1(\mathbb{R}^n) \quad \text{as} \quad \tau \downarrow 0,$$

$$\omega_{s_{k-1}}(\boldsymbol{x}) \to p(\boldsymbol{x}, s = (k-1)\sigma) \quad \text{in} \quad L^1(\mathbb{R}^n) \quad \text{as} \quad \sigma \downarrow 0.$$

Details:

— K.F. Caluya, and A.H., Wasserstein Proximal Algorithms for the Schrödinger Bridge Problem: Density Control with Nonlinear Drift, *IEEE Trans. Automatic Control, 2022*.

Feedback Density Control: Gradient Drift

Uncontrolled joint PDF evolution:

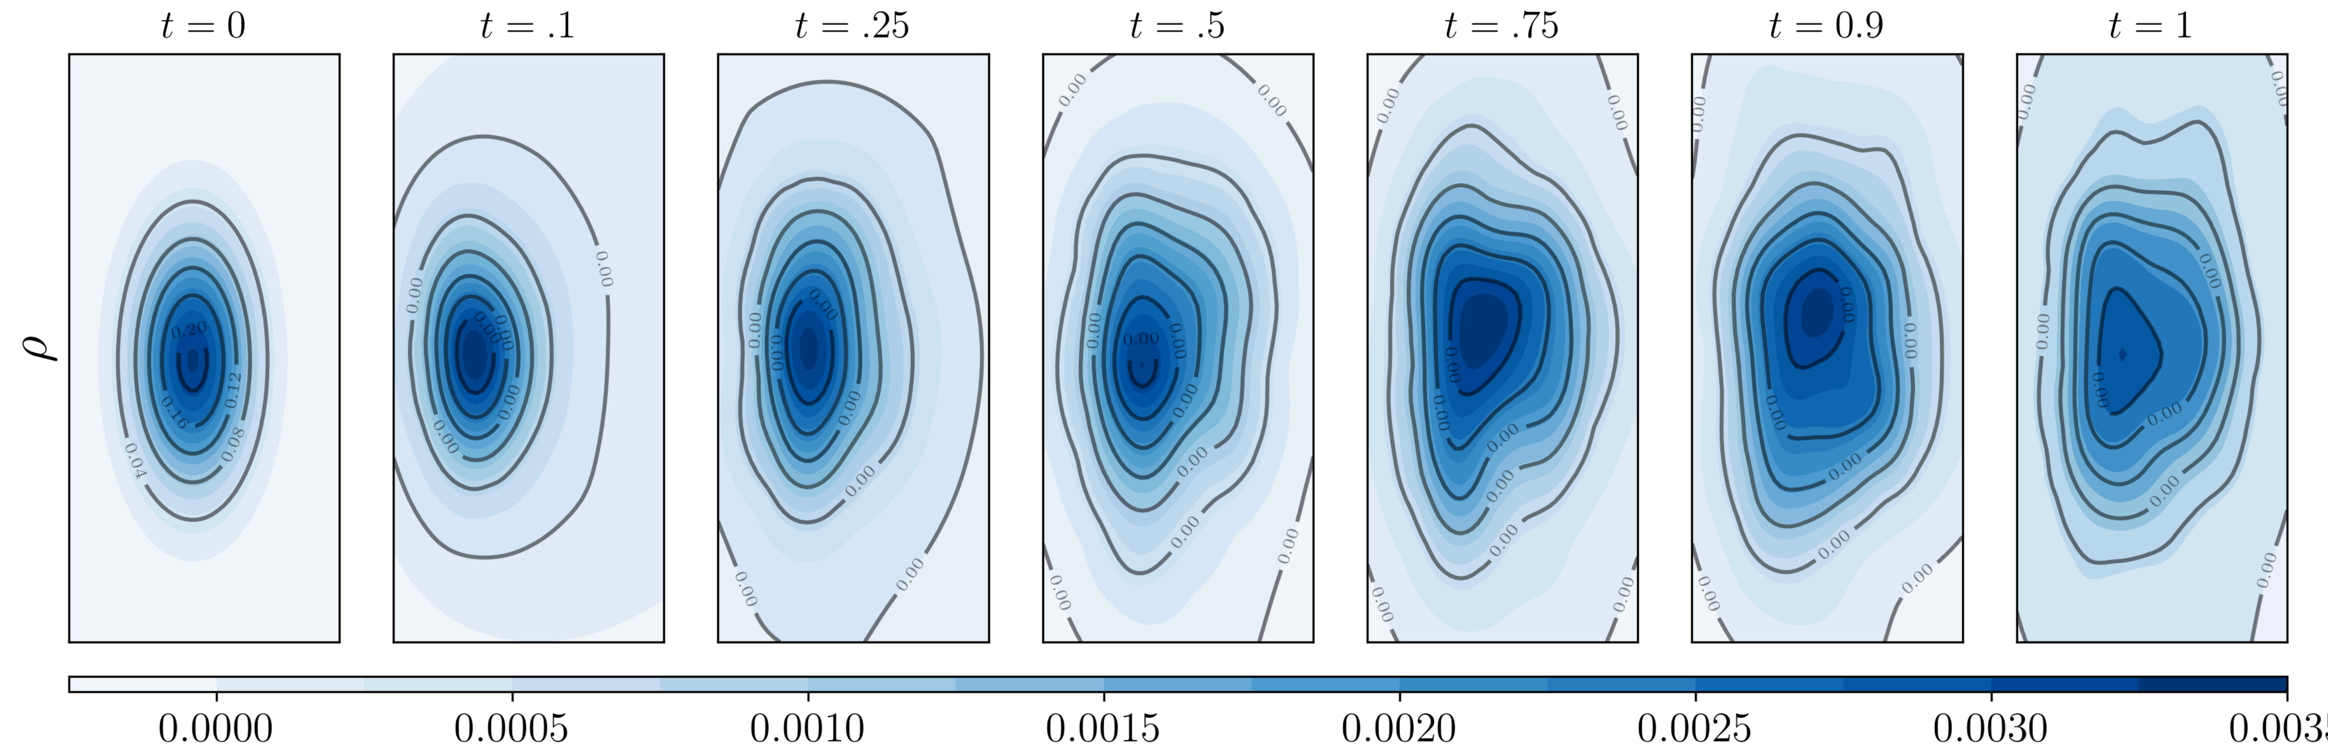

Optimal controlled joint PDF evolution:

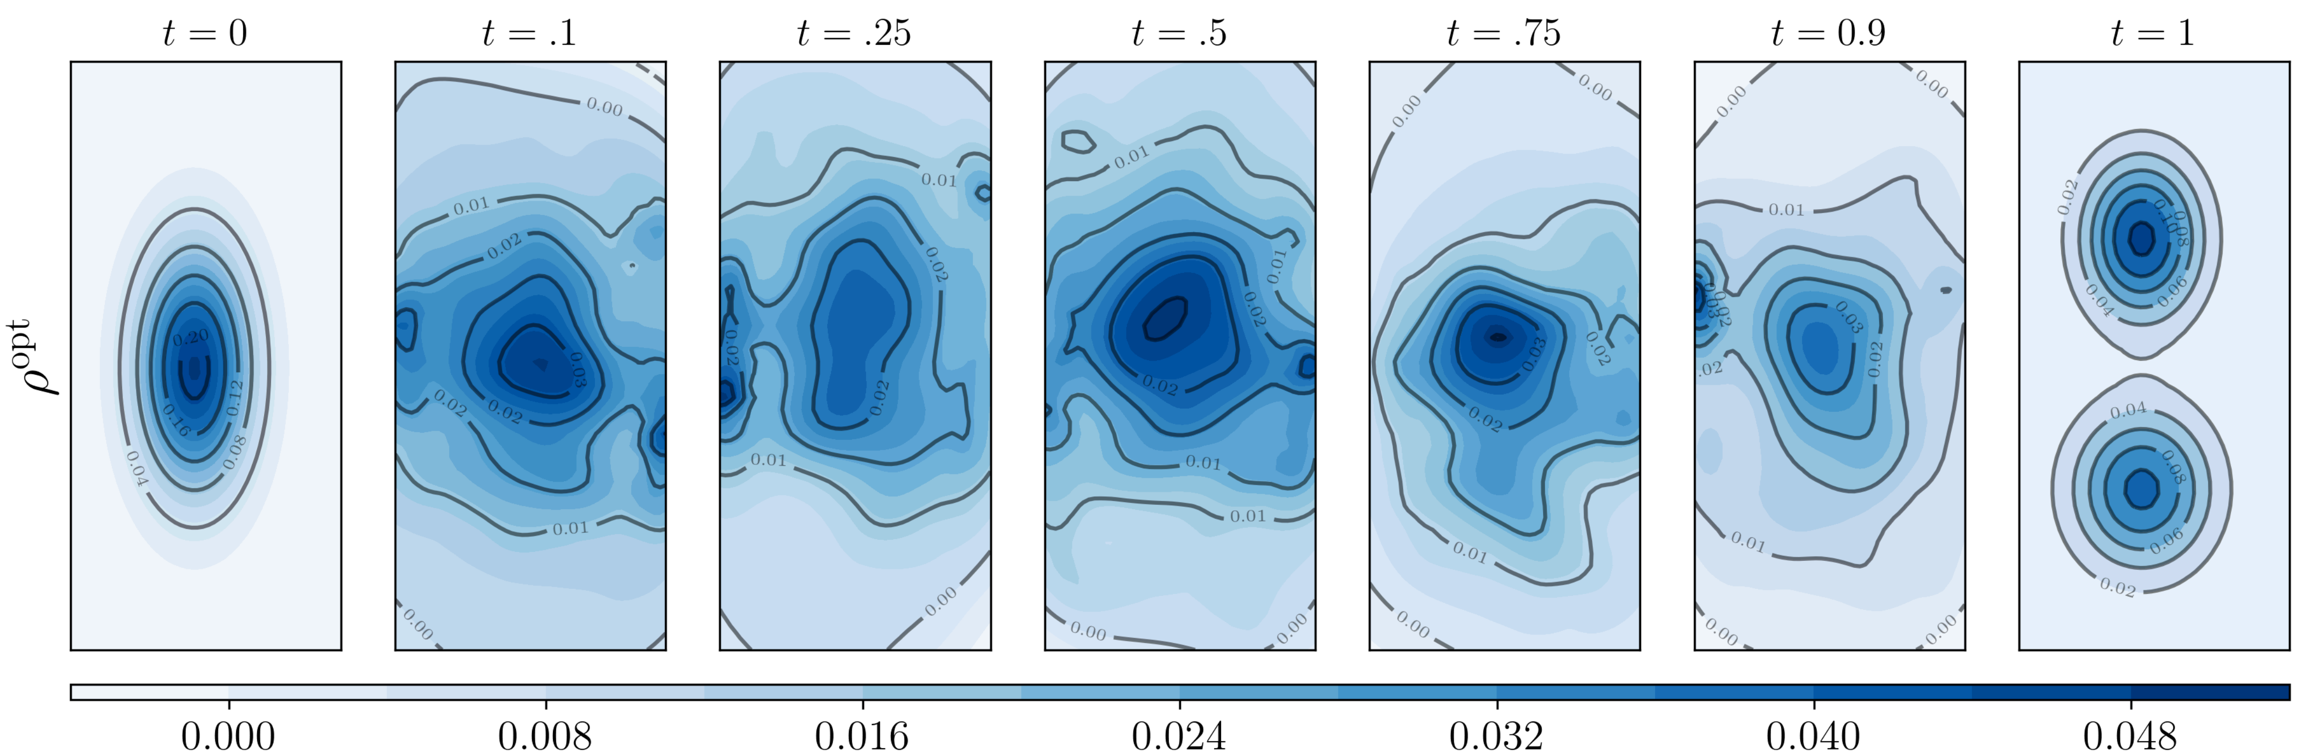

Feedback Density Control: Mixed Conservative-Dissipative Drift

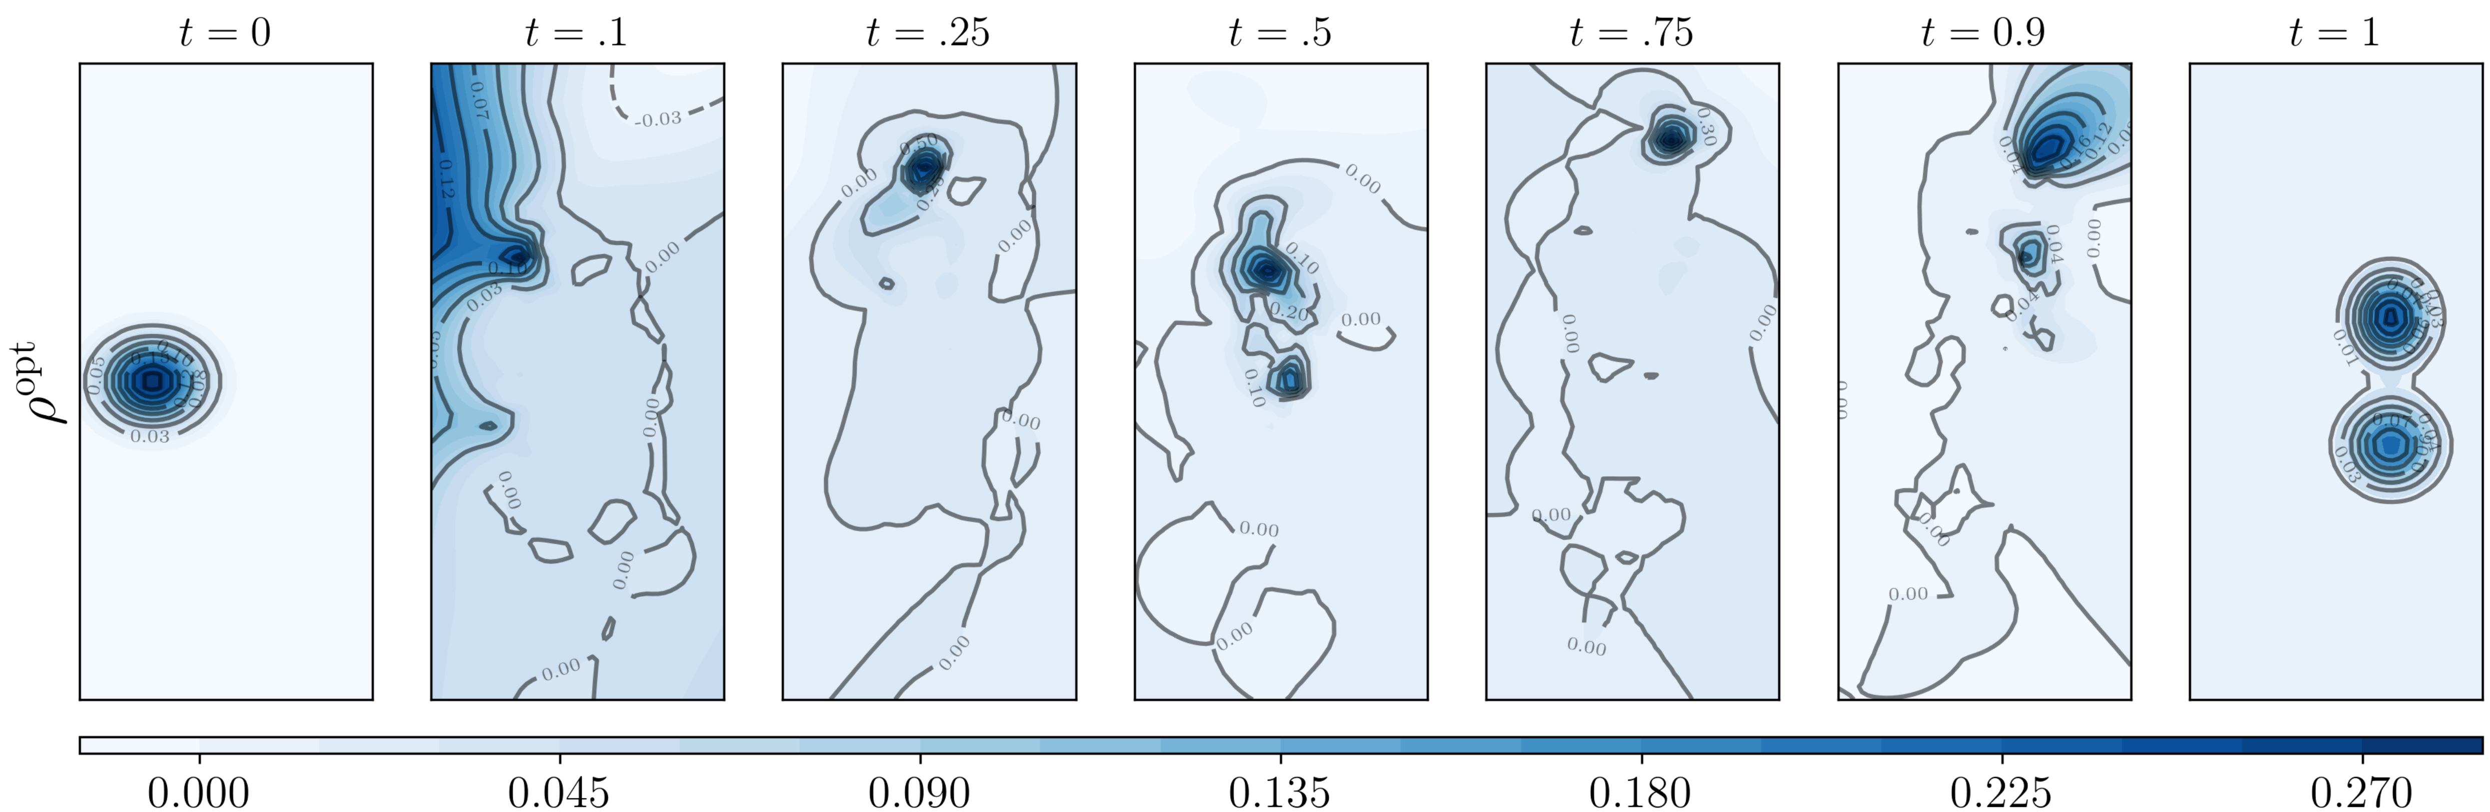

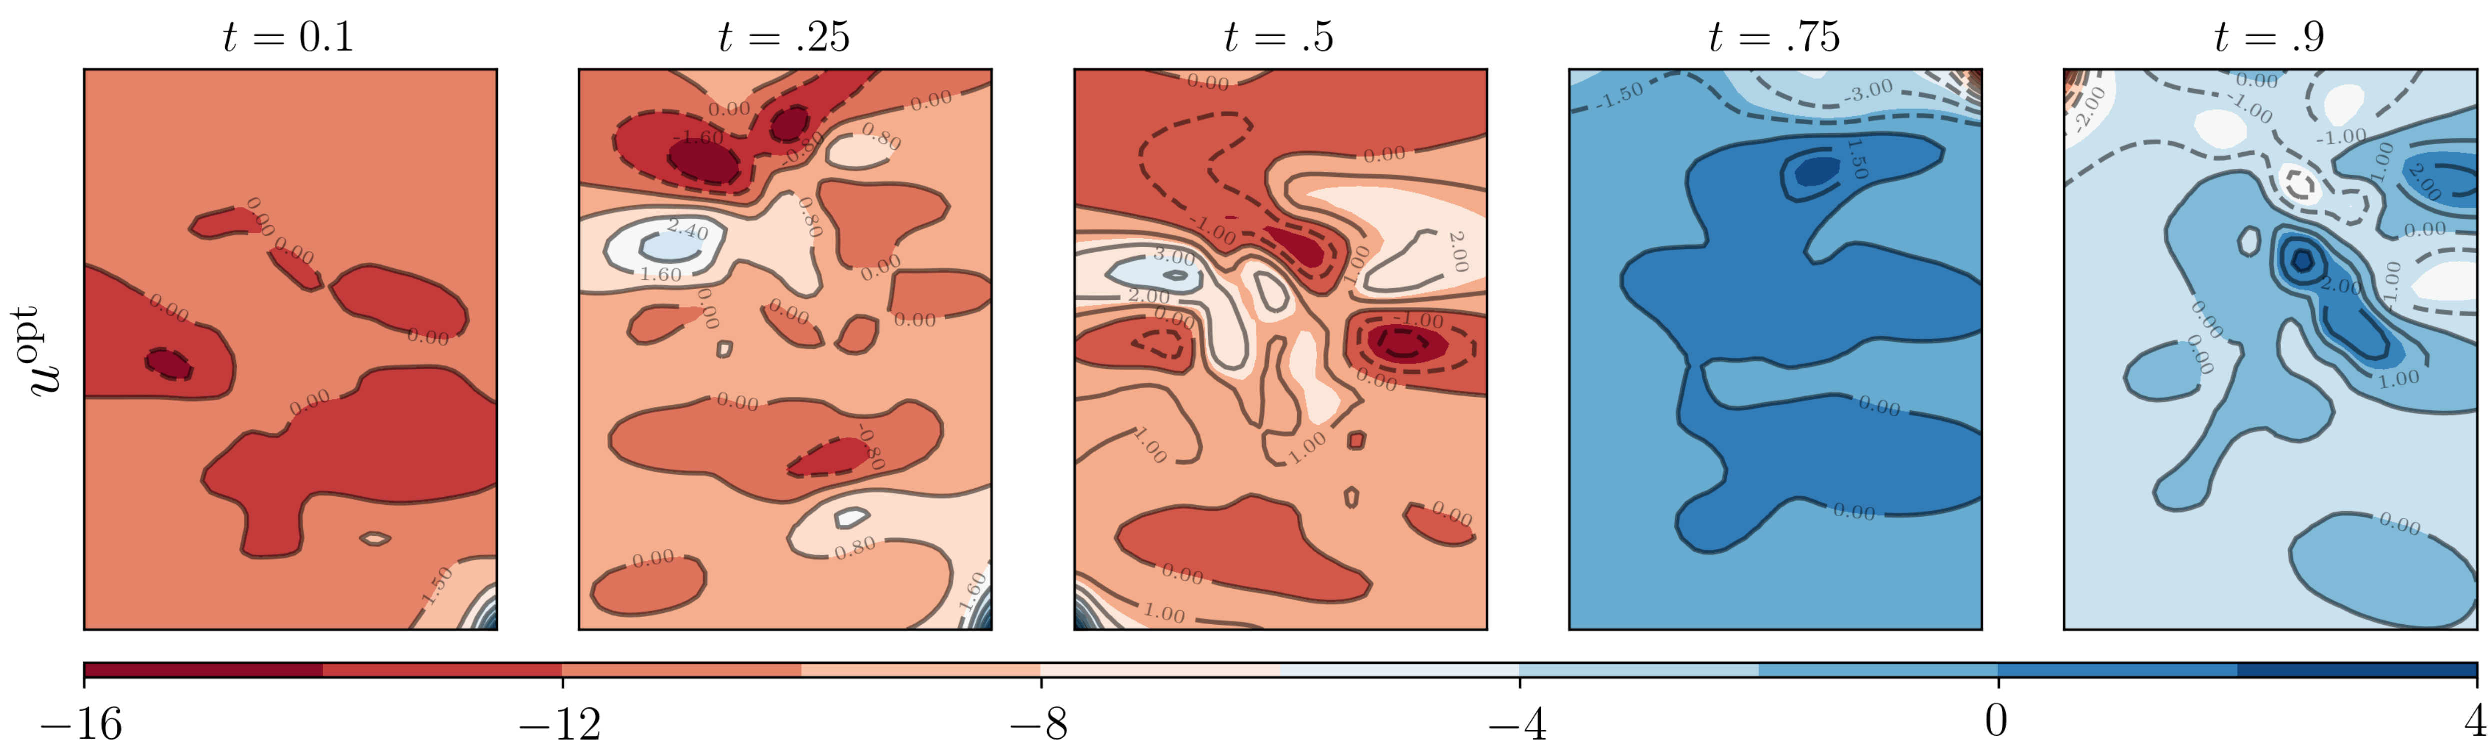

Nonlinear Density Steering with Deterministic Path Constraints

$$\int \text{with b.c. } \langle \boldsymbol{f}\hat{\varphi} - \theta\nabla\hat{\varphi}, \boldsymbol{n}\rangle\big|_{\partial\mathcal{X}} = 0$$

Robin b.c.

$$\begin{tikzcd}
\hat{\varphi}_0(\boldsymbol{x}) \arrow[rr] & & \hat{\varphi}_1(\boldsymbol{x}) \arrow[d, "\rho_1(\boldsymbol{x}) \oslash \hat{\varphi}_1(\boldsymbol{x})"] \\
\varphi_0(\boldsymbol{x}) \arrow[u, "\rho_0(\boldsymbol{x}) \oslash \varphi_0(\boldsymbol{x})"] & & \varphi_1(\boldsymbol{x}) \arrow[ll]
\end{tikzcd}$$

$$\int \text{with b.c. } \langle \nabla\varphi, \boldsymbol{n}\rangle\big|_{\partial\mathcal{X}} = 0$$

Neumann b.c.

— K.F. Caluya, and A.H., Reflected Schrödinger Bridge: Density Control with Path Constraints, *ACC 2021*.

Nonlinear Density Steering with Deterministic Path Constraints

Optimal controlled state PDFs: $V(x_1, x_2) = (x_1^2 + x_2^3)/5, \quad \overline{\mathcal{X}} = [-4, 4]^2$

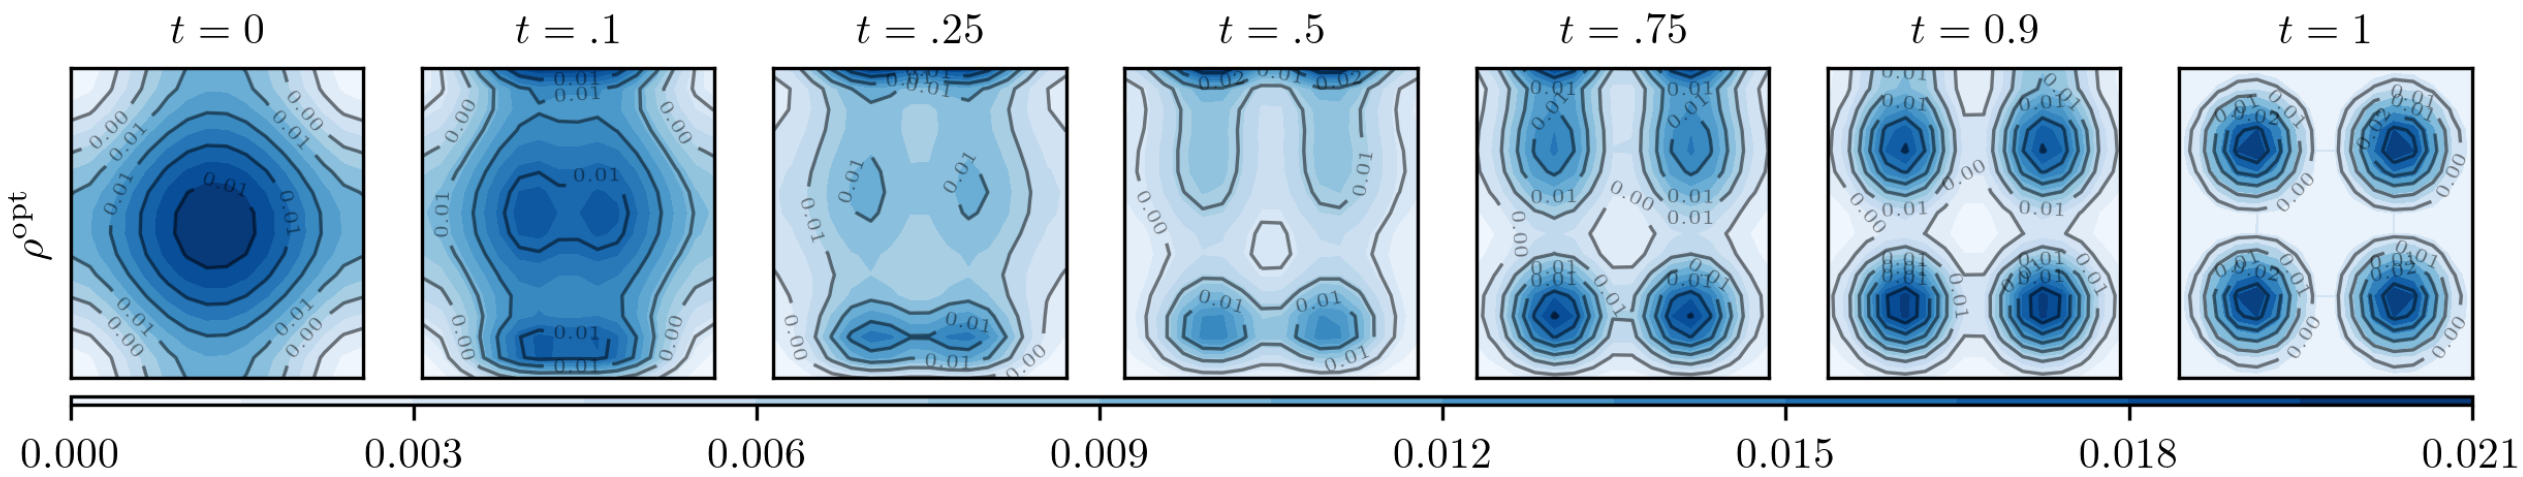

Uncontrolled state PDFs:

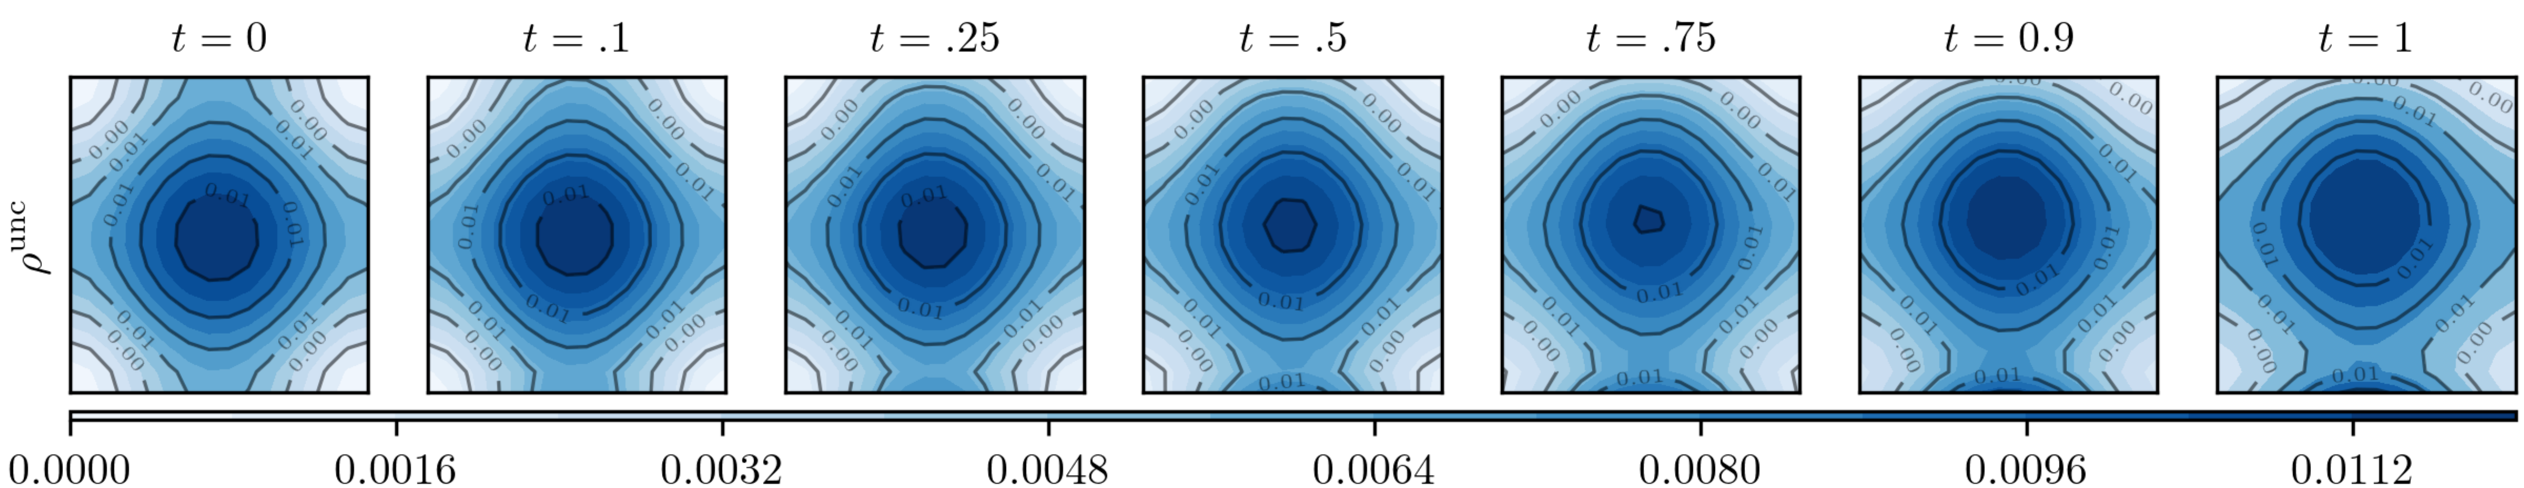

— K.F. Caluya, and A.H., Reflected Schrödinger Bridge: Density Control with Path Constraints, *ACC 2021*.

Application: Multi-lane Automated Driving

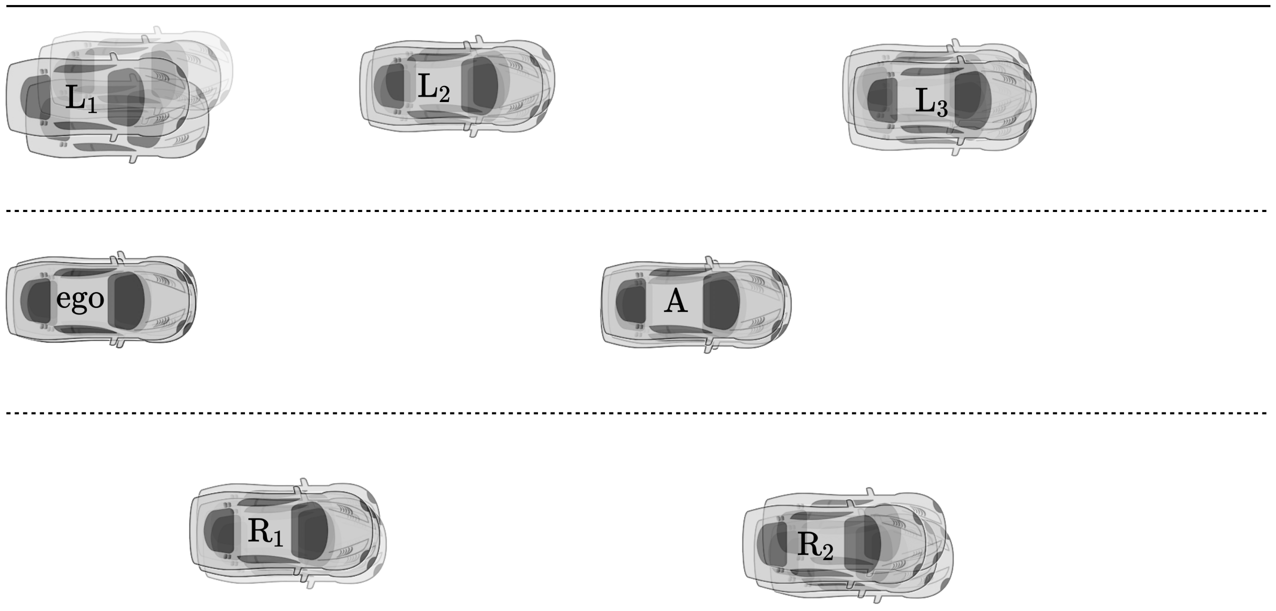

t_0

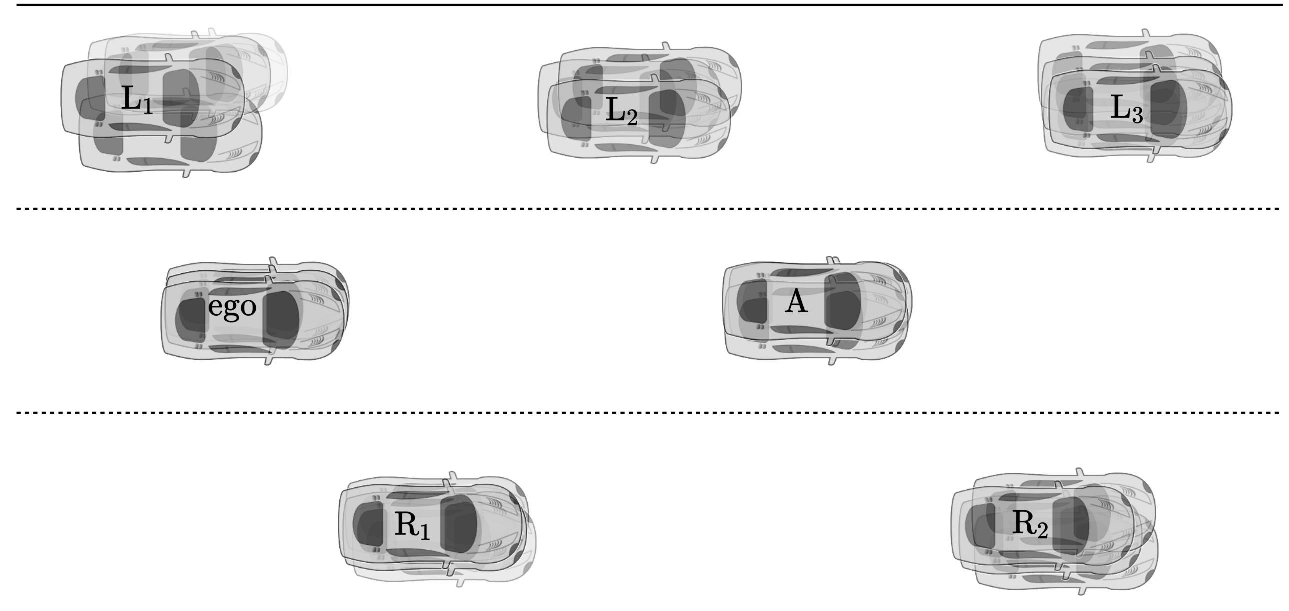

t_1

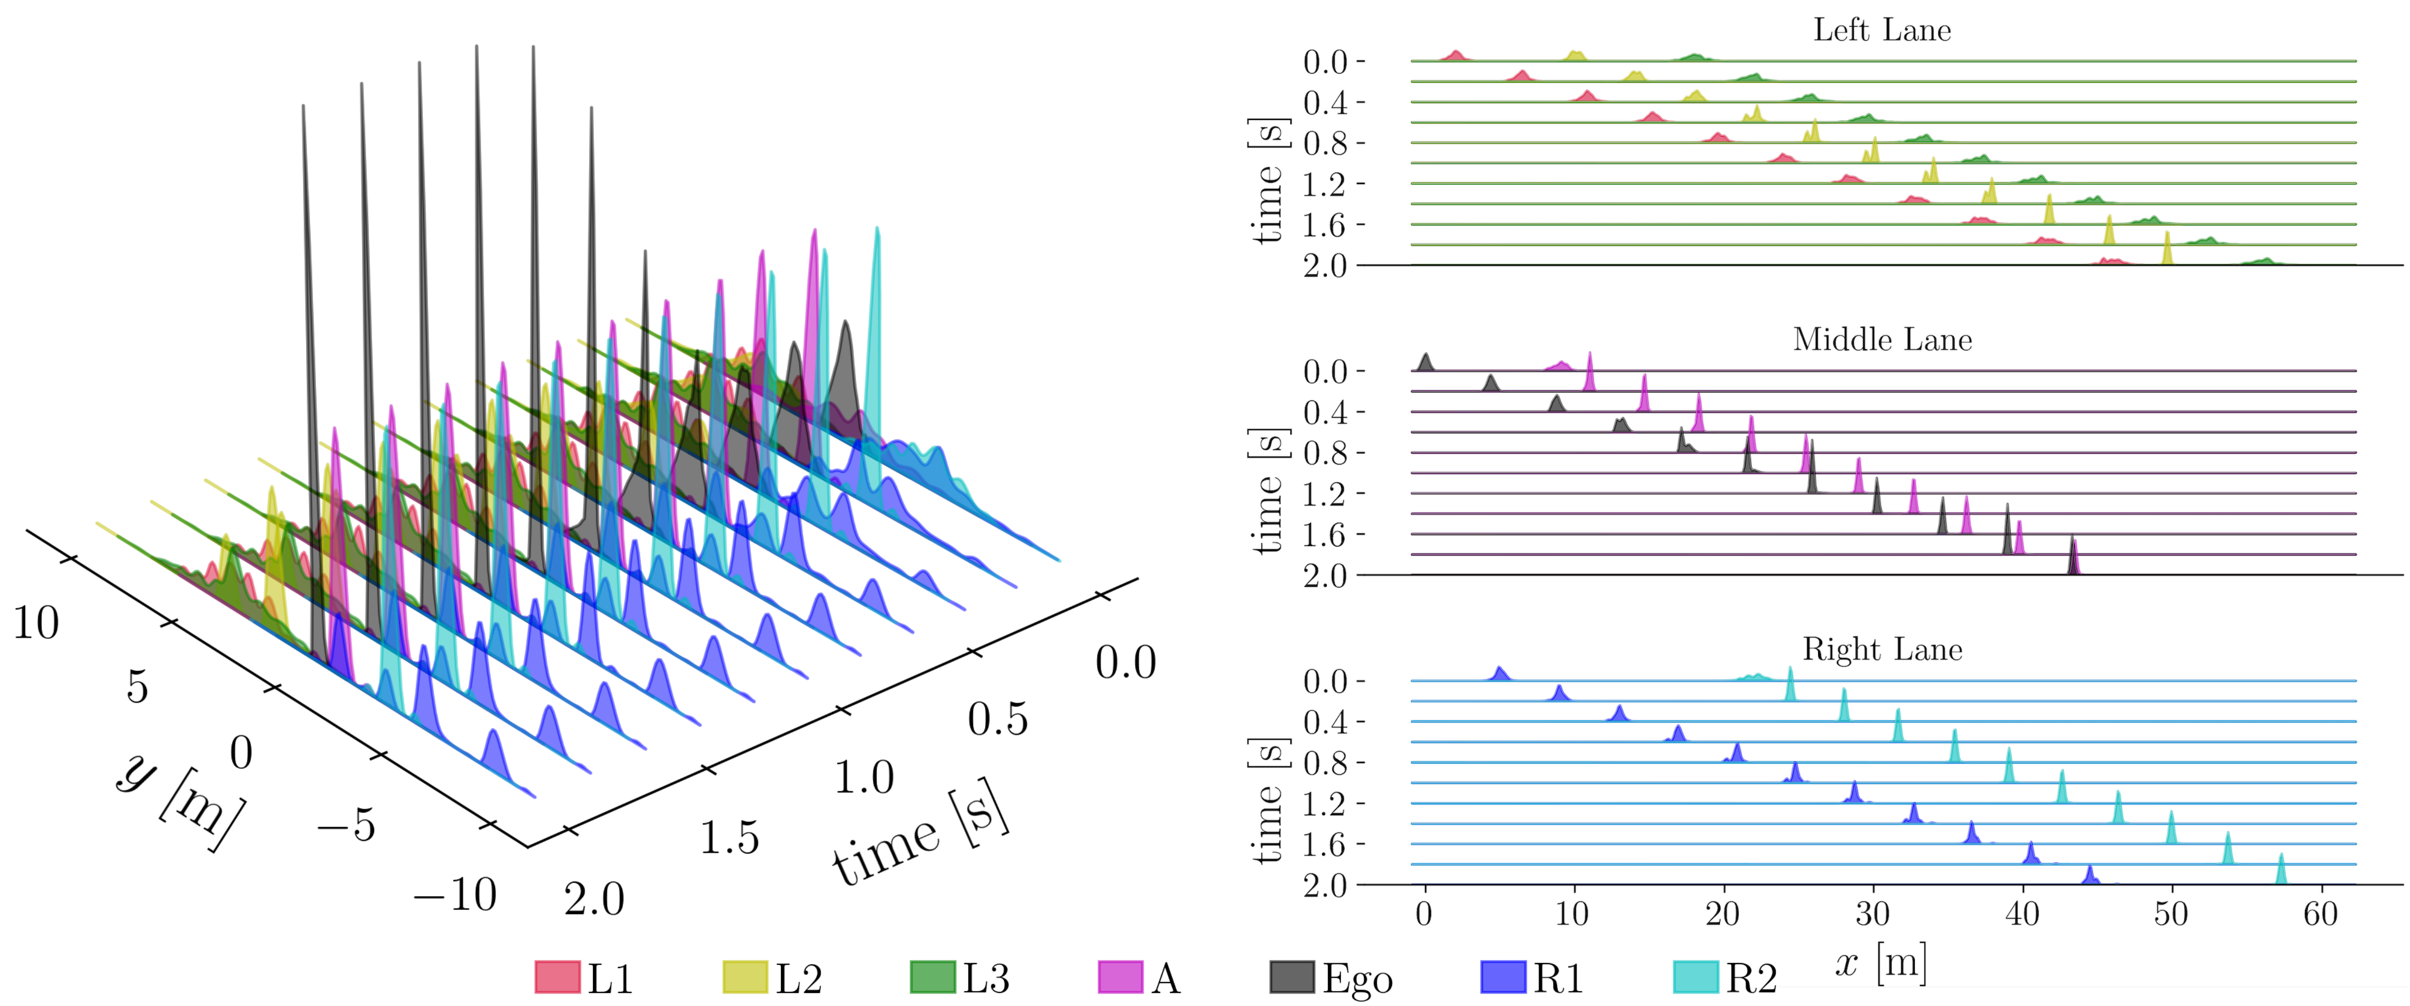

Application: Multi-lane Automated Driving

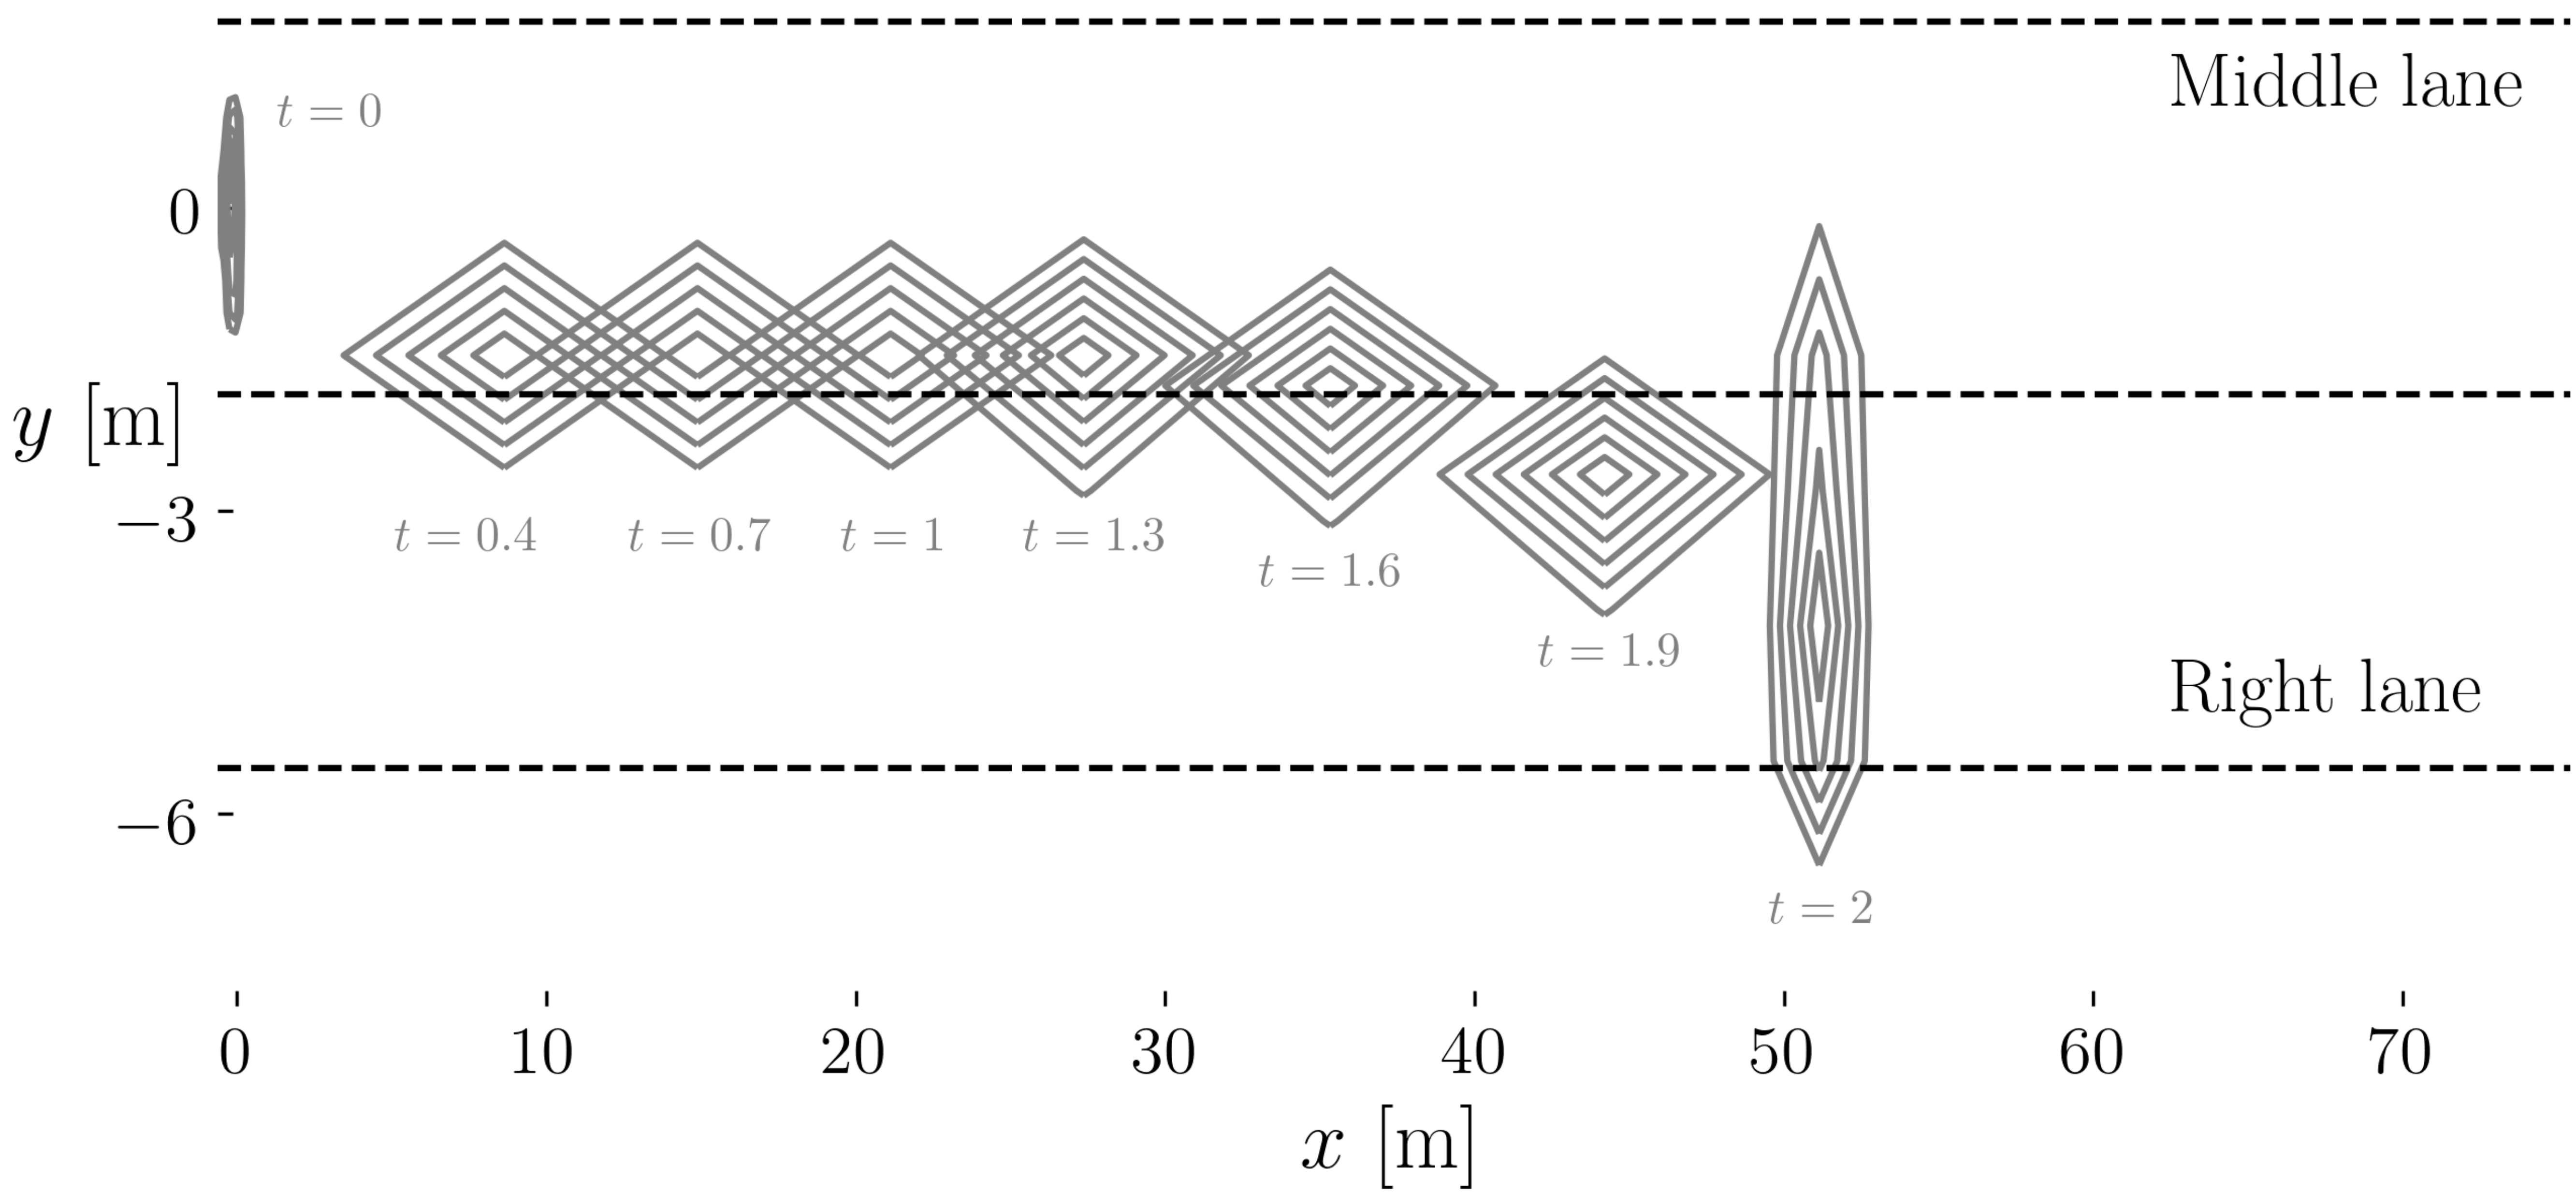

— S. Haddad, K.F. Caluya, A.H., and B. Singh, Prediction and Optimal Feedback Steering of Probability Density Functions for Safe Automated Driving, *IEEE Control Systems Letters, 2020*.

— S. Haddad, A.H., and B. Singh, Prediction and Optimal Feedback Steering of Probability Density Functions for Safe Automated Driving, *IEEE Trans. Control Systems Technology, 2022*.

Application: Multi-lane Automated Driving

Exploit differential flatness: density steering in (Brunovsky) normal coordinates

Markov kernel available but ill-conditioned controllability Gramian

Derived analytical formula for the elements of Gramian inverse

Vector relative degree $\boldsymbol{\pi}$	Computational time [s]	
	using Lyapunov ODE	using Theorem
$(2,2)^\top$	1.9556	0.2995
$(3,2)^\top$	49.7869	6.9294

Details

— S. Haddad, K.F. Caluya, A.H., and B. Singh, Prediction and Optimal Feedback Steering of Probability Density Functions for Safe Automated Driving, *IEEE Control Systems Letters, 2020*.

Application: Controlled Self-assembly (SA)

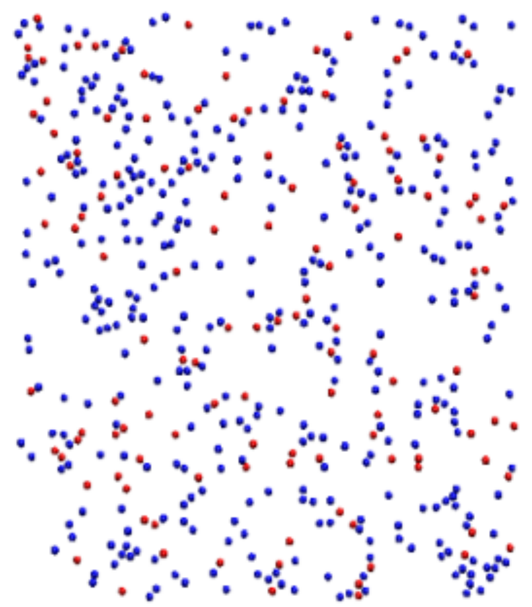

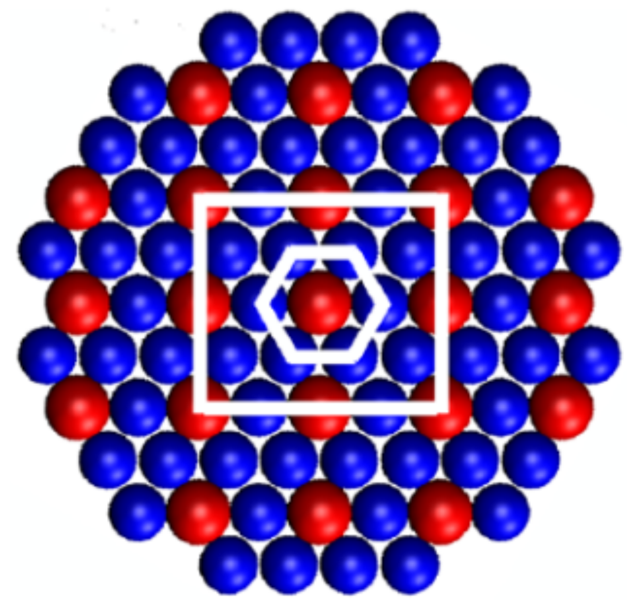

Dispersed particles **Ordered structure**

Applications:

Precision (e.g., sub nm scale) manufacturing of materials with advanced electrical, magnetic or optical properties

Application: Controlled SA

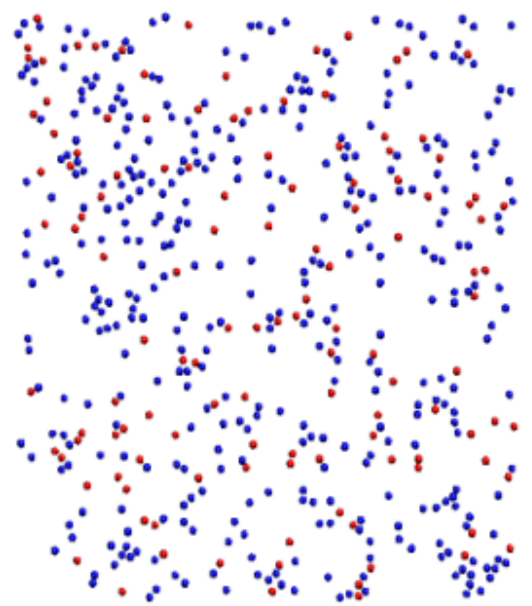

Dispersed particles

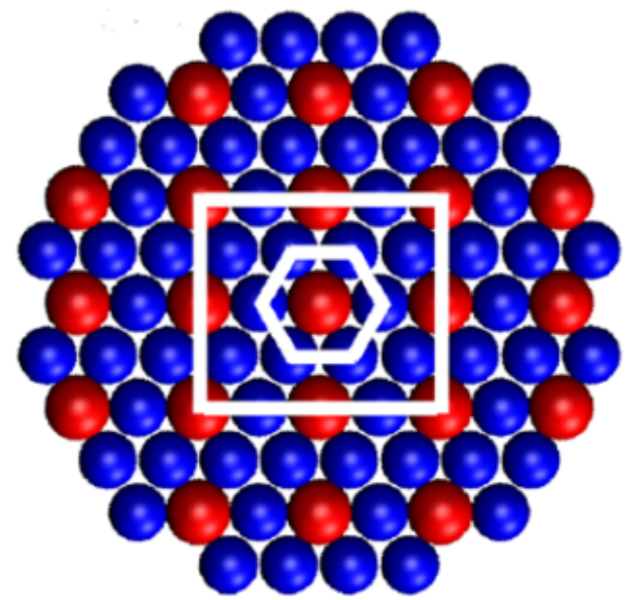

Ordered structure

Typical state variable: $\langle C_6 \rangle \in (0,6)$

Average number of hexagonally close packed neighboring particles in 2D assembly $\rightsquigarrow$ measure of crystallinity order

Typical control variable: u

Electric field voltage

Technical challenge:

Nonlinear + noisy molecular dynamics

$\langle C_6 \rangle$ is a controlled stochastic process

Applications:

Precision (e.g., sub nm scale) manufacturing of materials with advanced electrical, magnetic or optical properties

Controlled SA as PDF Steering

Intuition: $\langle C_6 \rangle \approx 0 \;\Leftrightarrow$ Crystalline disorder

$\langle C_6 \rangle \approx 5 \;\Leftrightarrow$ Crystalline order

Steer the PDF of the stochastic state $\langle C_6 \rangle$ from disordered at $t = t_0 \equiv 0$ to ordered at $t = T \equiv 200\,\mathrm{s}$

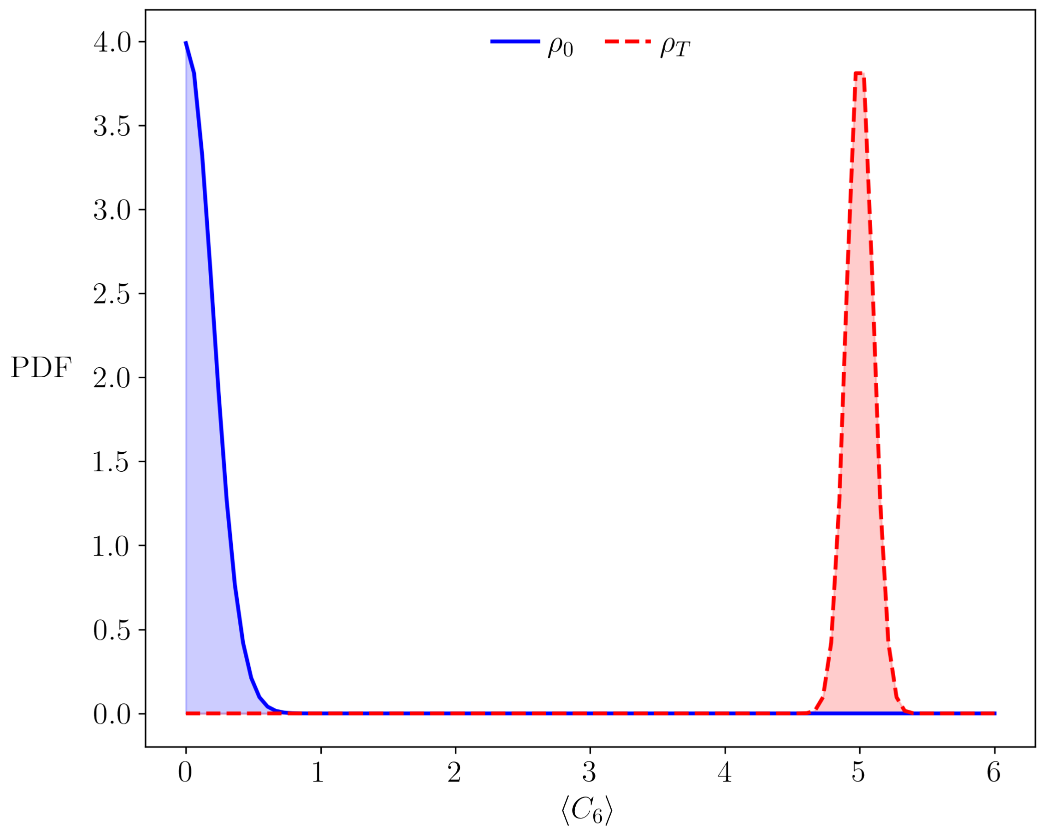

Typical prescribed finite horizon for controlled self-assembly

Endpoint PDF constraints: $\langle C_6 \rangle (t = t_0) \sim \rho_0$ (given)

$\langle C_6 \rangle (t = T) \sim \rho_T$ (given)

Control policy to accomplish the PDF steering: $u = \pi\left(\langle C_6 \rangle, t\right)$

Underdetermined

Minimum Effort SA

Proposed formulation:

$$\inf_{u \in \mathscr{U}} \quad \mathbb{E}_{\mu^u}\left[\int_0^T \frac{1}{2}u^2 \,\mathrm{d}t\right], \quad \mu^u \ll \mathrm{d}x^u$$

drift landscape — diffusion landscape — free energy landscape

$$D_1(x^u, u) := \frac{\partial}{\partial x} D_2(x^u, u) - \frac{\partial}{\partial x} F(x^u, u) \frac{D_2(x^u, u)}{k_B \theta}$$

either from model or learnt from MD simulation data

subject to $\mathrm{d}x^u = D_1(x^u, u)\,\mathrm{d}t + \sqrt{2D_2(x^u, u)}\,\mathrm{d}w,$

$\langle C_6 \rangle$

standard Wiener process

$$x^u(t=0) \sim \mathrm{d}\mu_0 = \rho_0\,\mathrm{d}x^u, \quad x^u(t=T) \sim \mathrm{d}\mu_T = \rho_T\,\mathrm{d}x^u$$

Nonlinear in state, non-affine in control

Conditions for Optimality

$$\frac{\partial \psi}{\partial t} = \frac{1}{2}\left(\pi^{\text{opt}}\right)^2 - \frac{\partial \psi}{\partial x} D_1 - \frac{\partial^2 \psi}{\partial x^{u2}} D_2$$ **HJB PDE**

$$\frac{\partial \rho^u}{\partial t} = -\frac{\partial}{\partial x^u}\left(D_1 \rho^u\right) + \frac{\partial^2}{\partial x^{u2}}\left(D_2 \rho^u\right)$$ **Controlled FPK PDE**

$$\pi^{\text{opt}}(x^u, t) = \frac{\partial \psi}{\partial x^u}\frac{\partial D_1}{\partial u} + \frac{\partial^2 \psi}{\partial x^{u2}}\frac{\partial D_2}{\partial u}$$ **Optimal policy**

$$\rho^u(x^u, t=0) = \rho_0, \quad \rho^u(x^u, t=T) = \rho_T$$ **Boundary conditions**

value function | optimally controlled PDF | optimal policy

to be solved for the triple: $\psi\left(x^u, t\right), \rho^u\left(x^u, t\right), \pi^{\text{opt}}\left(x^u, t\right)$

Solve via PINN: Losses for Training

Loss term for HJB PDE

$$\mathscr{L}_{\psi} = \frac{1}{n}\sum_{i=1}^{n}\left(\frac{\partial \psi}{\partial t}\bigg|_{\boldsymbol{x}_i} - \frac{1}{2}(\pi^{\text{opt}})^2\bigg|_{\boldsymbol{x^u}_i} - +\frac{\partial \psi}{\partial x^u}D_1\bigg|_{\boldsymbol{x}_i^u} - +\frac{\partial^2 \psi}{\partial x^{u2}}D_2\bigg|_{\boldsymbol{x}_i^u}\right)^2$$

Loss term for FPK PDE

$$\mathscr{L}_{\rho^u} = \frac{1}{n}\sum_{i=1}^{n}\left(\frac{\partial \rho^u}{\partial t}\bigg|_{\boldsymbol{x}_i^u} + \frac{\partial}{\partial x^u}\left(D_1\rho^u\right)\bigg|_{\boldsymbol{x}_i^u} - \frac{\partial^2}{\partial x^{u2}}\left(D_2\rho^u\right)\bigg|_{\boldsymbol{x}_i^u}\right)^2$$

Loss term for policy equation

$$\mathscr{L}_{\pi^{\text{opt}}} = \frac{1}{n}\sum_{i=1}^{n}\left(\pi^{\text{opt}}\Big|_{\boldsymbol{x}_i^u} - \frac{\partial \psi}{\partial x^u}\frac{\partial D_1}{\partial u}\bigg|_{\boldsymbol{x}_i^u} - \frac{\partial^2 \psi}{\partial x^{u2}}\frac{\partial D_2}{\partial u}\bigg|_{\boldsymbol{x}_i^u}\right)^2$$

Loss term for initial condition

$$\mathscr{L}_{\rho_0^u} = \frac{1}{n}\sum_{i=1}^{n}\left(\rho^u\Big|_{t=0} - \rho_0^u(x)\right)^2$$

Loss term for terminal condition

$$\mathscr{L}_{\rho_T^u} = \frac{1}{n}\sum_{i=1}^{n}\left(\rho^u\Big|_{t=T} - \rho_T^u(x)\right)^2$$

PINN Architecture

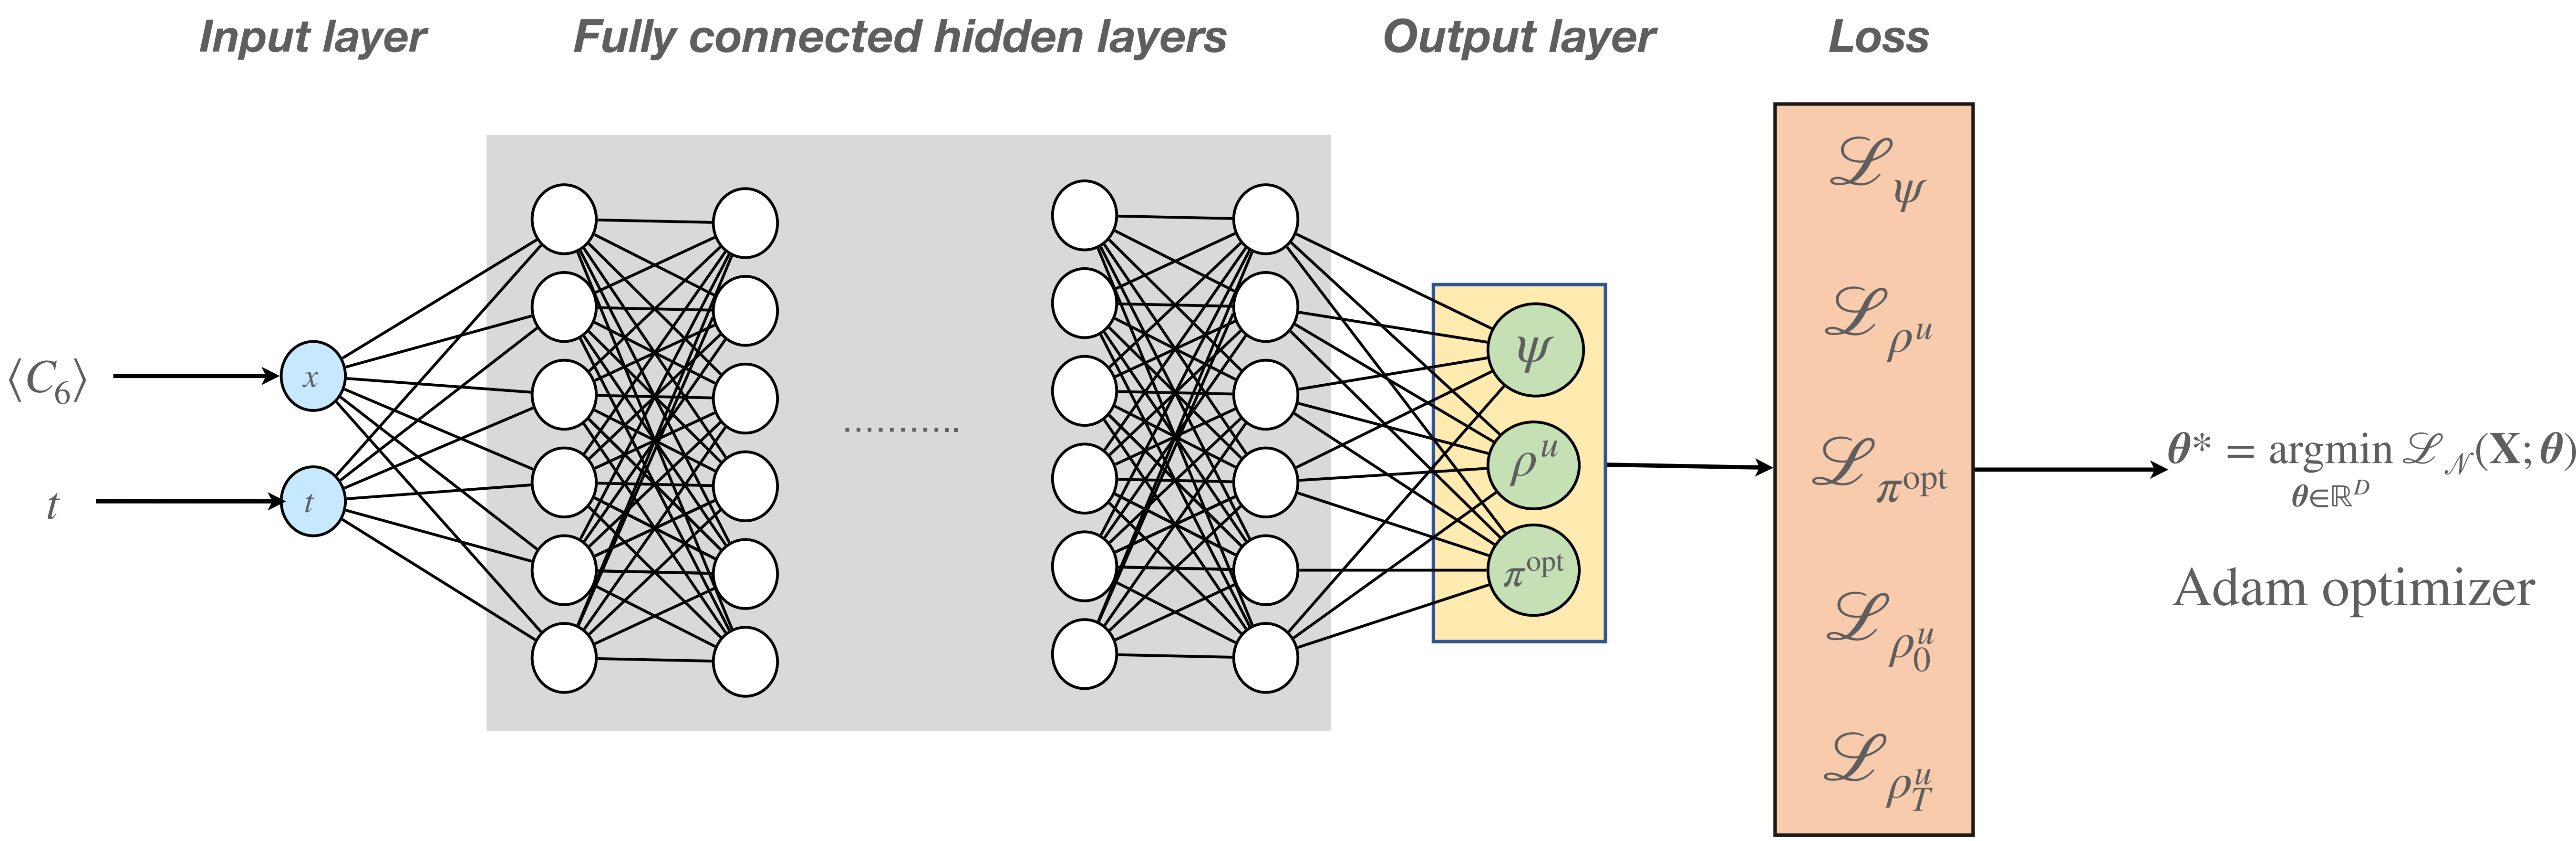

$$\mathscr{L}_{\mathscr{N}} = \mathscr{L}_{\psi} + \mathscr{L}_{\rho^u} + \mathscr{L}_{\pi^{\text{opt}}} + \mathscr{L}_{\rho_0^u} + \mathscr{L}_{\rho_T^u}$$

[Lu Lu et al, 2021] [Niaki et al, 2021]

Training of the PINN

Benchmark controlled self-assembly system: [Y Xue, et al, *IEEE Trans. Control Sys. Technology*, 2014]

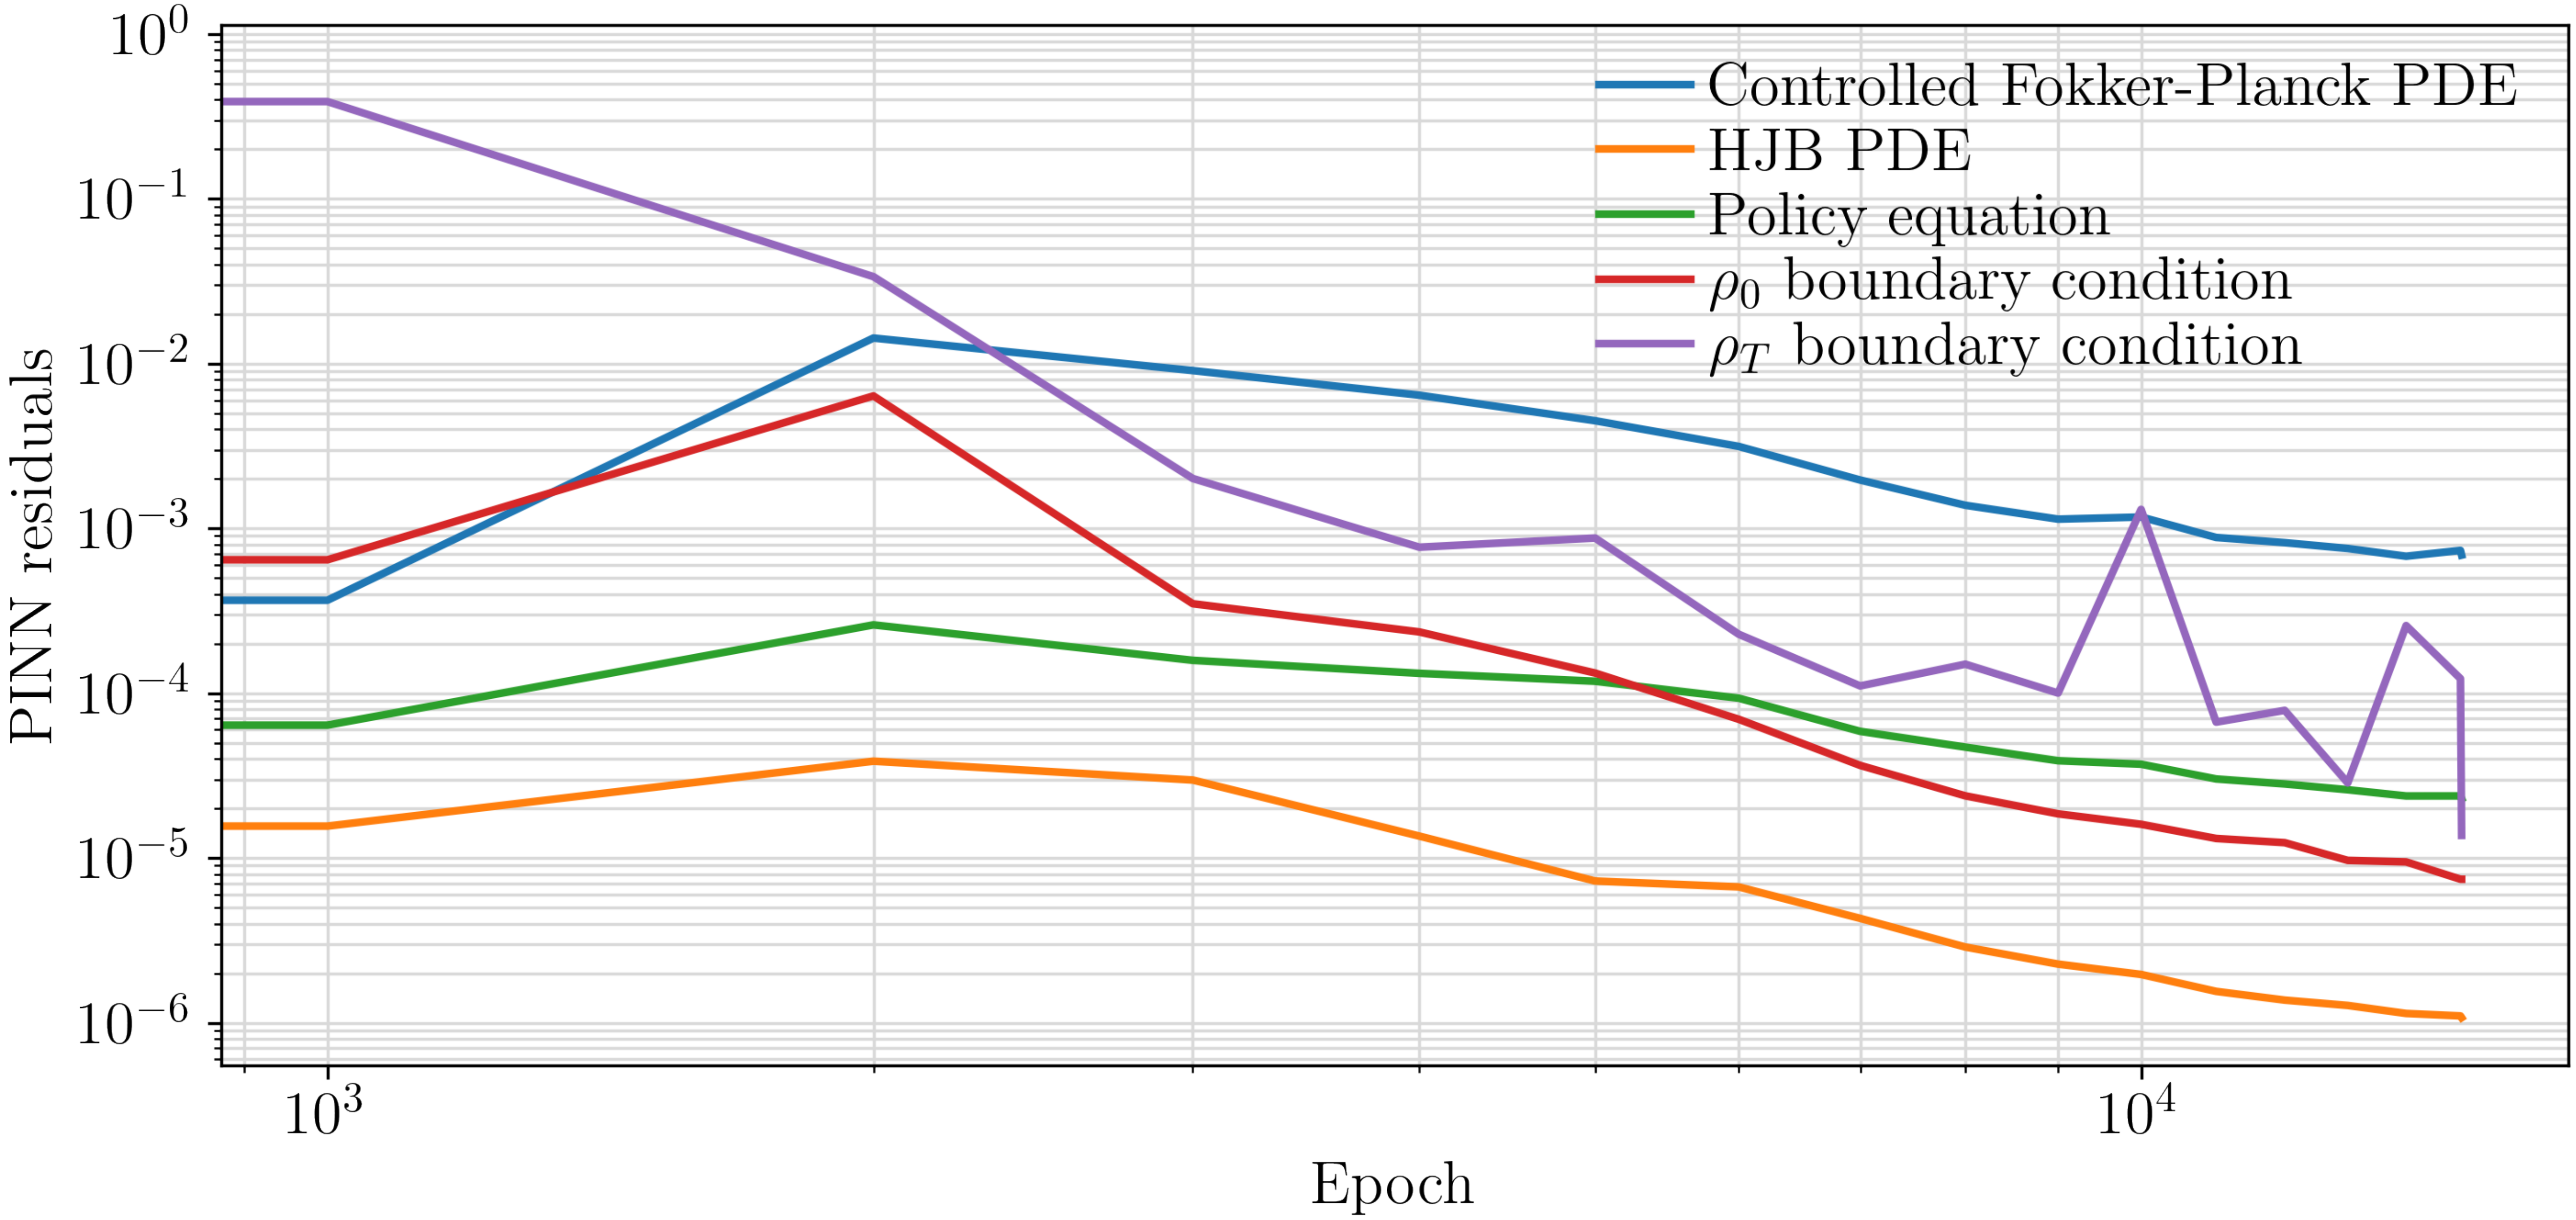

Optimal Policy

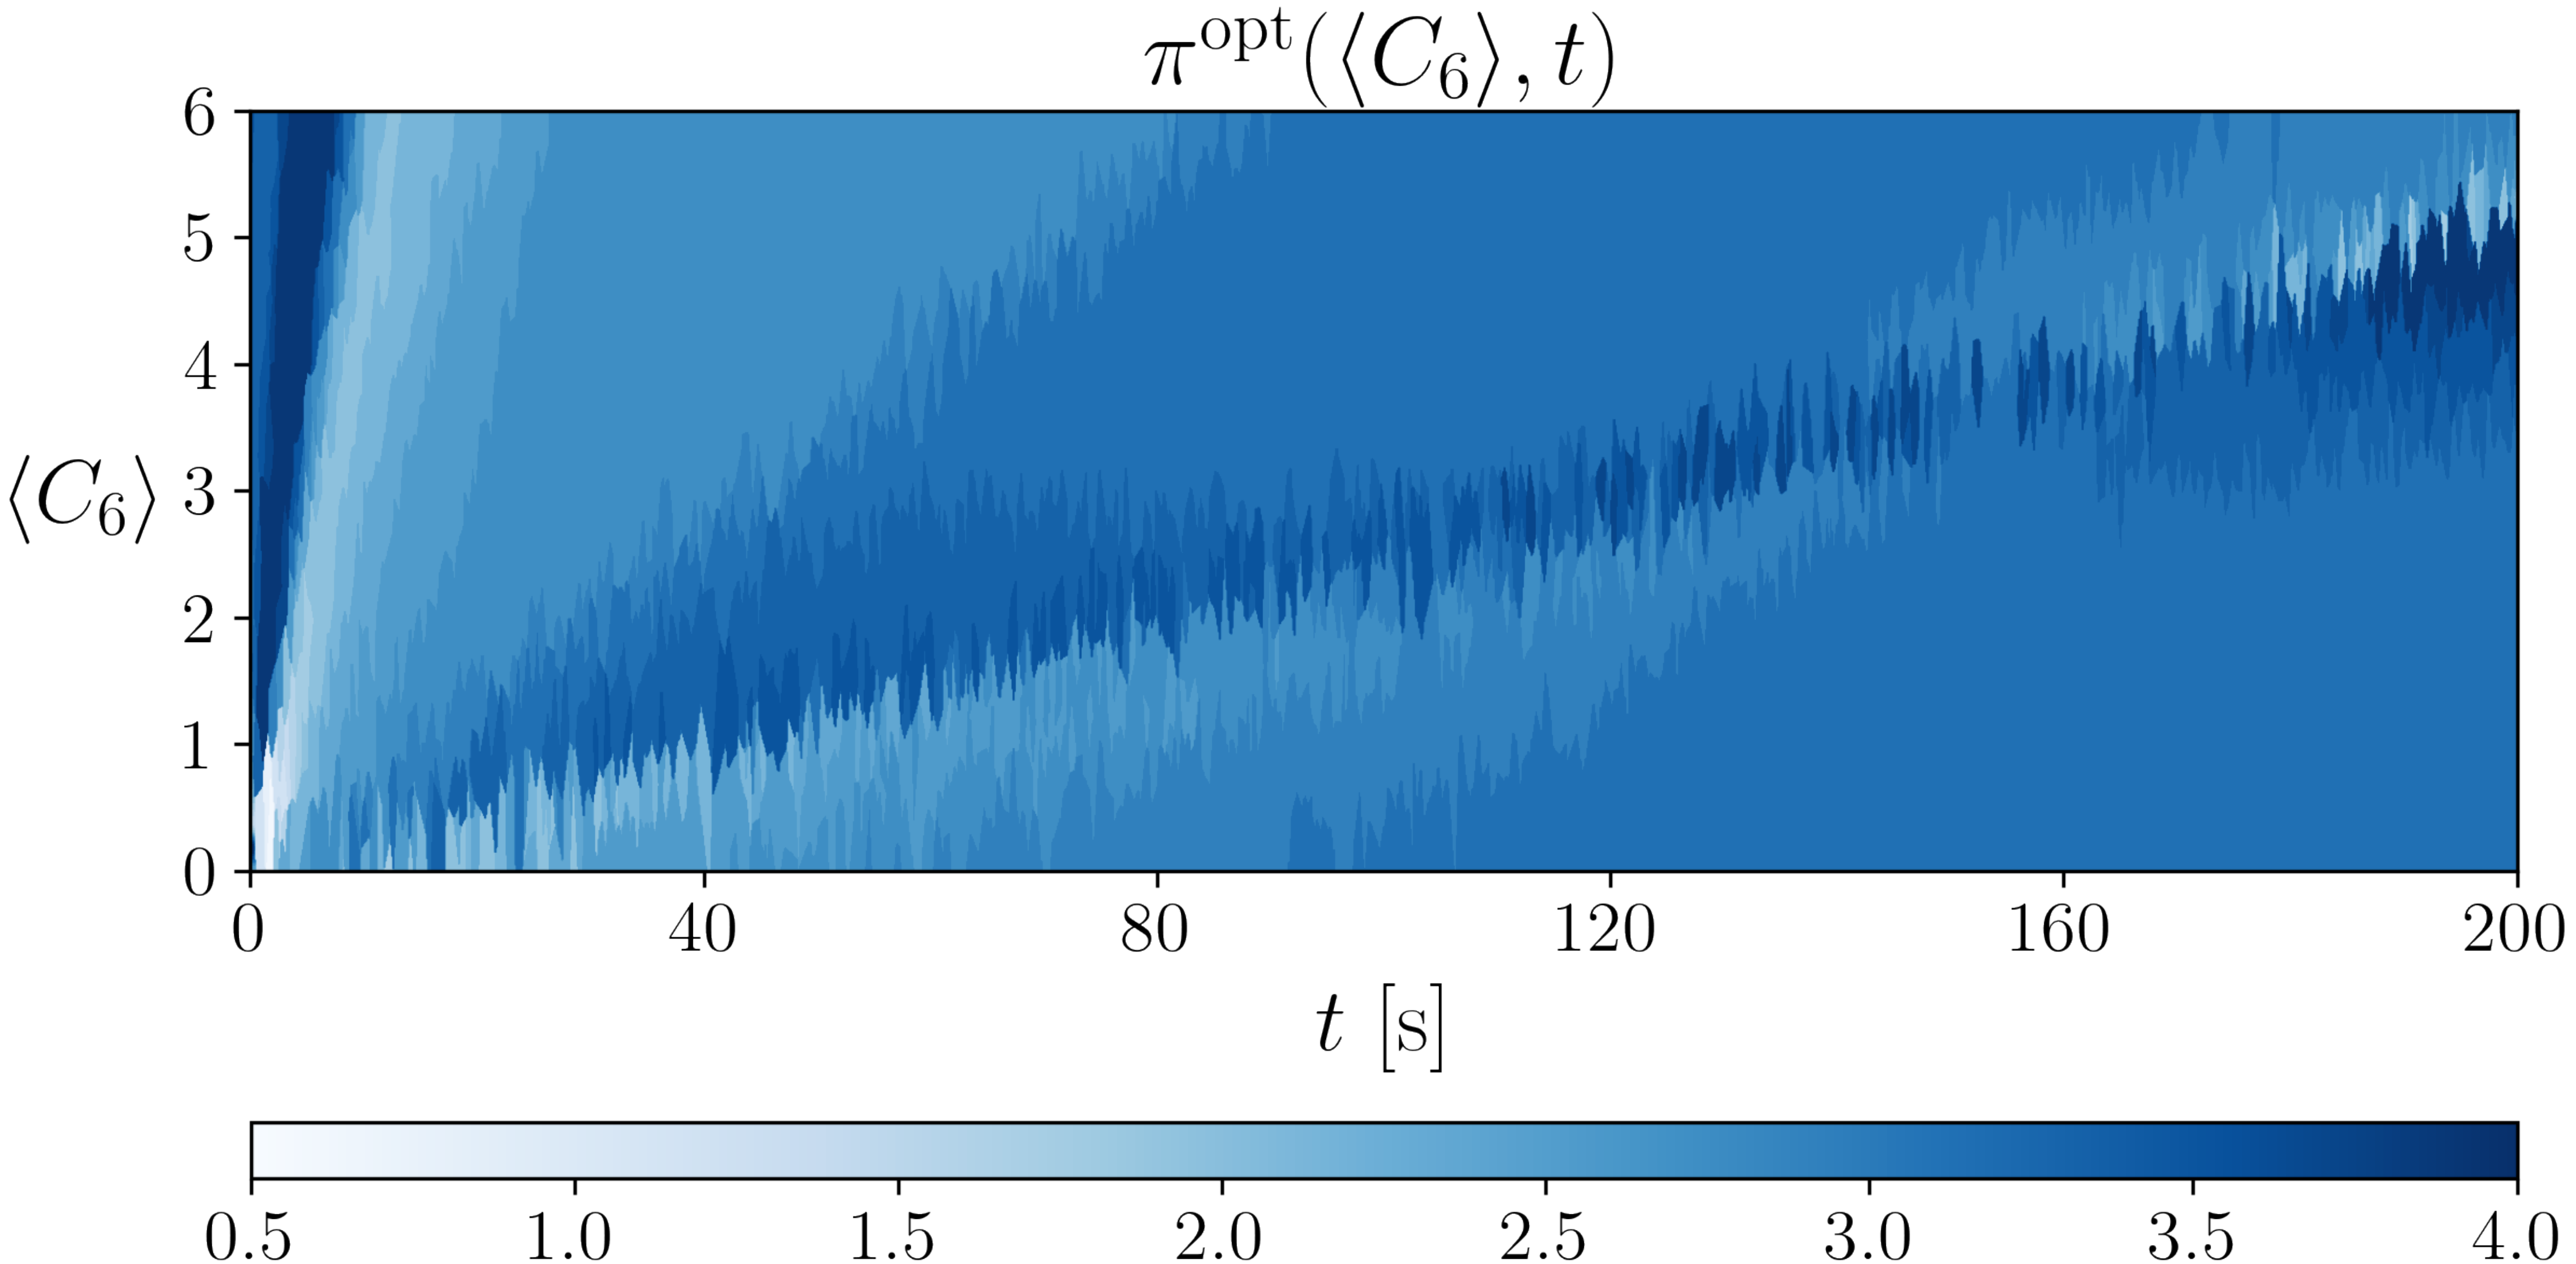

Value Function

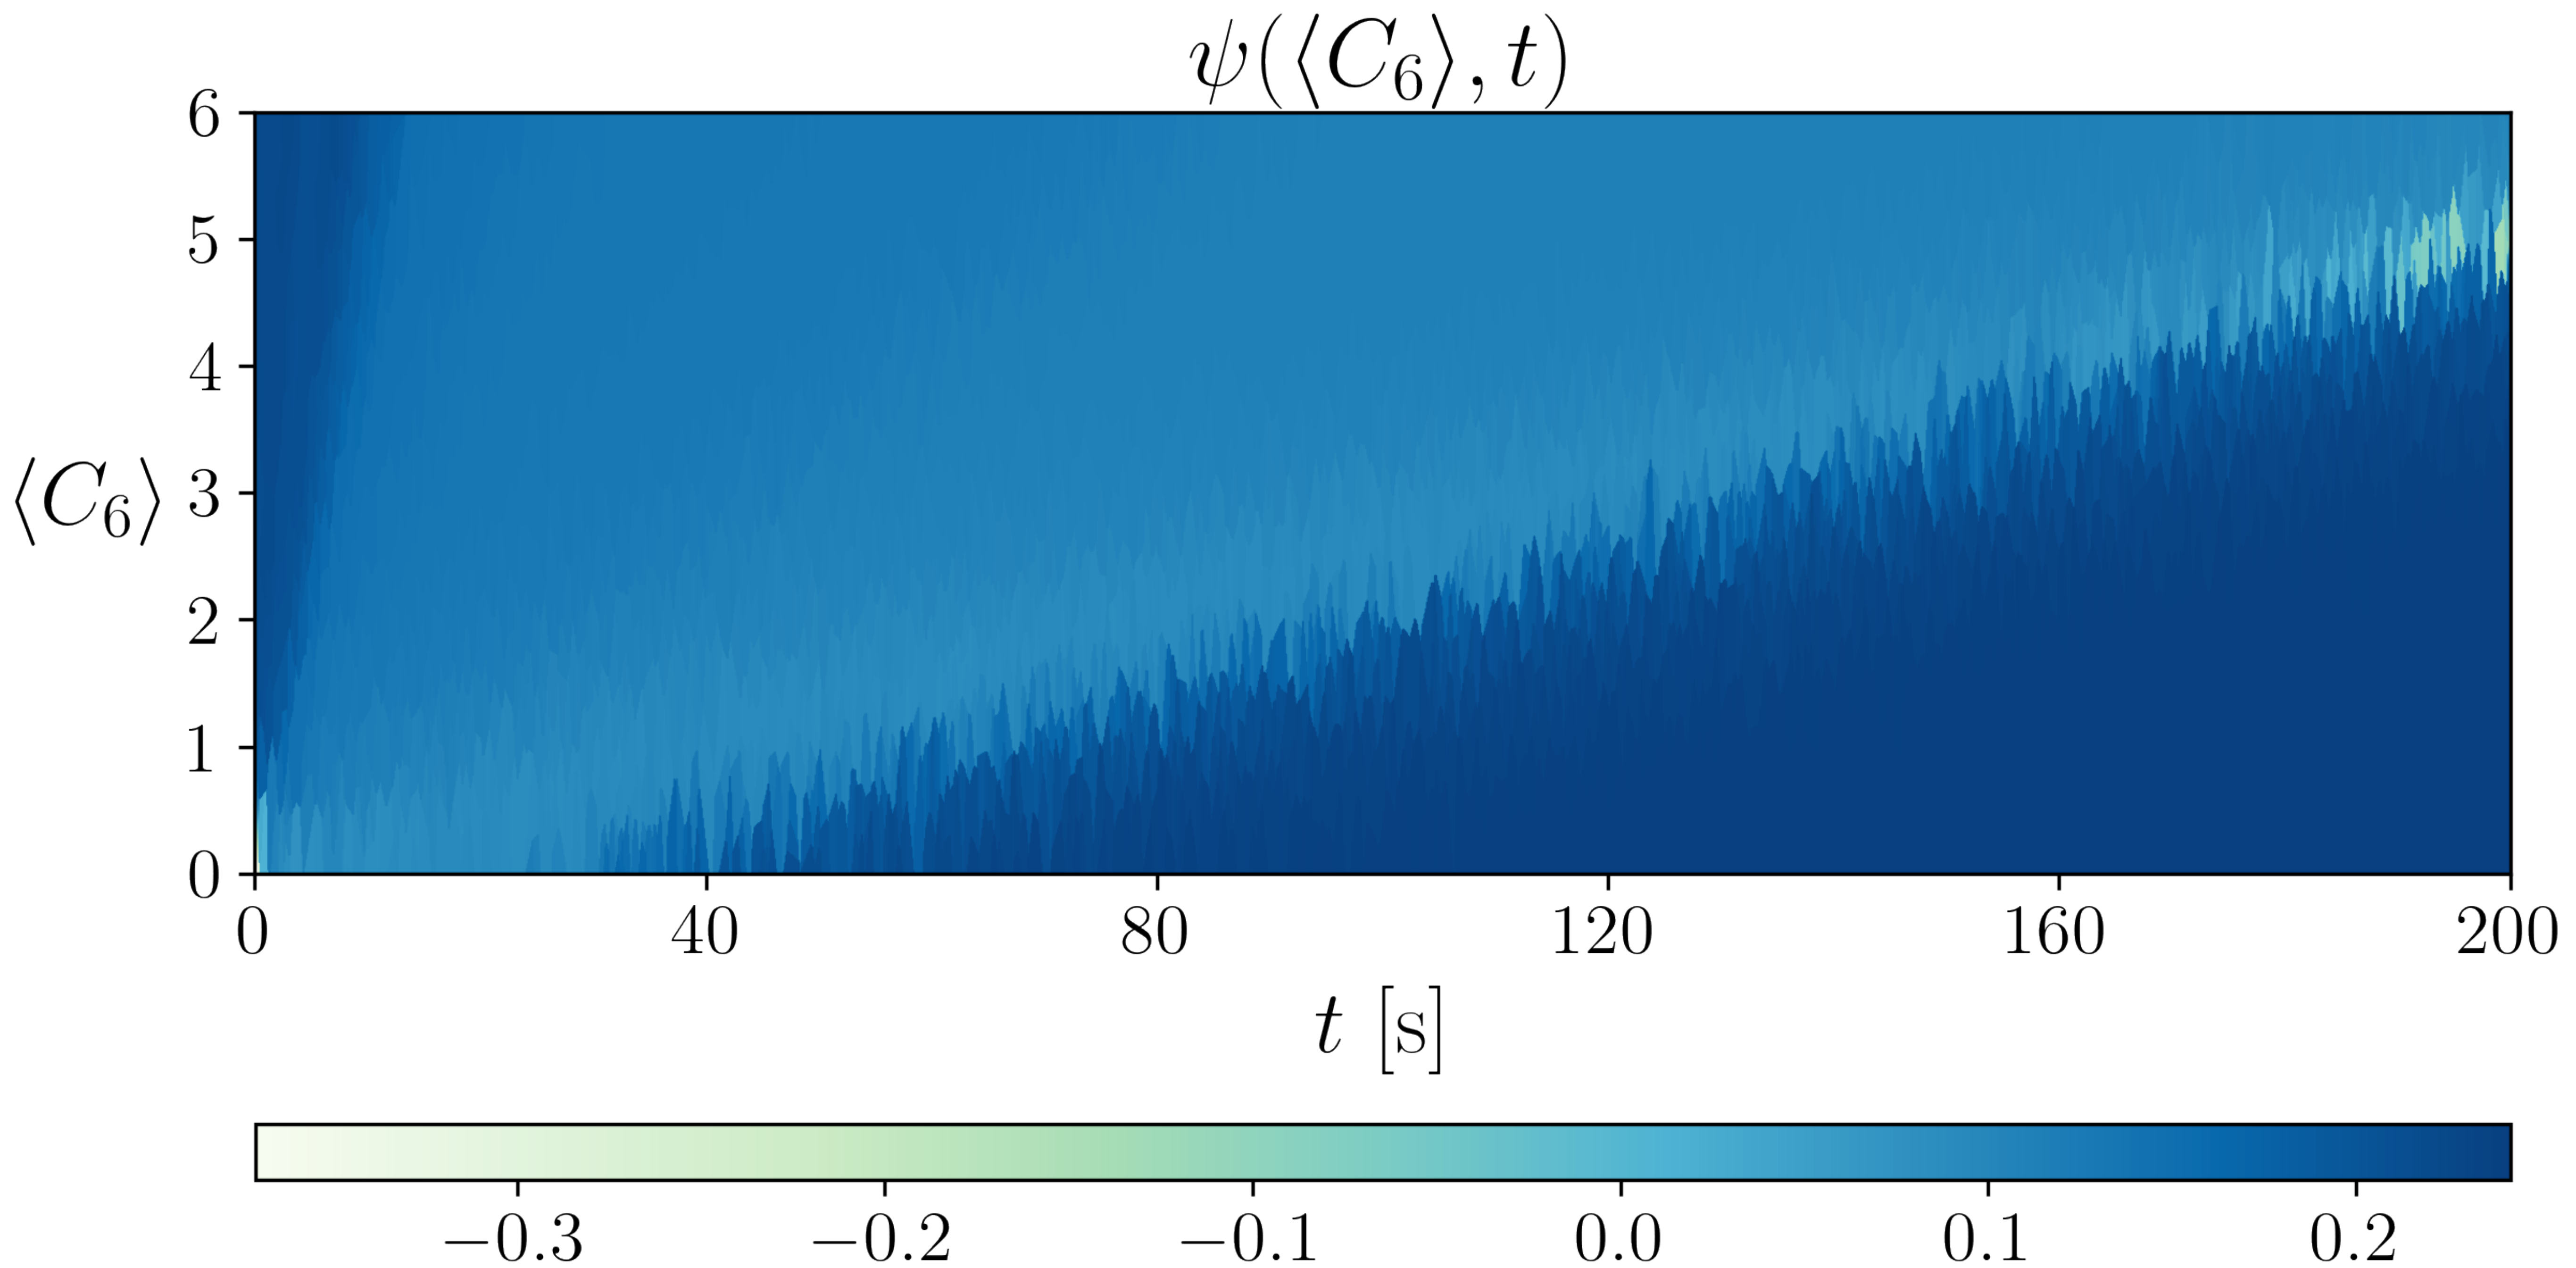

Optimally Controlled State PDFs

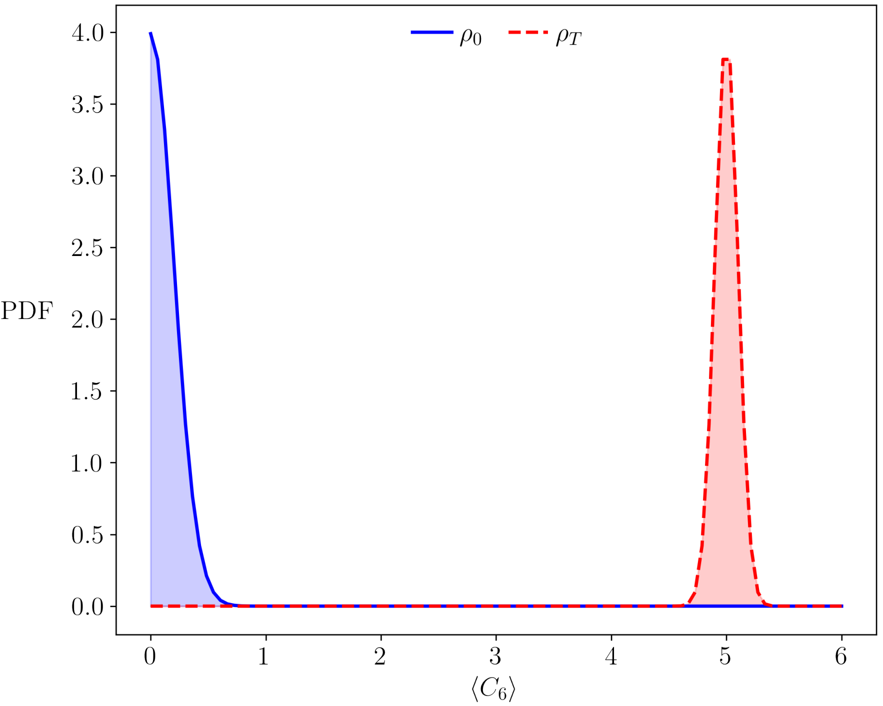

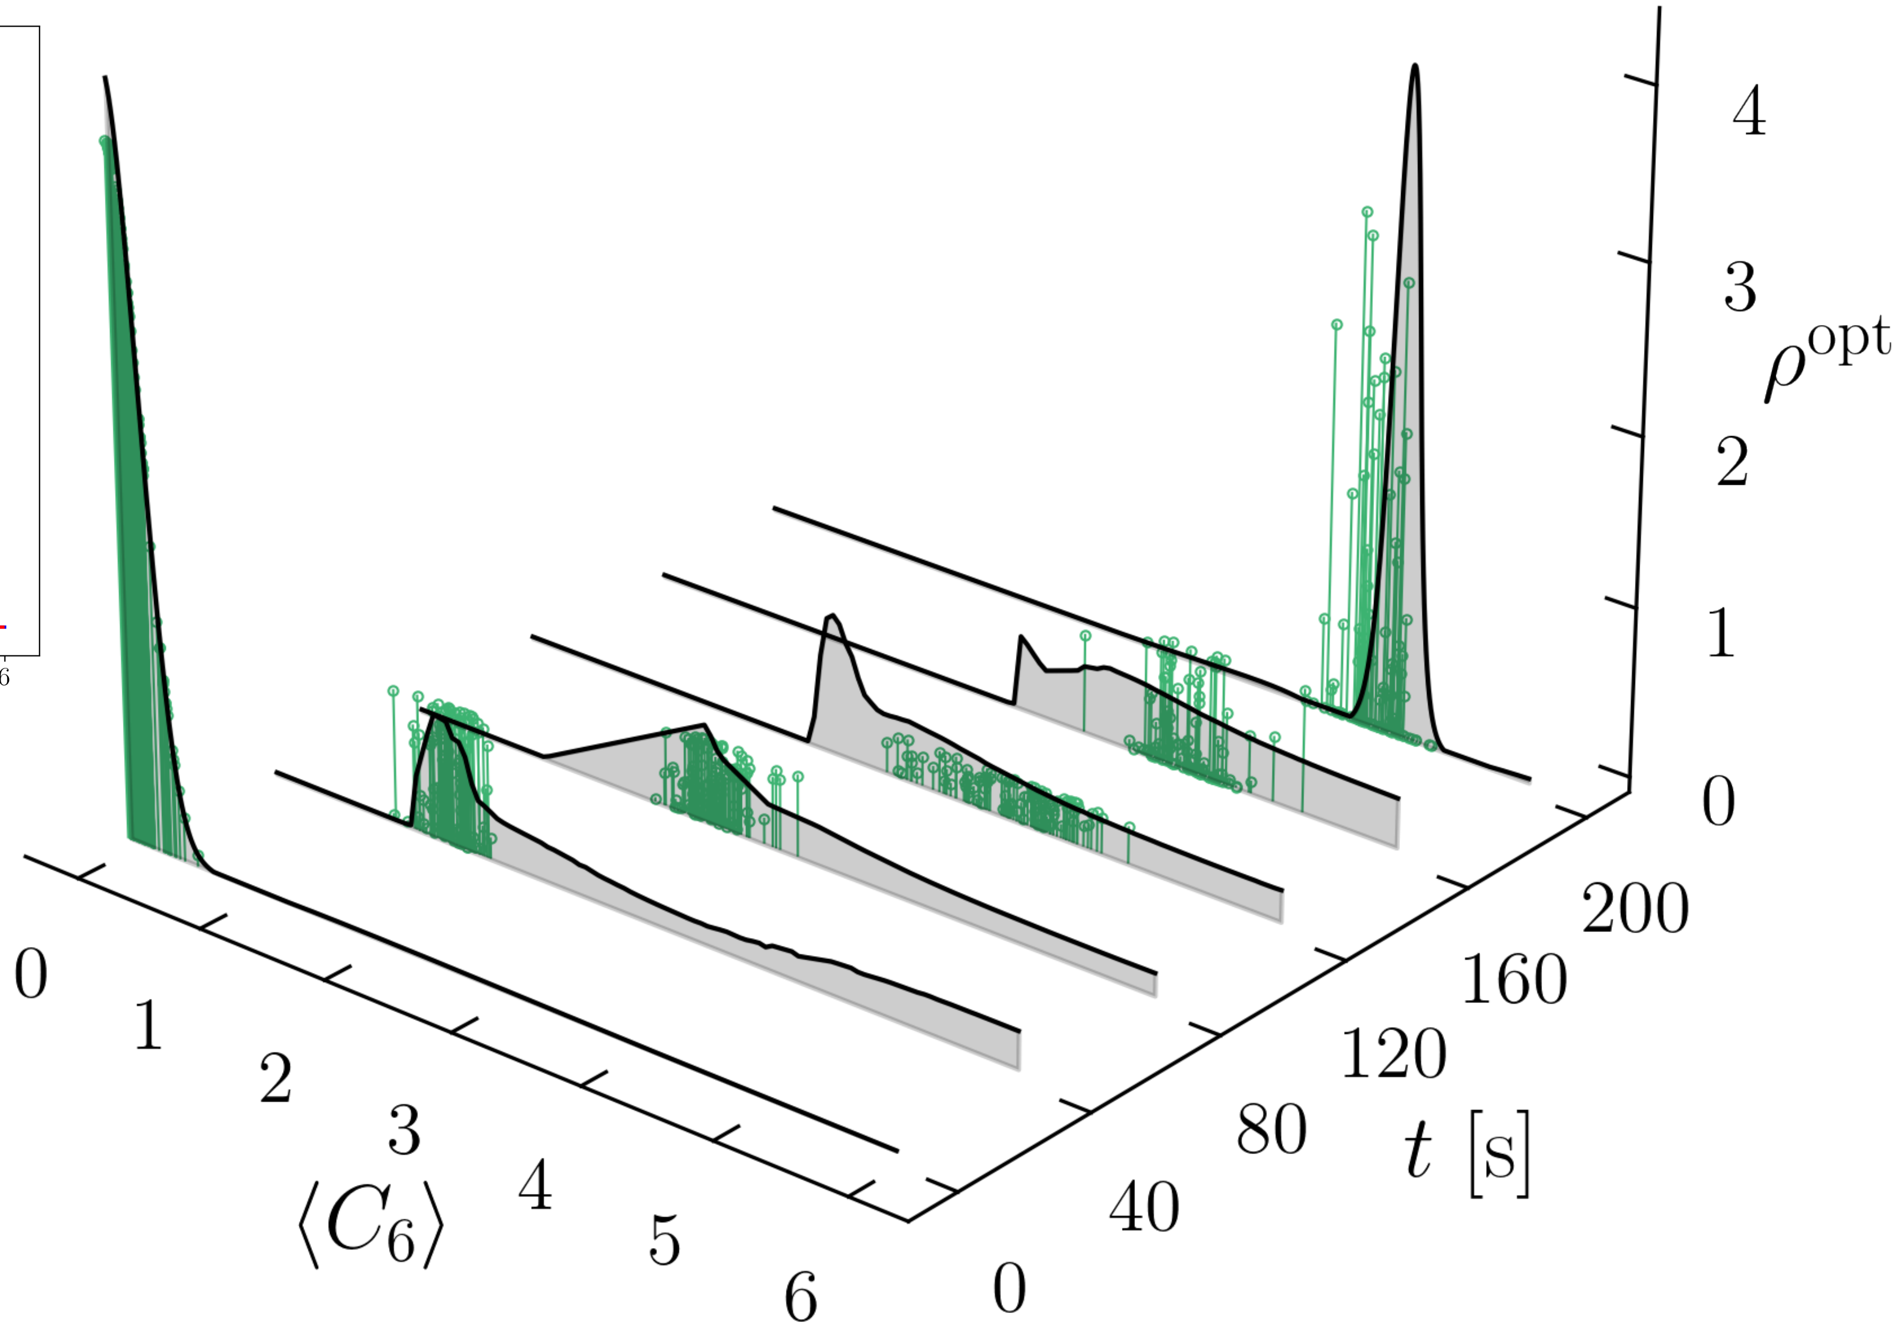

Optimal State and Optimal Control Sample Paths

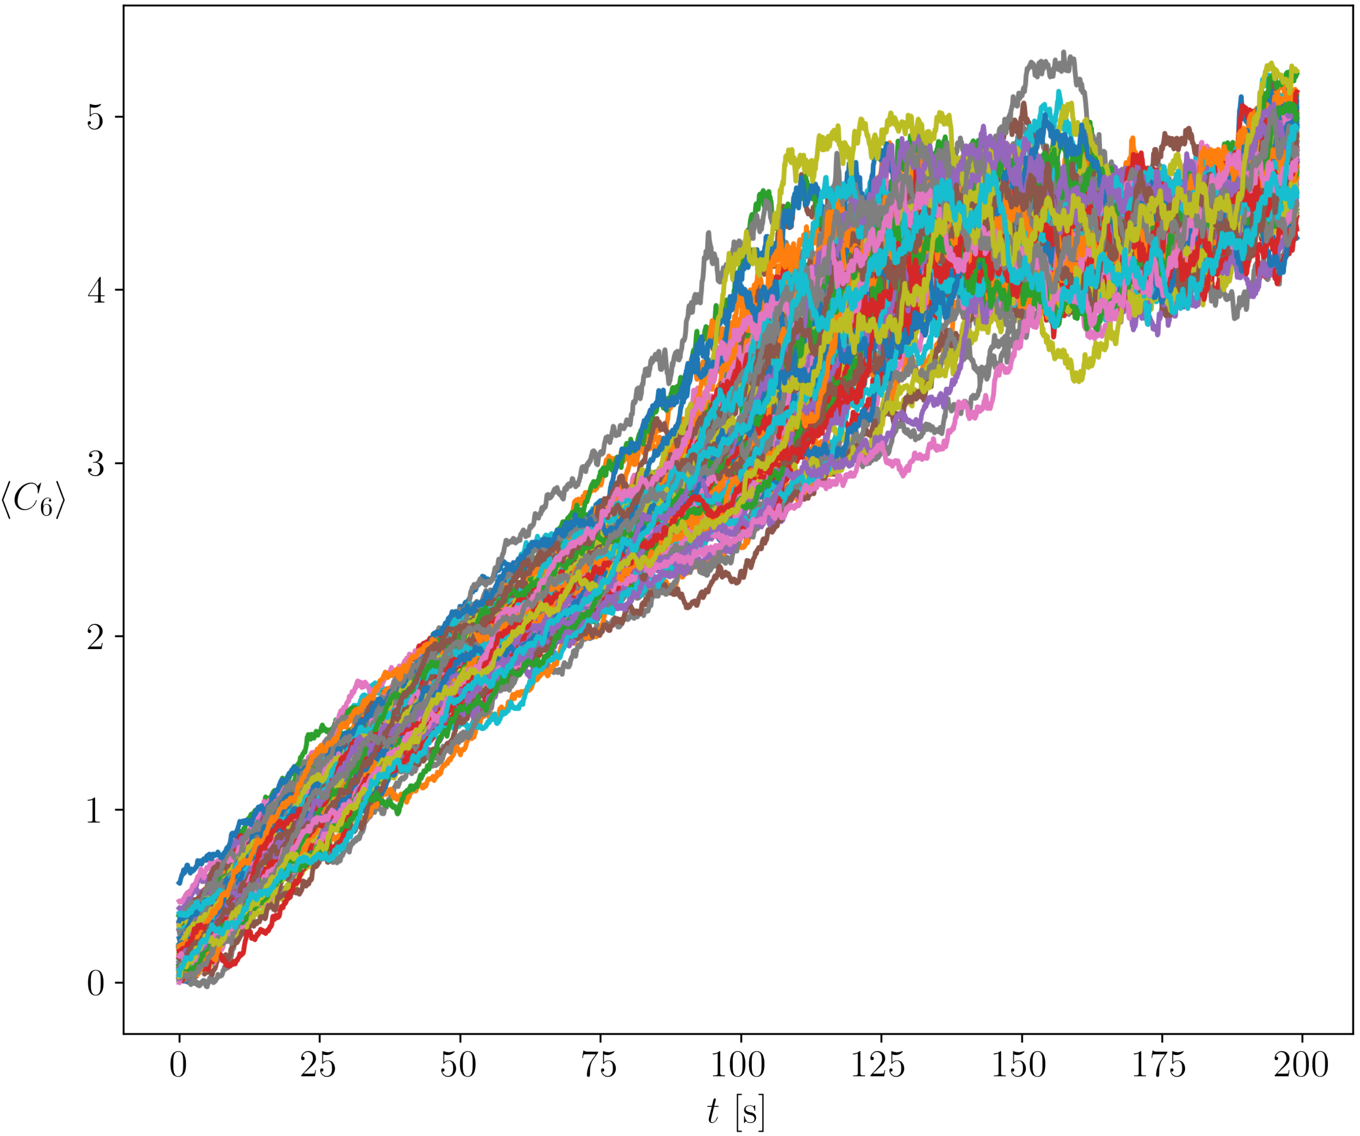

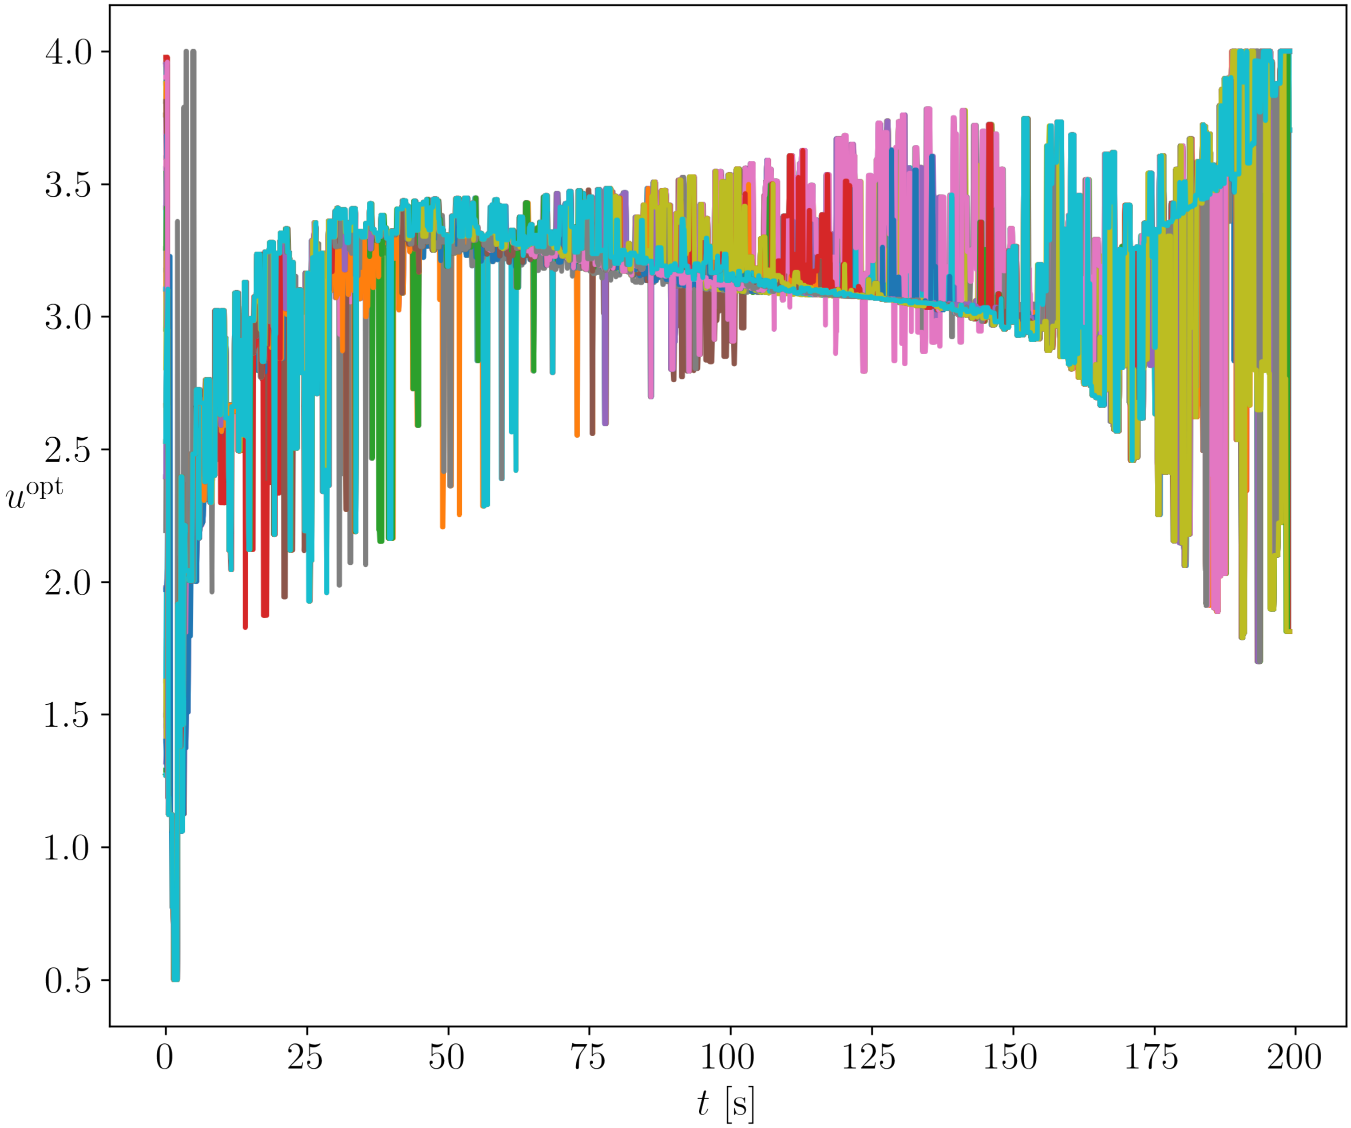

GSBP Conditions for Optimality with m Inputs

$m+2$ coupled PDEs with endpoint boundary conditions:

Drift coefficient Diffusion tensor

$$\frac{\partial \psi}{\partial t} = \frac{1}{2}\|\boldsymbol{u}_{\text{opt}}\|_2^2 - \langle \nabla_{\boldsymbol{x}}\psi, \boldsymbol{f}\rangle - \langle \boldsymbol{G}, \mathbf{Hess}(\psi)\rangle,$$

$$\frac{\partial \rho^{\boldsymbol{u}}_{\text{opt}}}{\partial t} = -\nabla.(\rho^{\boldsymbol{u}}_{\text{opt}}\boldsymbol{f}) + \langle \boldsymbol{G}, \mathbf{Hess}(\rho^{\boldsymbol{u}}_{\text{opt}})\rangle,$$

$$\boldsymbol{u}_{\text{opt}} = \nabla_{\boldsymbol{u}_{\text{opt}}}\left(\langle \nabla_{\boldsymbol{x}}\psi, \boldsymbol{f}\rangle + \langle \boldsymbol{G}, \mathbf{Hess}(\psi)\rangle\right),$$

$$\rho^{\boldsymbol{u}}_{\text{opt}}(0,\boldsymbol{x}) = \rho_0, \qquad \rho^{\boldsymbol{u}}_{\text{opt}}(T,\boldsymbol{x}) = \rho_T,$$

Cf. classical SBP: two coupled PDEs + optimal policy explicit in value fn ψ

Data-driven GSBP for Colloidal SA

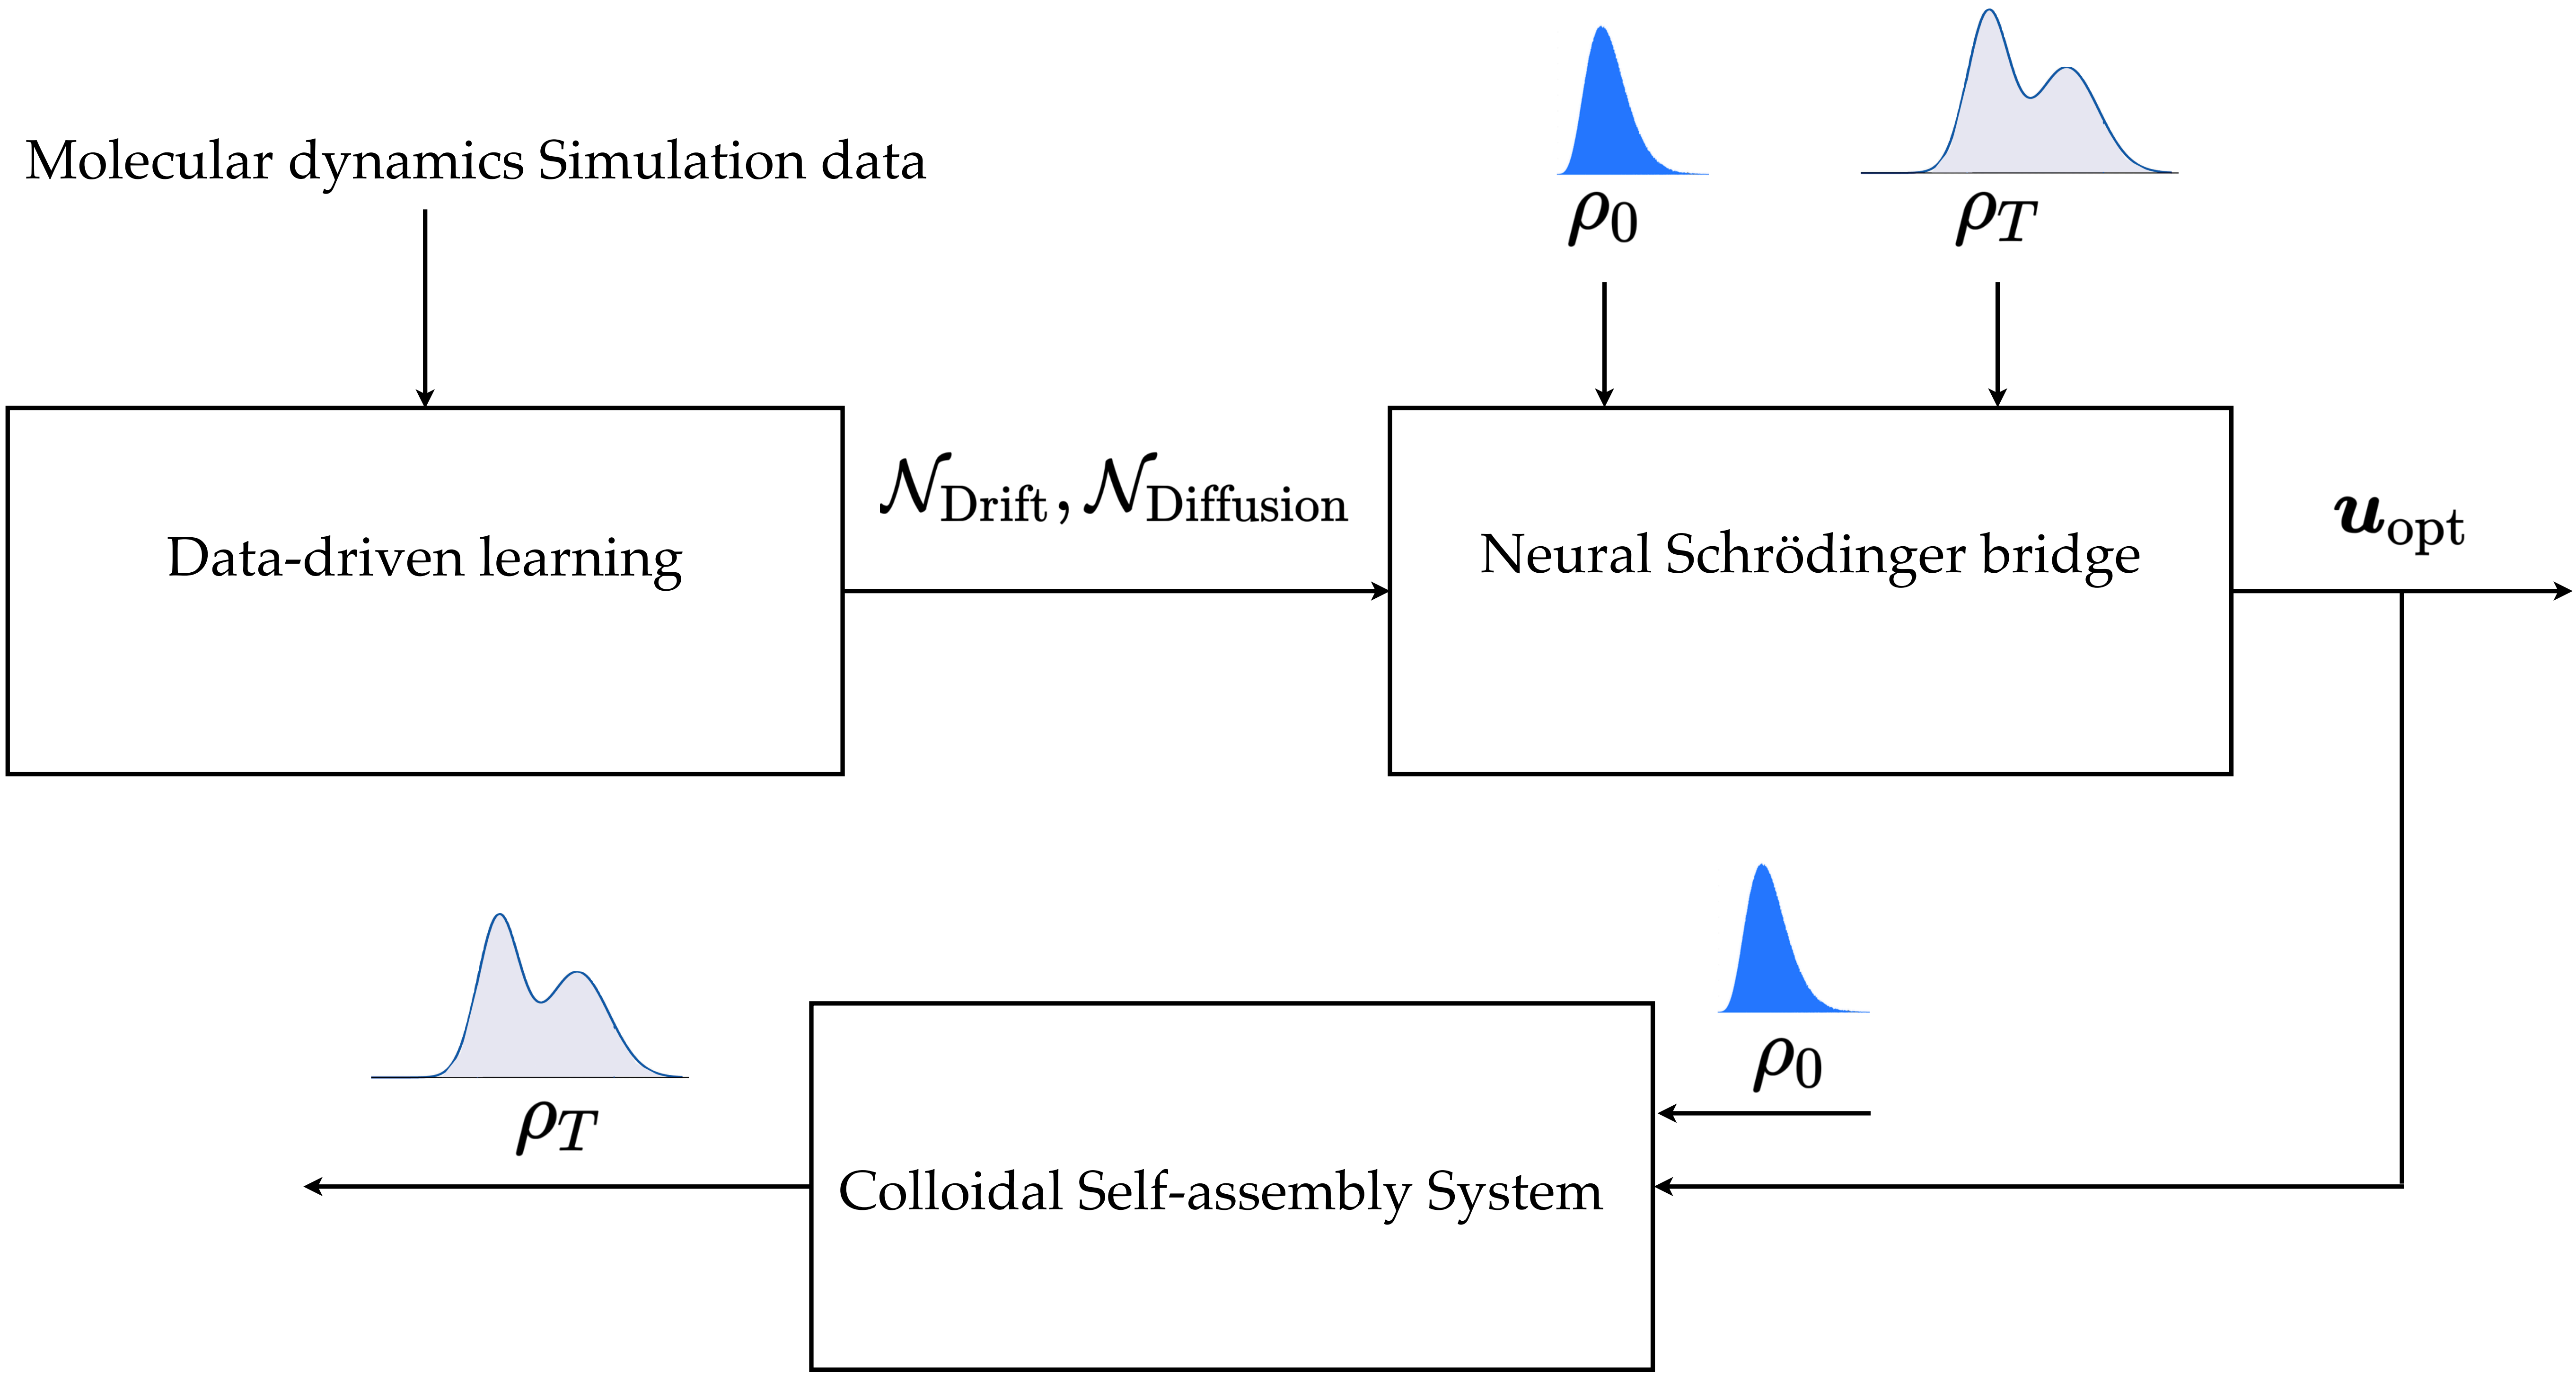

Architecture for Data-driven GSBP

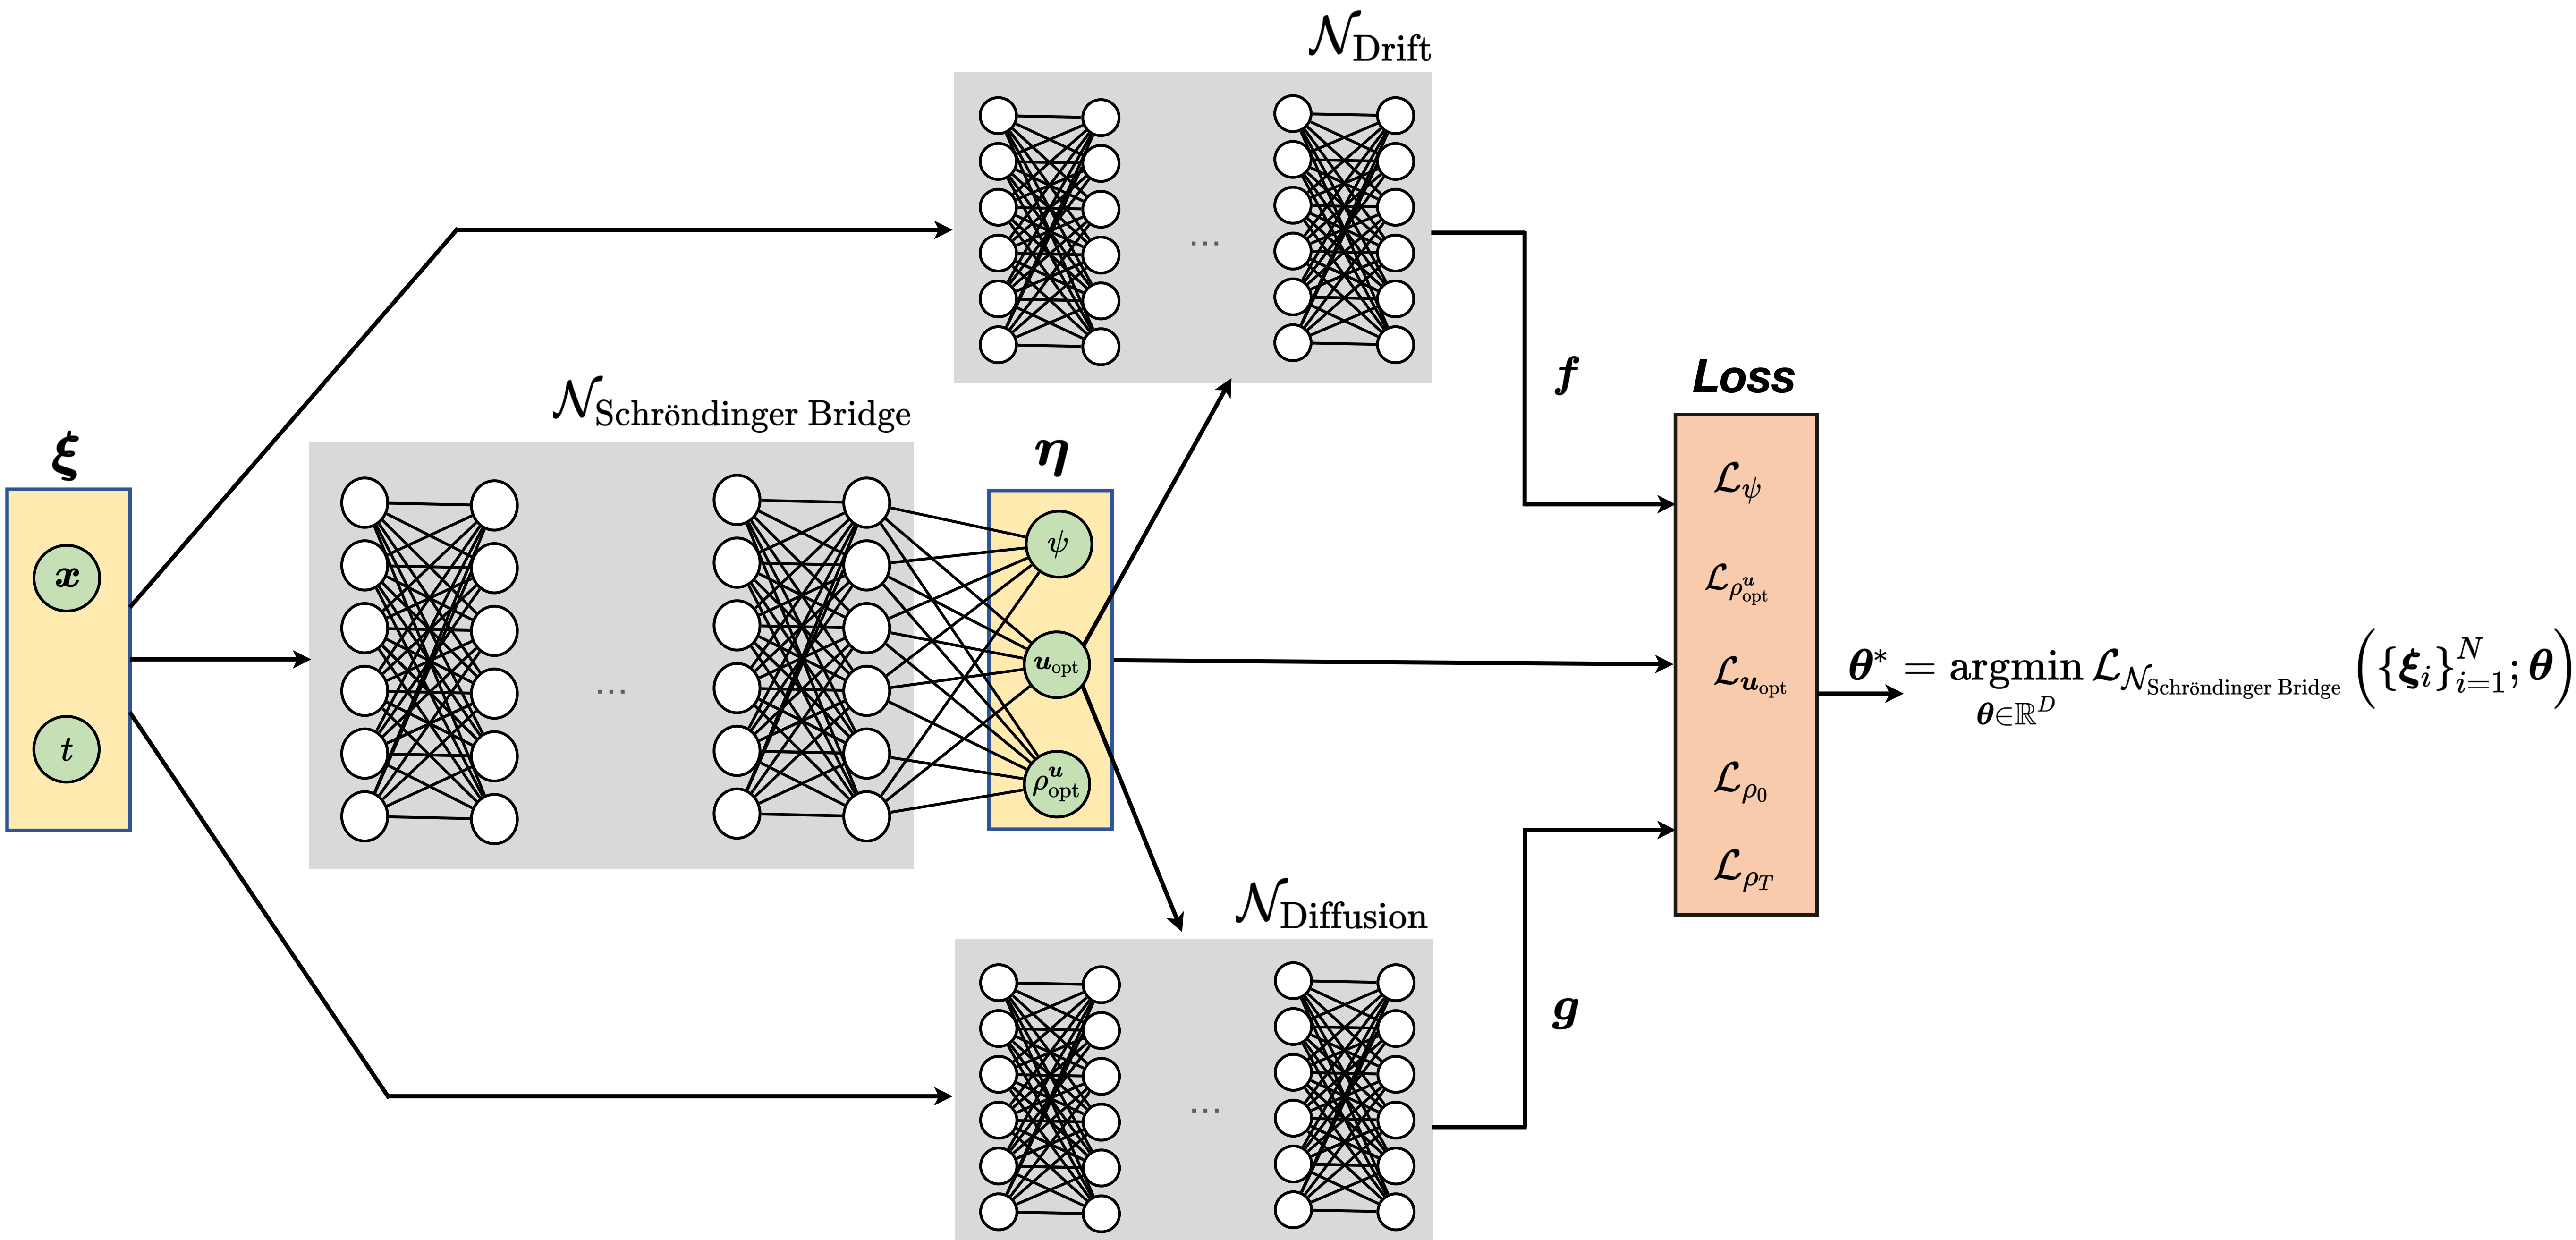

Sinkhorn Losses for Boundary Conditions

$$W_{\varepsilon}^{2}(\mu_0,\mu_1) := \inf_{\pi\in\Pi_2(\mu_0,\mu_T)} \int_{\mathbb{R}^n\times\mathbb{R}^n} \left\{\|\boldsymbol{x}-\boldsymbol{y}\|_2^2 + \varepsilon\log\pi(\boldsymbol{x},\boldsymbol{y})\right\}\mathrm{d}\pi(\boldsymbol{x},\boldsymbol{y})$$

For boundary conditions, use Sinkhorn losses: $\mathcal{L}_{\rho_i} := W_{\varepsilon}^{2}\left(\rho_i, \rho_i^{\text{epoch index}}(\boldsymbol{\theta})\right)$

Implementation friendly for PINN training:

$$\texttt{Autodiff}_{\boldsymbol{\theta}} W_{\varepsilon}^{2}\left(\rho_i, \rho_i^{\text{epoch index}}(\boldsymbol{\theta})\right) \quad \forall i \in \{0, T\}$$

Case Study: Synthesize BCC Crystalline Structure by PDF Steering in $(\langle C_{10}\rangle, \langle C_{12}\rangle)$ Space

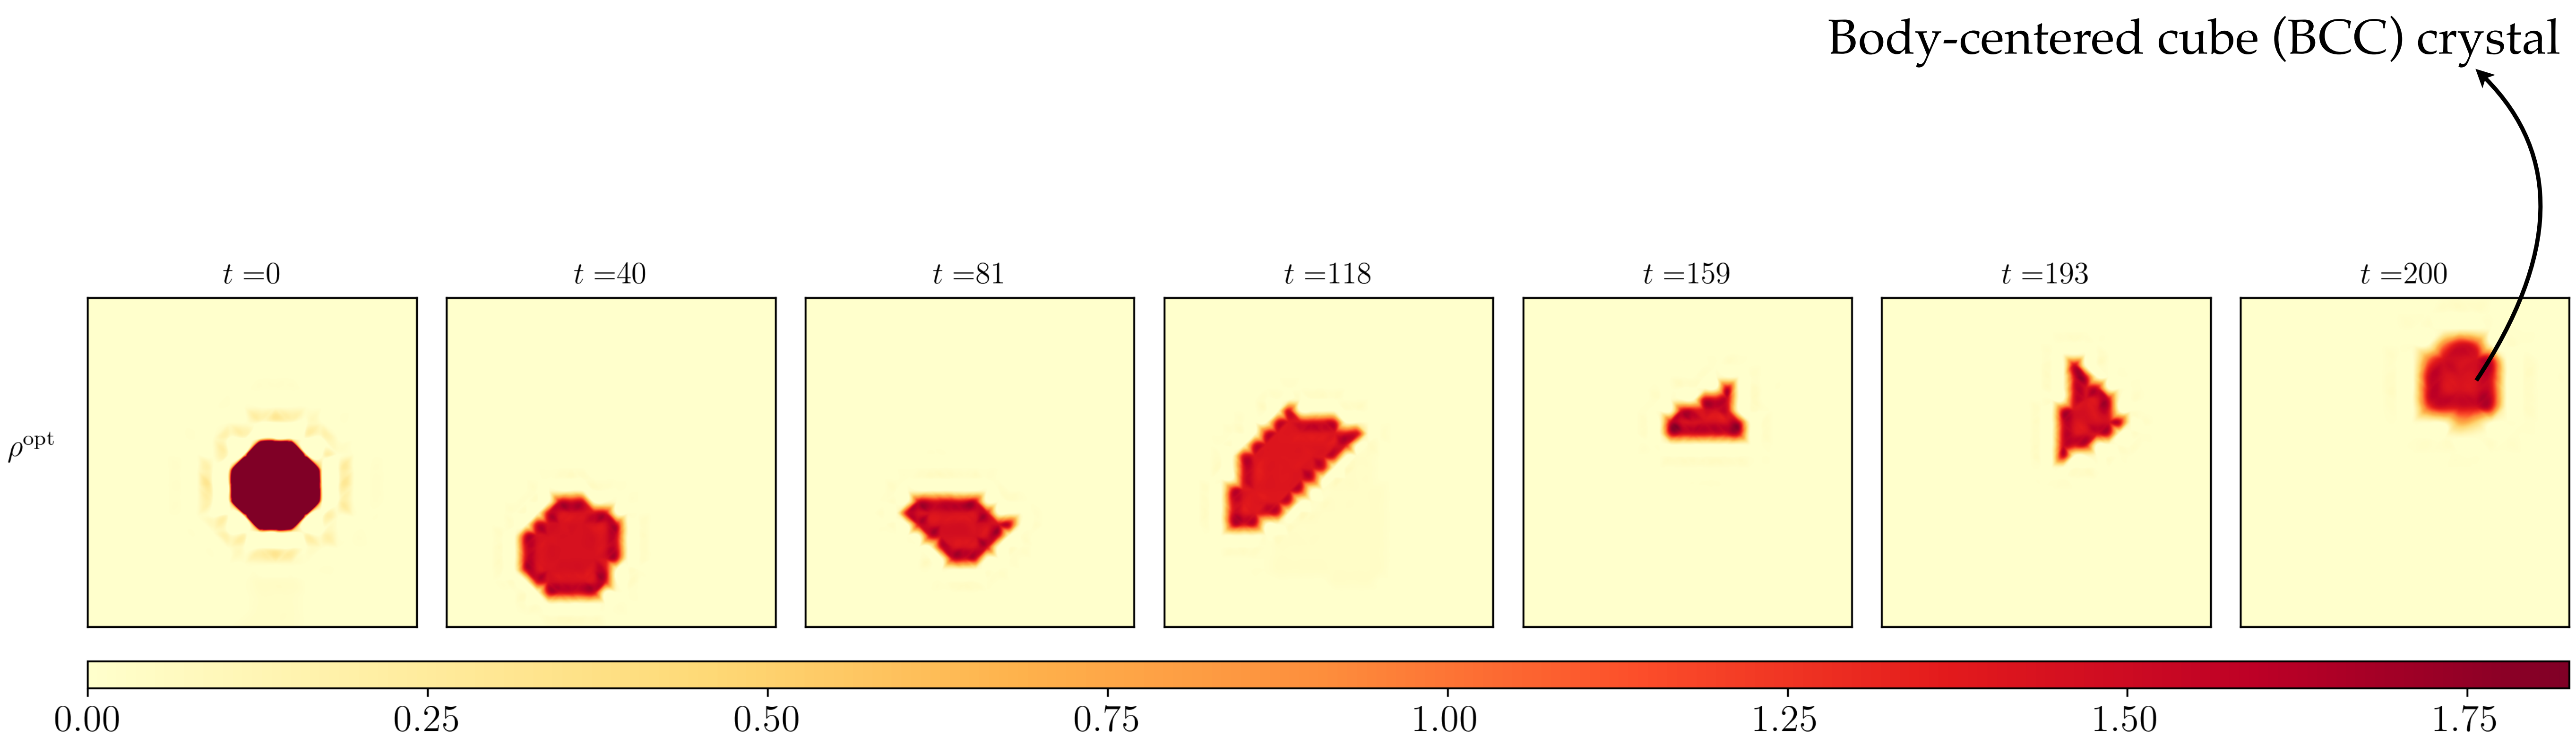

Data-driven:

Uses PINN with Sinkhorn losses + the drift-diffusion are themselves NNs

Outlook

- Lots of interesting theory, algorithms and applications to be done

- Excellent intersections with related communities: learning, information theory, robotics, systems biology, smart manufacturing

PhD and Postdoc openings on Optimal Transport, Schrödinger bridge, and Stochastic Control/ML

Thank You

Support: