Finite Horizon Optimal Density Regulation for Nonlinear Systems

Abhishek Halder

Department of Applied Mathematics
University of California, Santa Cruz
Santa Cruz, CA 95064

Joint work with S. Haddad and K.F. Caluya (UC Santa Cruz),
B. Singh (Ford Greenfield Labs)

ACC 2021 Workshop: Control of Distributions

Density Regulation via State Feedback

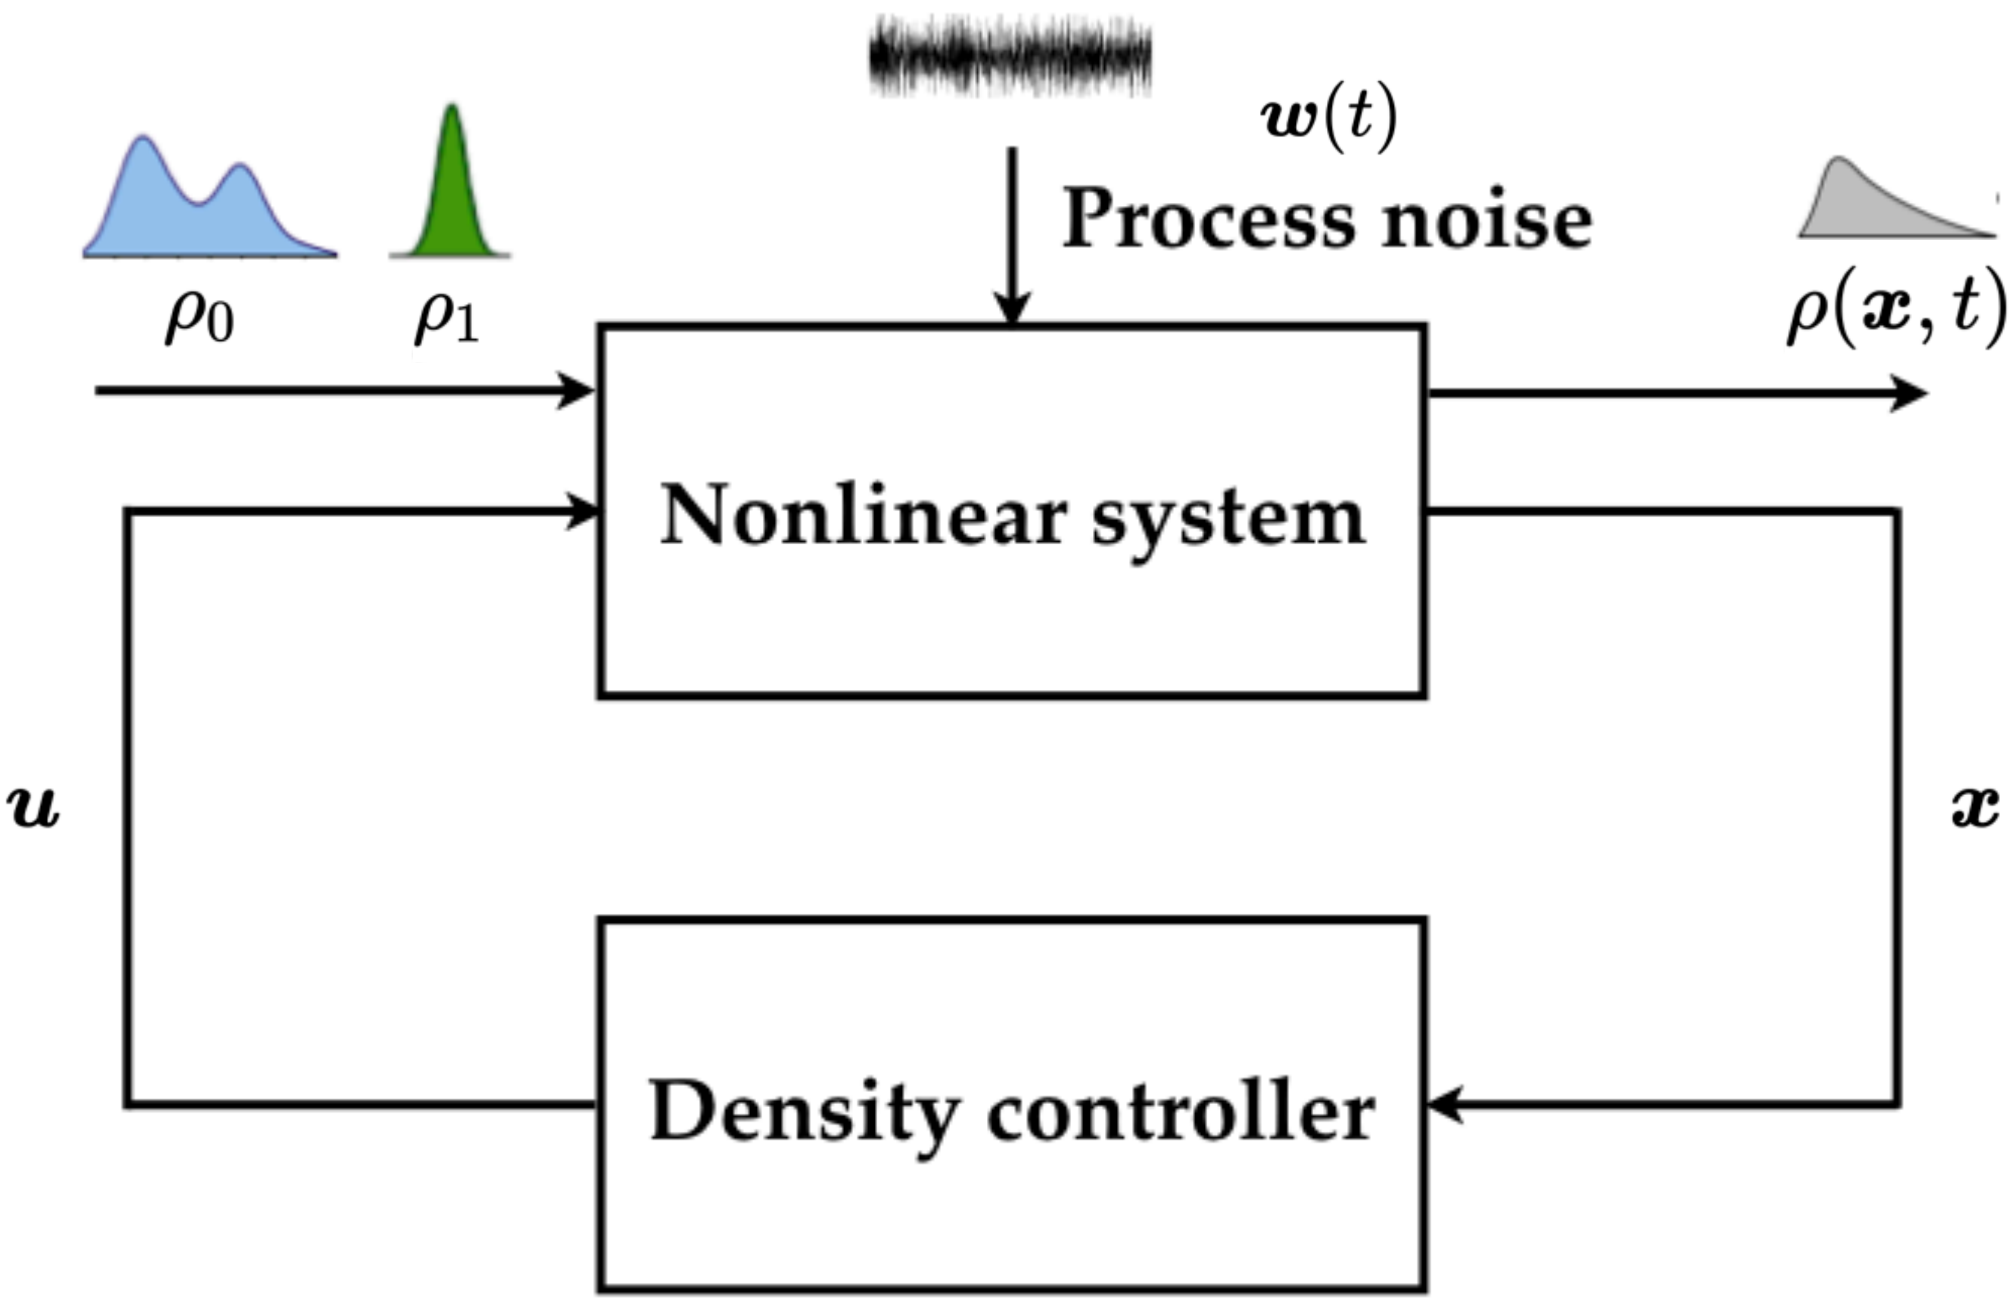

Motivating Applications

Distribution $\sim$ Probability

Distribution $\sim$ Population

Motivating Applications

Distribution ~ Probability

Distribution ~ Population

Spacecraft landing with desired statistical accuracy

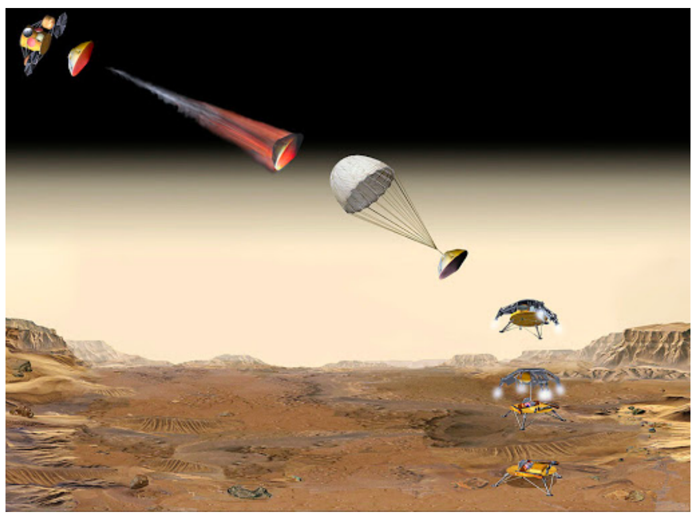

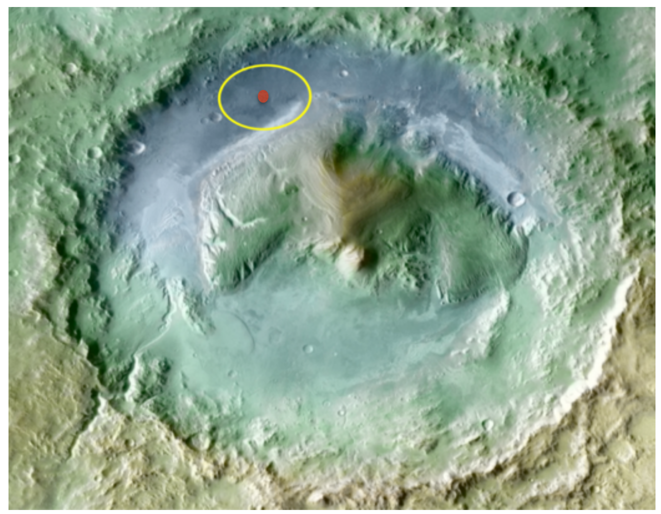

Gale Crater (4.49S, 137.42E)

Risk management for automated driving in multi-lane highways

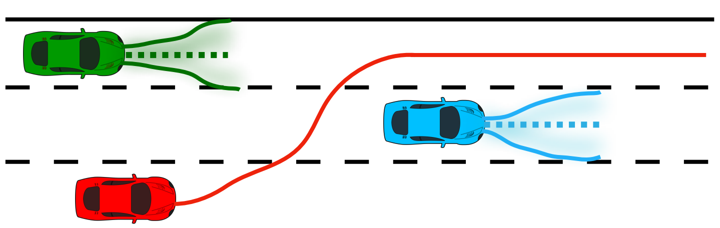

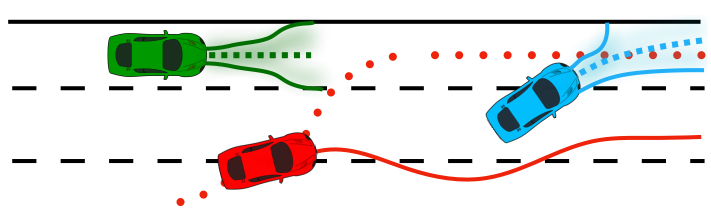

Control of uncertainties

Motivating Applications

Distribution ~ Probability

Spacecraft landing with desired statistical accuracy

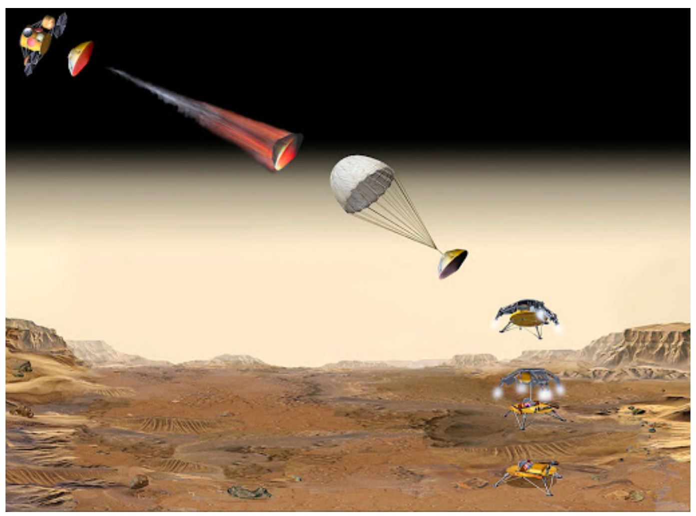

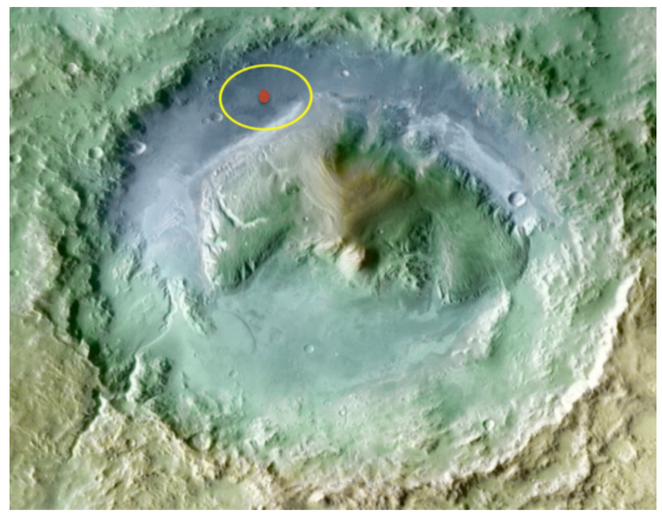

Gale Crater (4.49S, 137.42E)

Risk management for automated driving in multi-lane highways

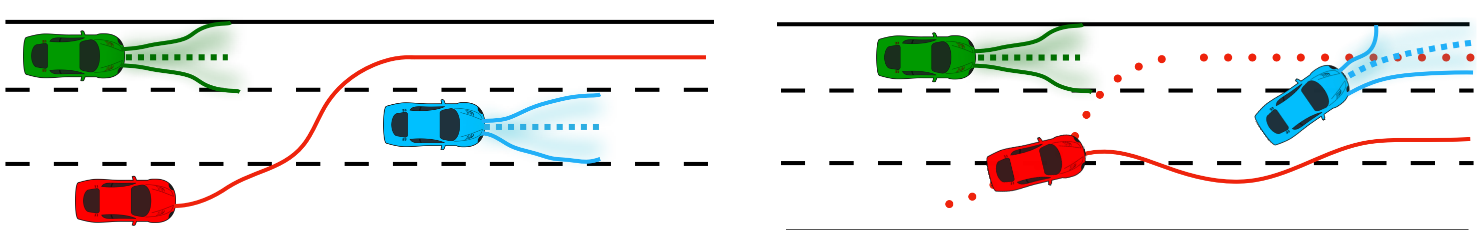

Control of uncertainties

Distribution ~ Population

Dynamic shaping of swarms

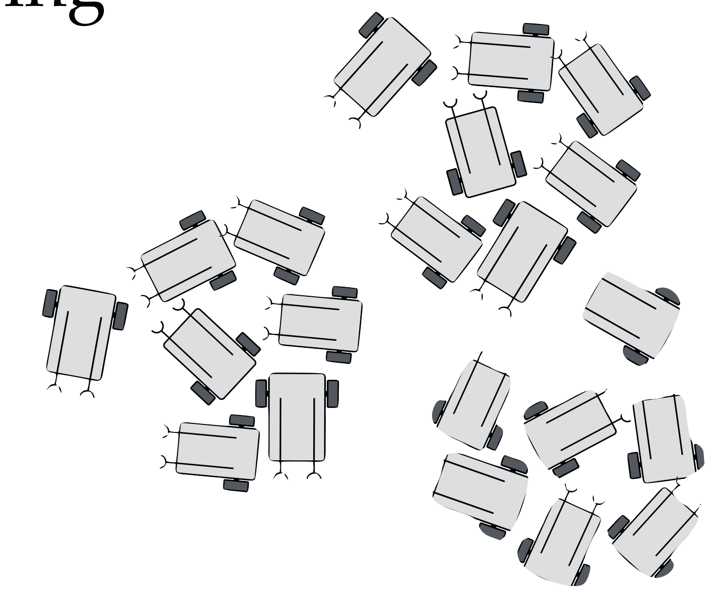

Feedback sync. and desync. of neuronal population

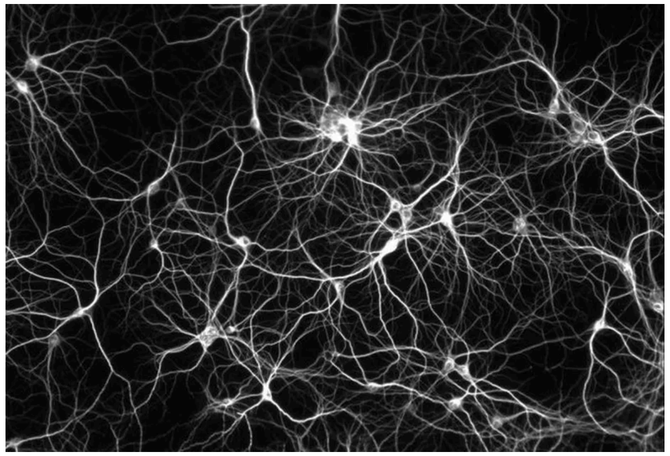

Control of ensemble

State Feedback Density Steering

Steer joint state PDF via feedback control over finite time horizon

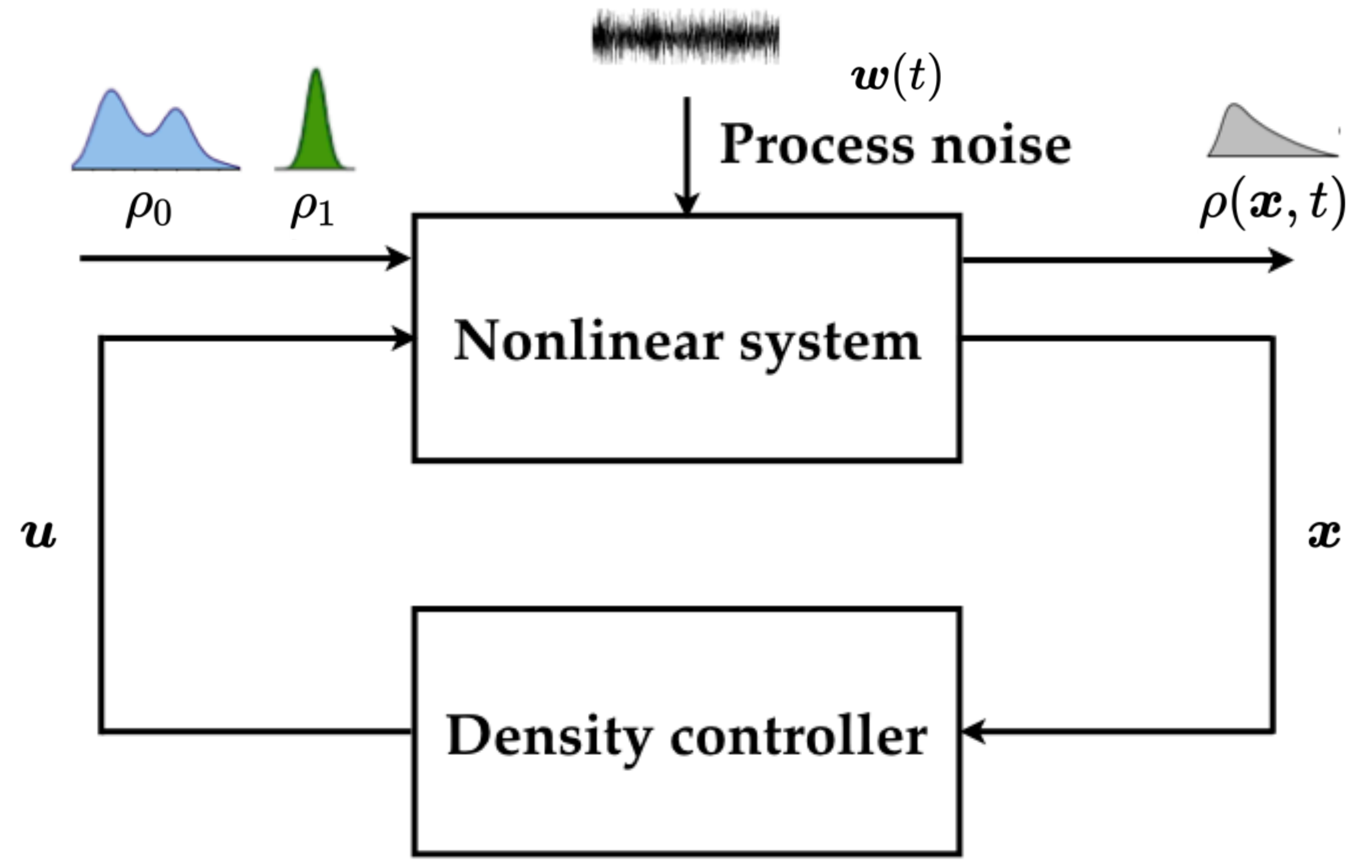

Common scenario: $\boldsymbol{G} \equiv \boldsymbol{B}$

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

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

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

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

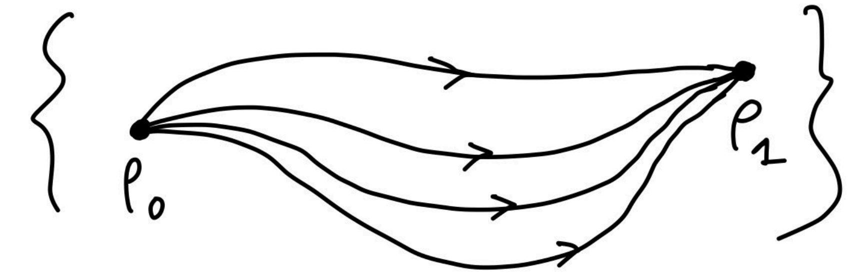

Hessian operator w.r.t. state: Hess

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

subject to

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

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

Optimal Control Problem over PDFs

Existence-uniqueness needs regularity assumptions

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

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

Hamilton-Jacobi-Bellman-like PDE

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

Boundary conditions:

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

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

Feedback Synthesis via the Schrödinger System

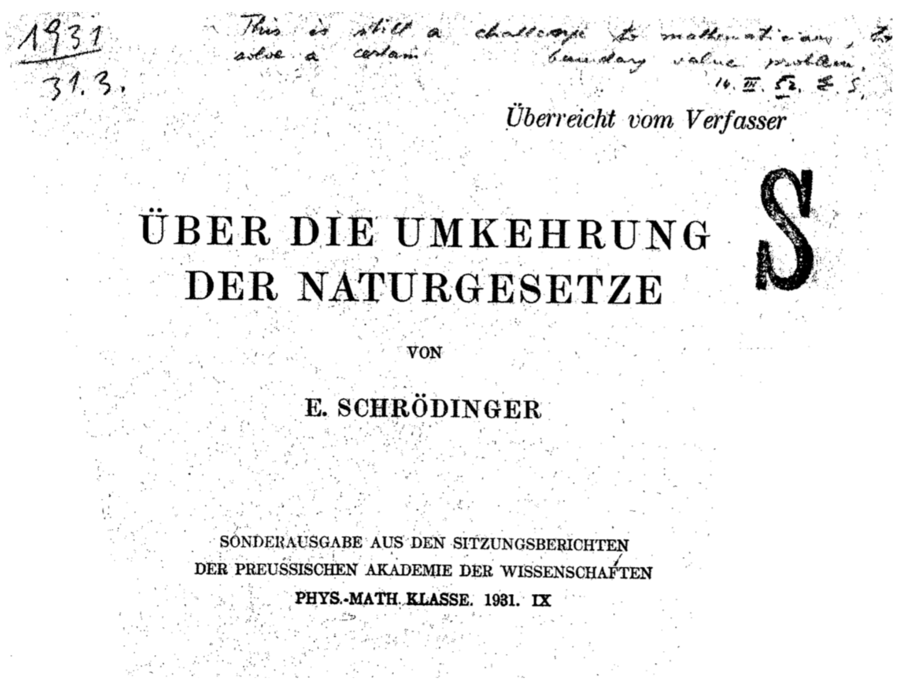

1931
31.3.

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

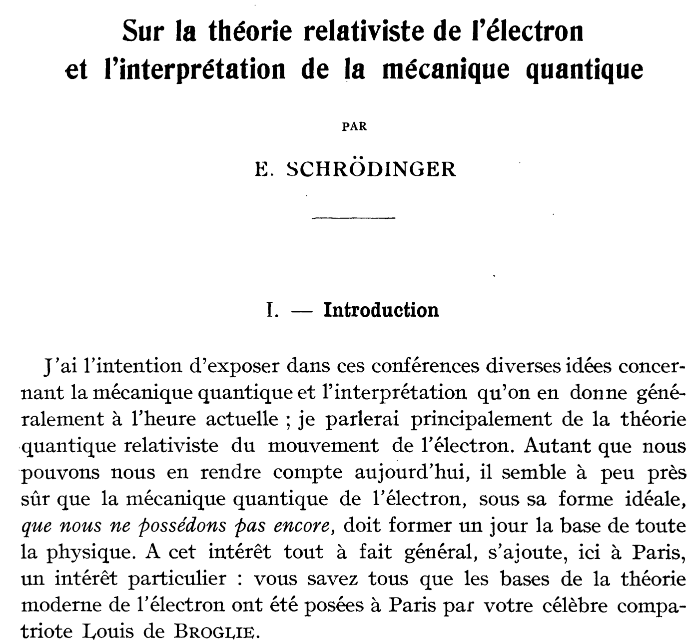

Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique

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:

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

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

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

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

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

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

This recursion is contractive in the Hilbert metric!!

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

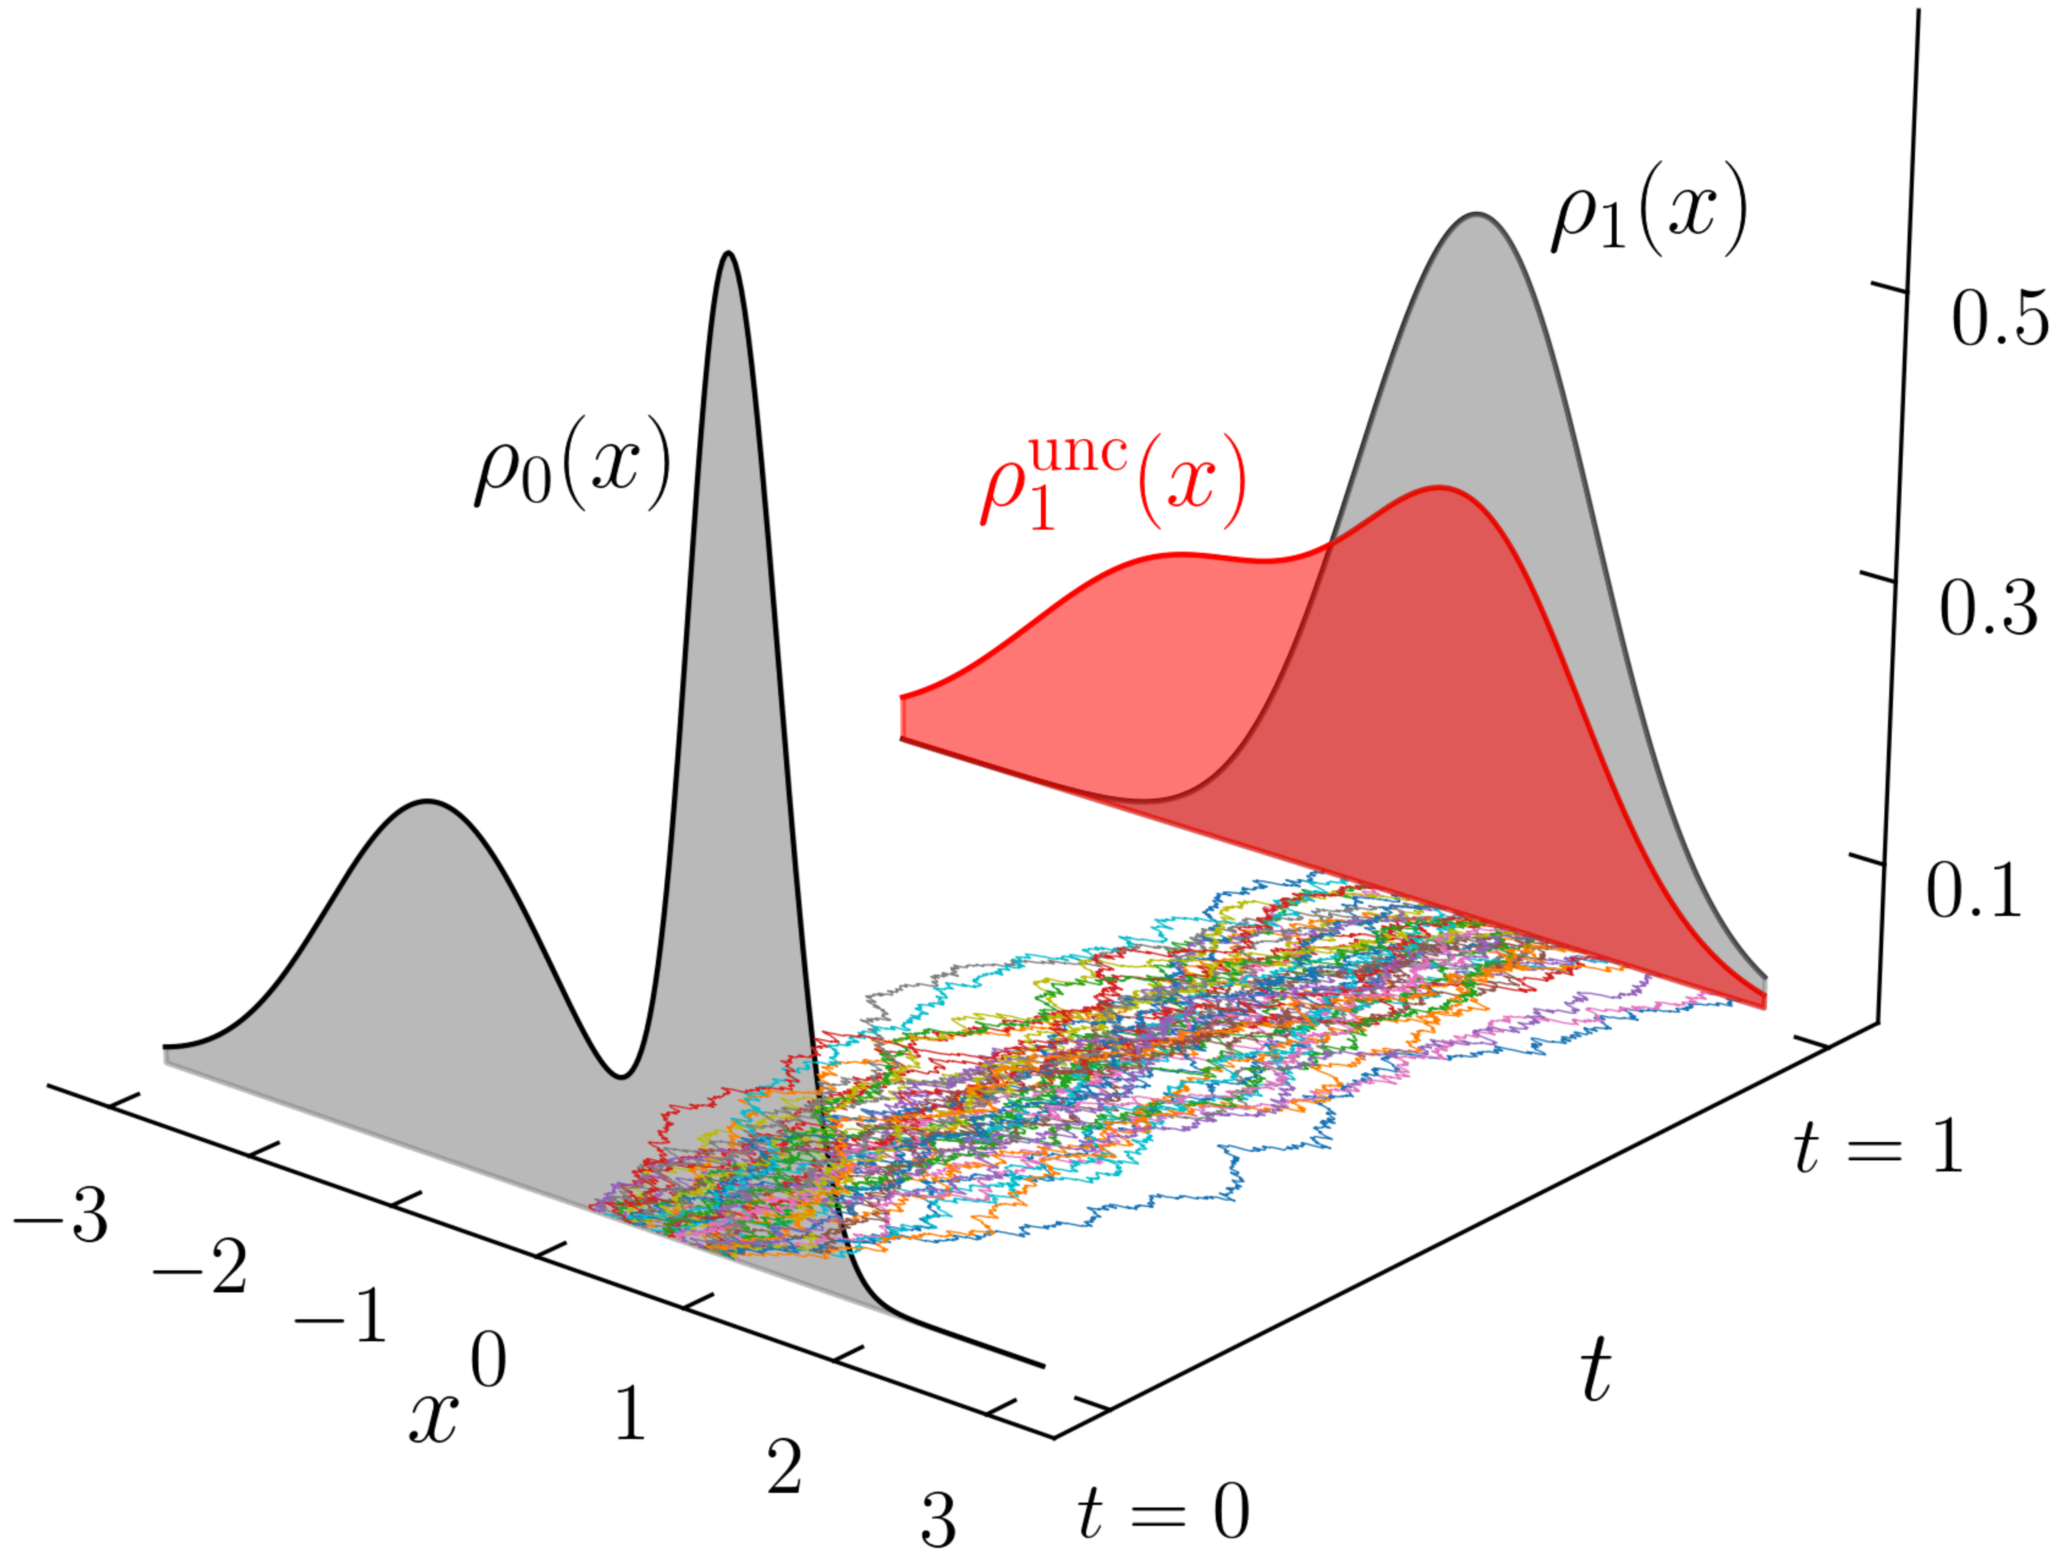

Zero prior dynamics

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

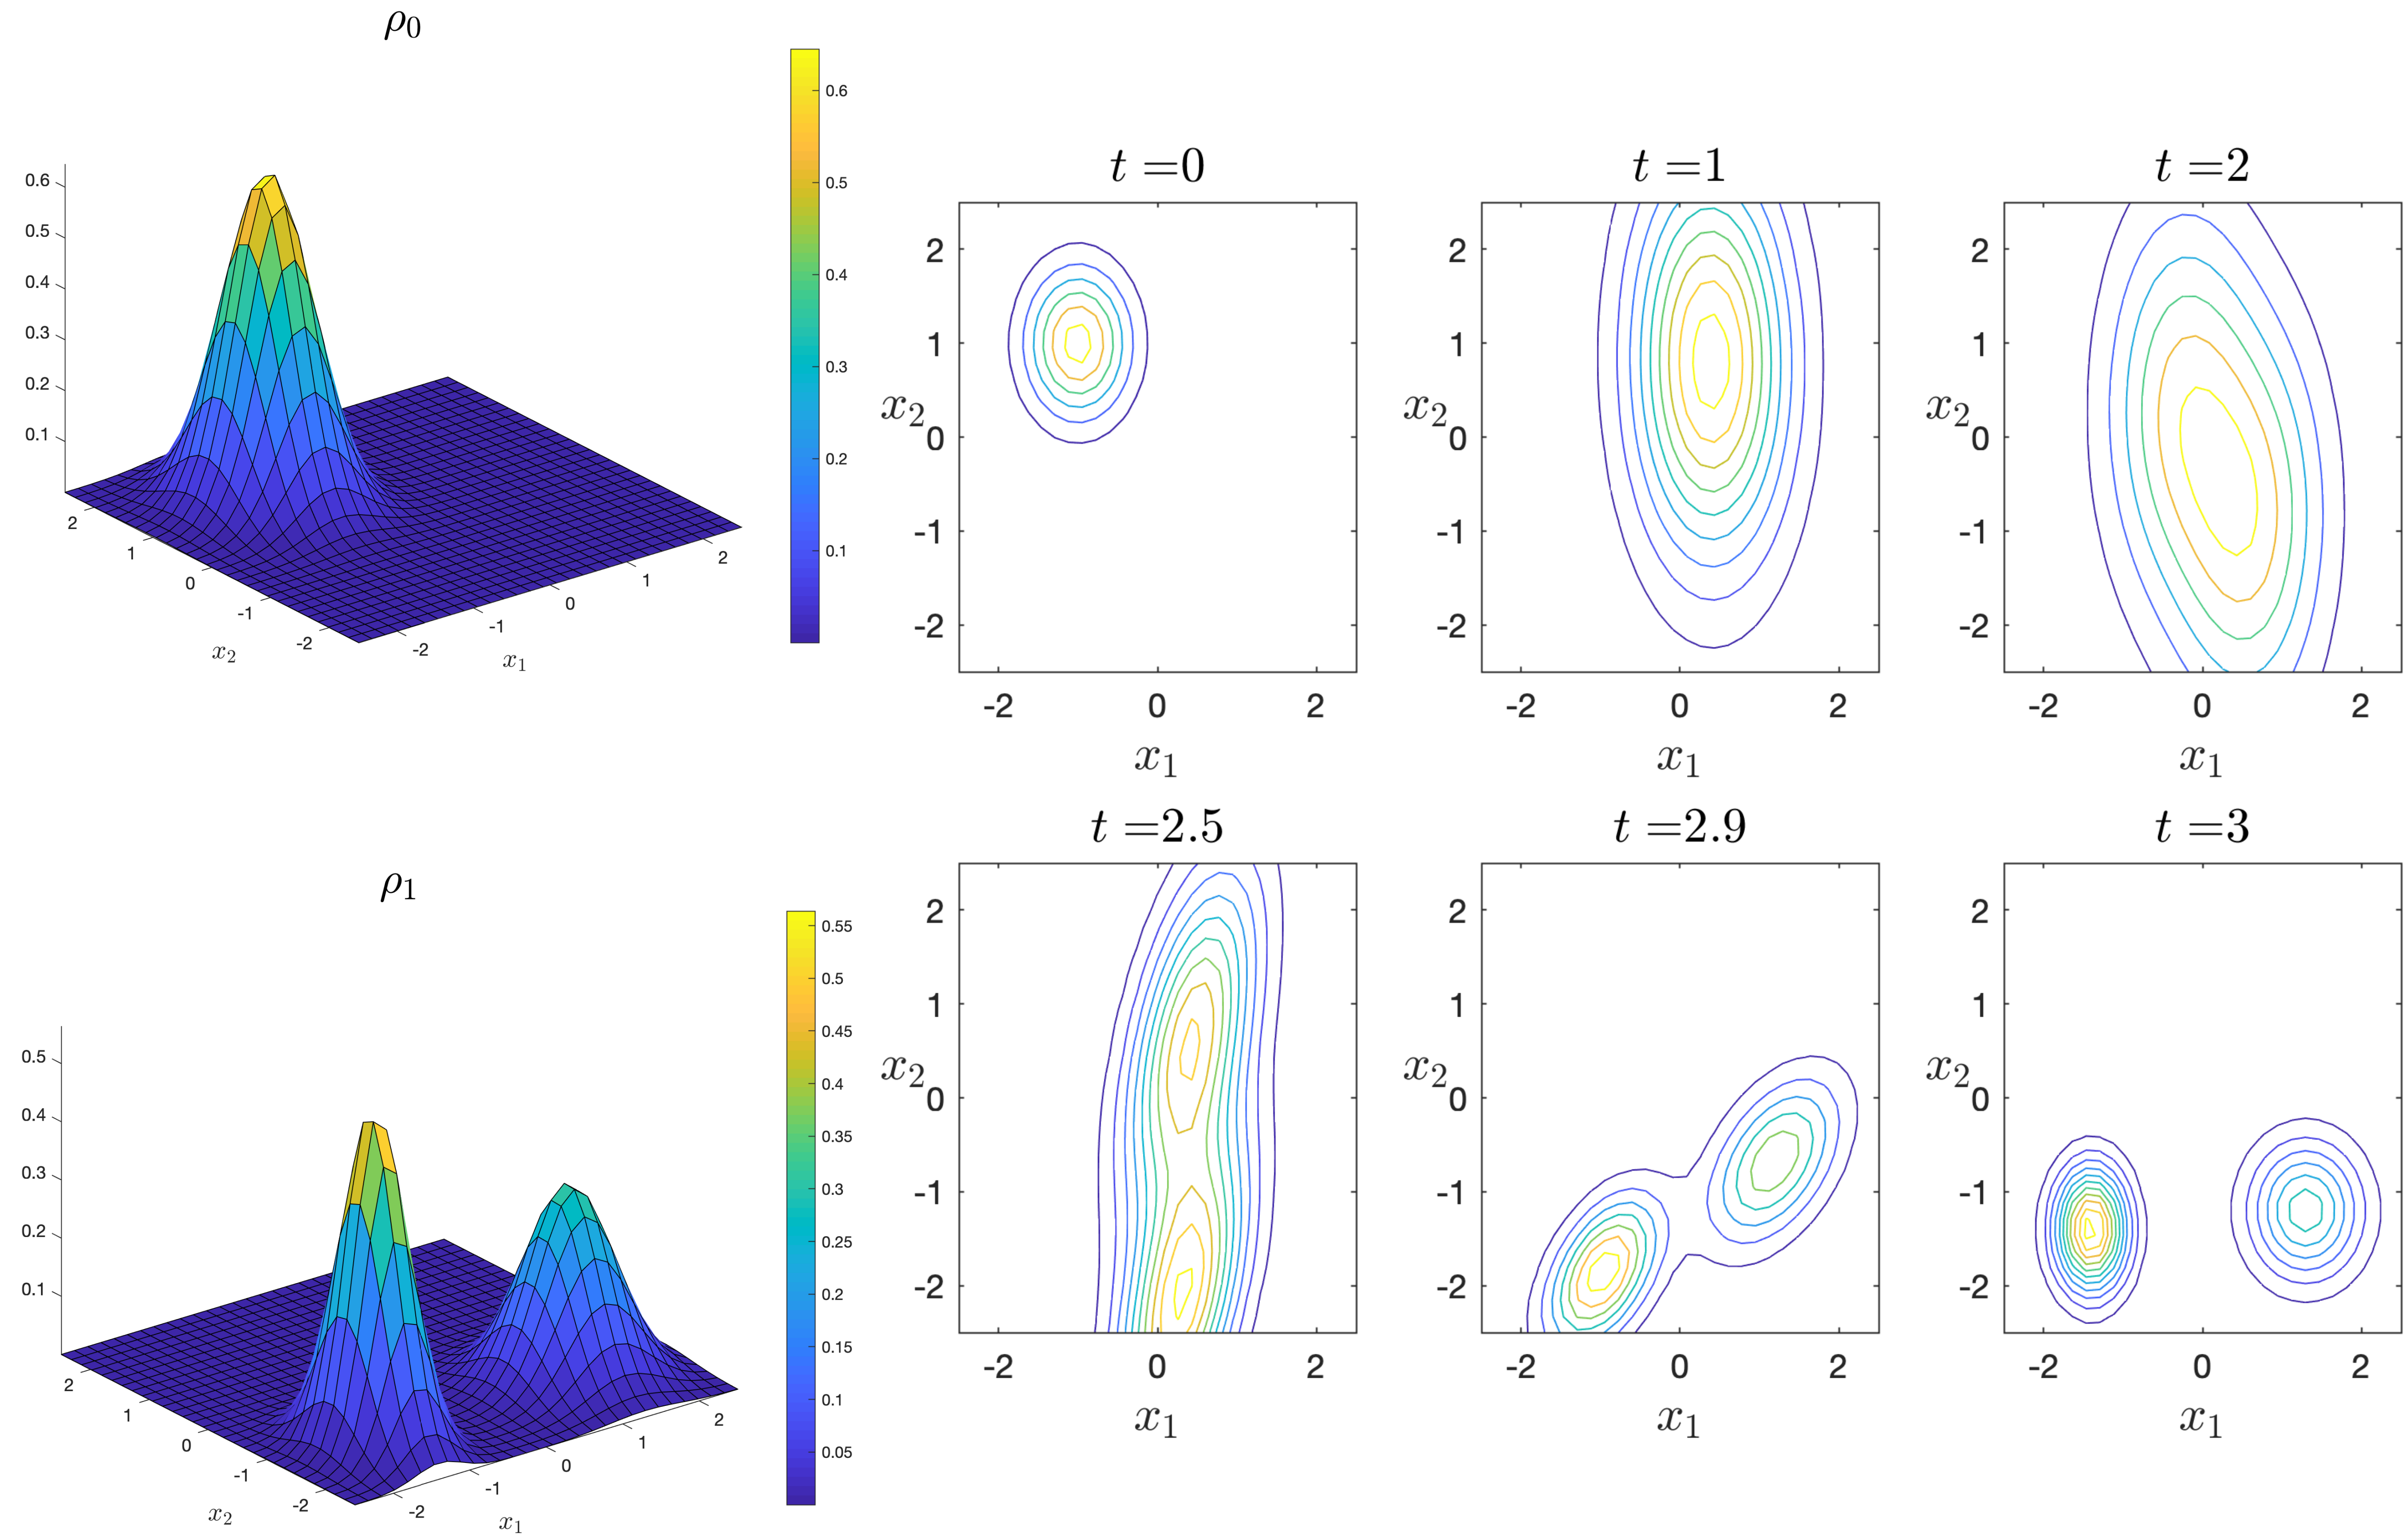

Linear prior dynamics

In general …

Need (uncontrolled) forward AND backward Kolmogorov solvers

Bad news: need two different solvers

Good news: sometimes one solver suffices!

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:

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

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

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

Brief Detour: Generalized Gradient Flow

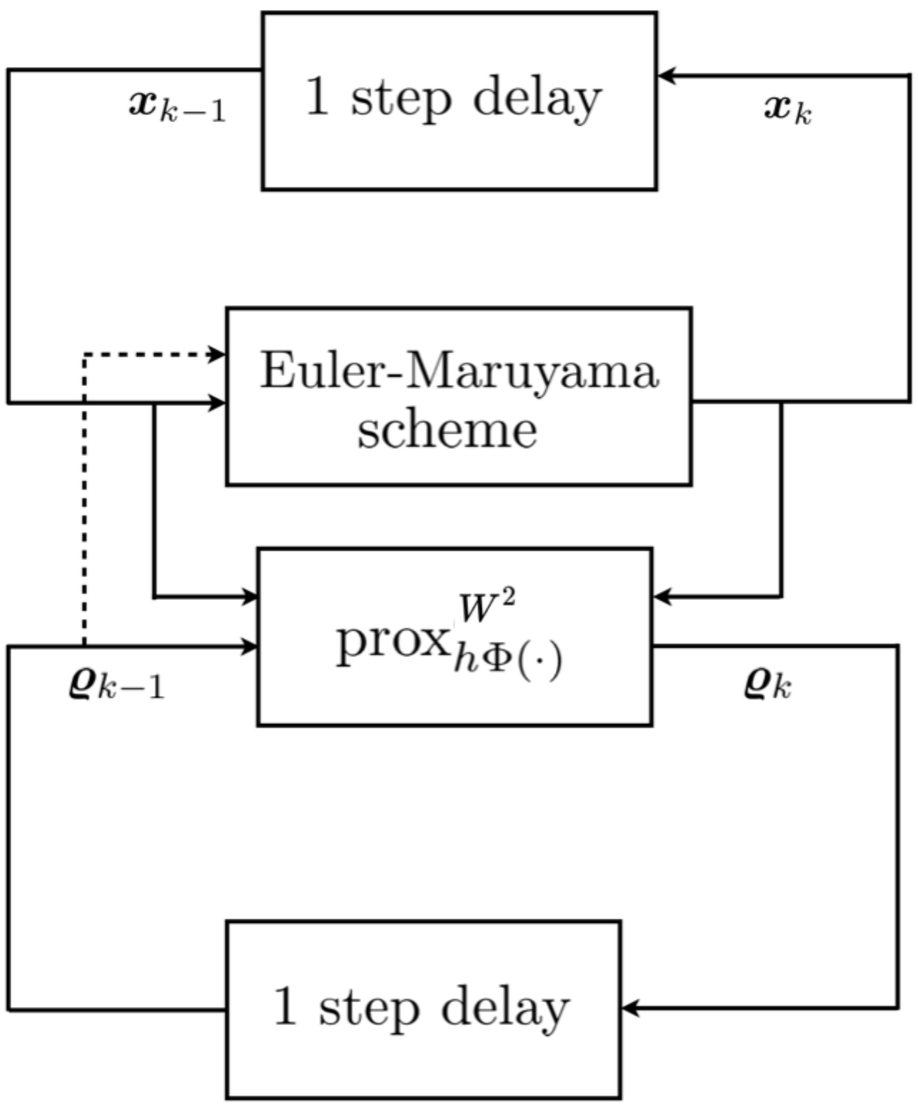

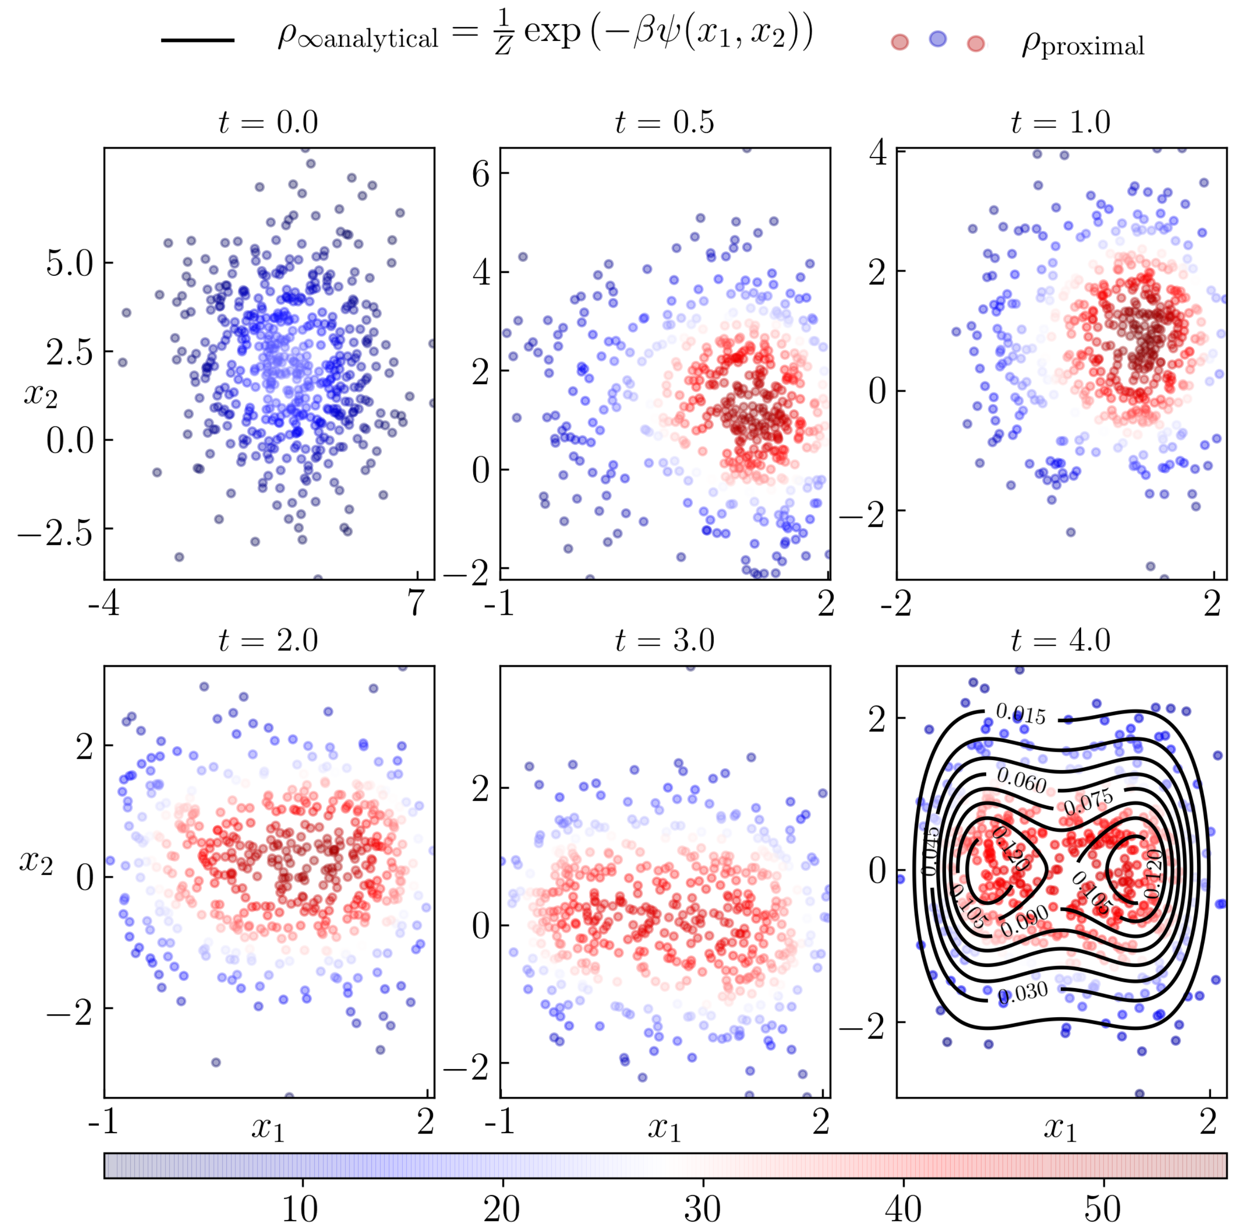

Details:

— 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: $\boldsymbol{\xi}, \boldsymbol{\eta} \in \mathbb{R}^m, \boldsymbol{x} := (\boldsymbol{\xi}, \boldsymbol{\eta})^\top \in \mathbb{R}^n$, $n = 2m$, $V \in C^2(\mathbb{R}^m)$, $\inf V > -\infty$, 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 TAC 2021*.

Feedback Density Control: Gradient Drift

Uncontrolled joint PDF evolution:

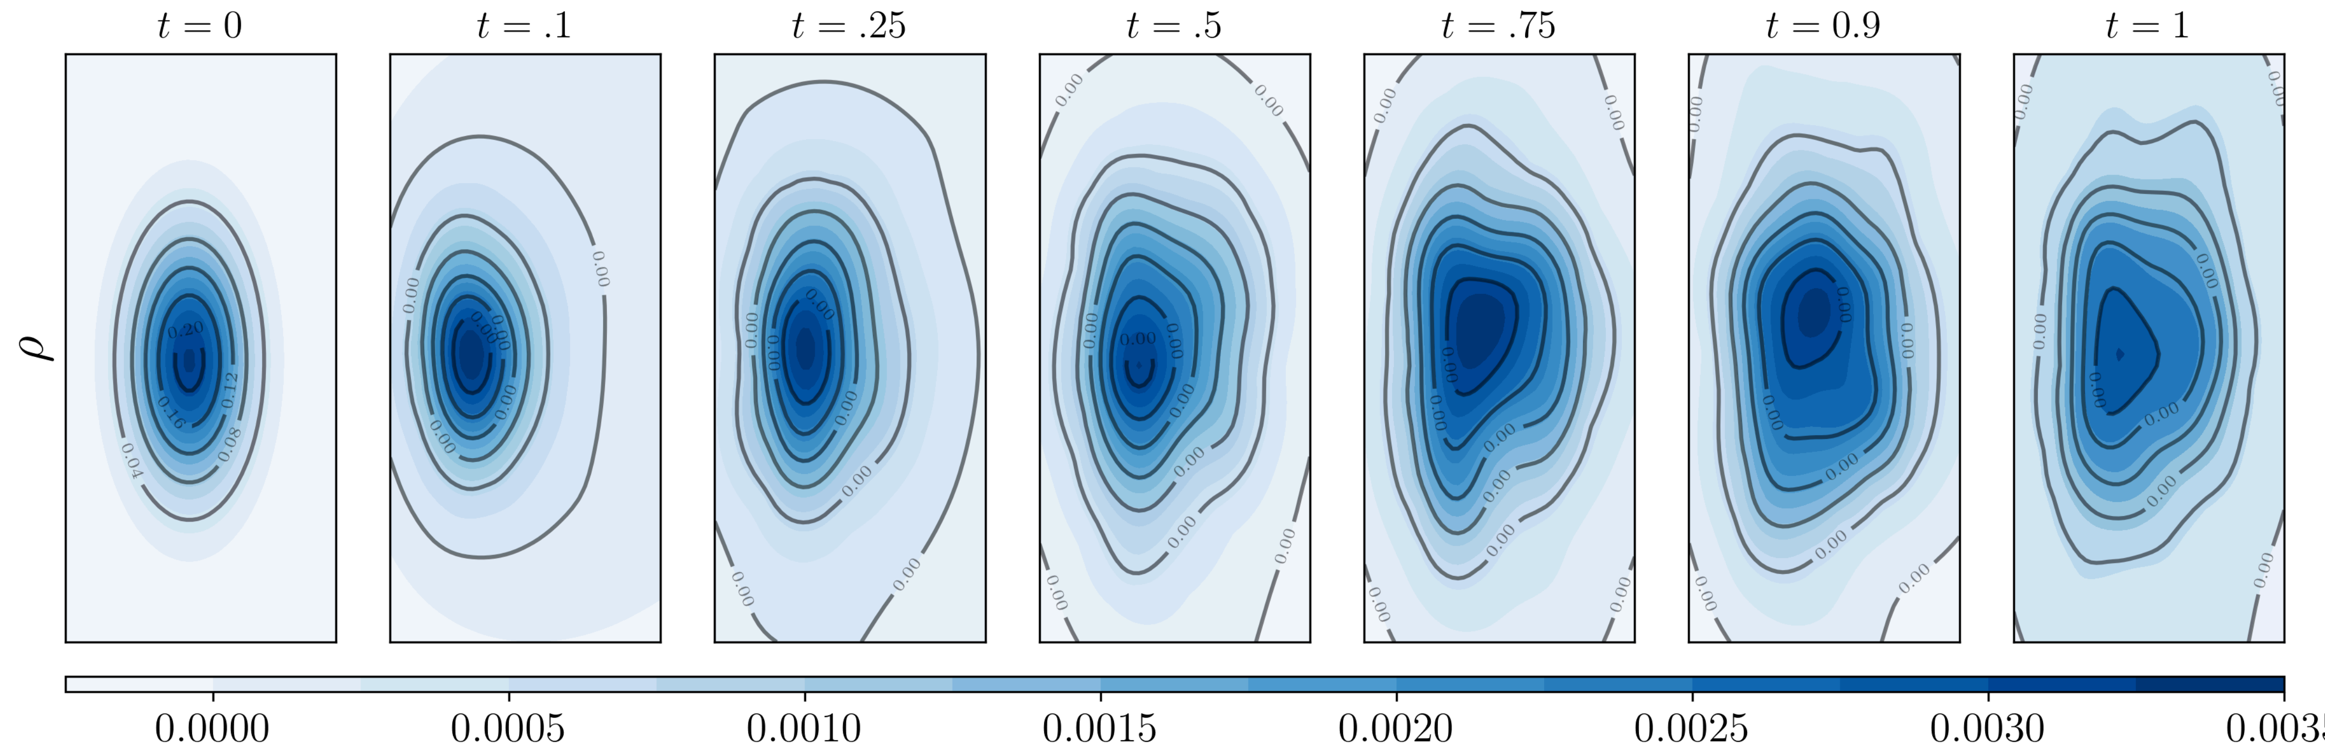

Optimal controlled joint PDF evolution:

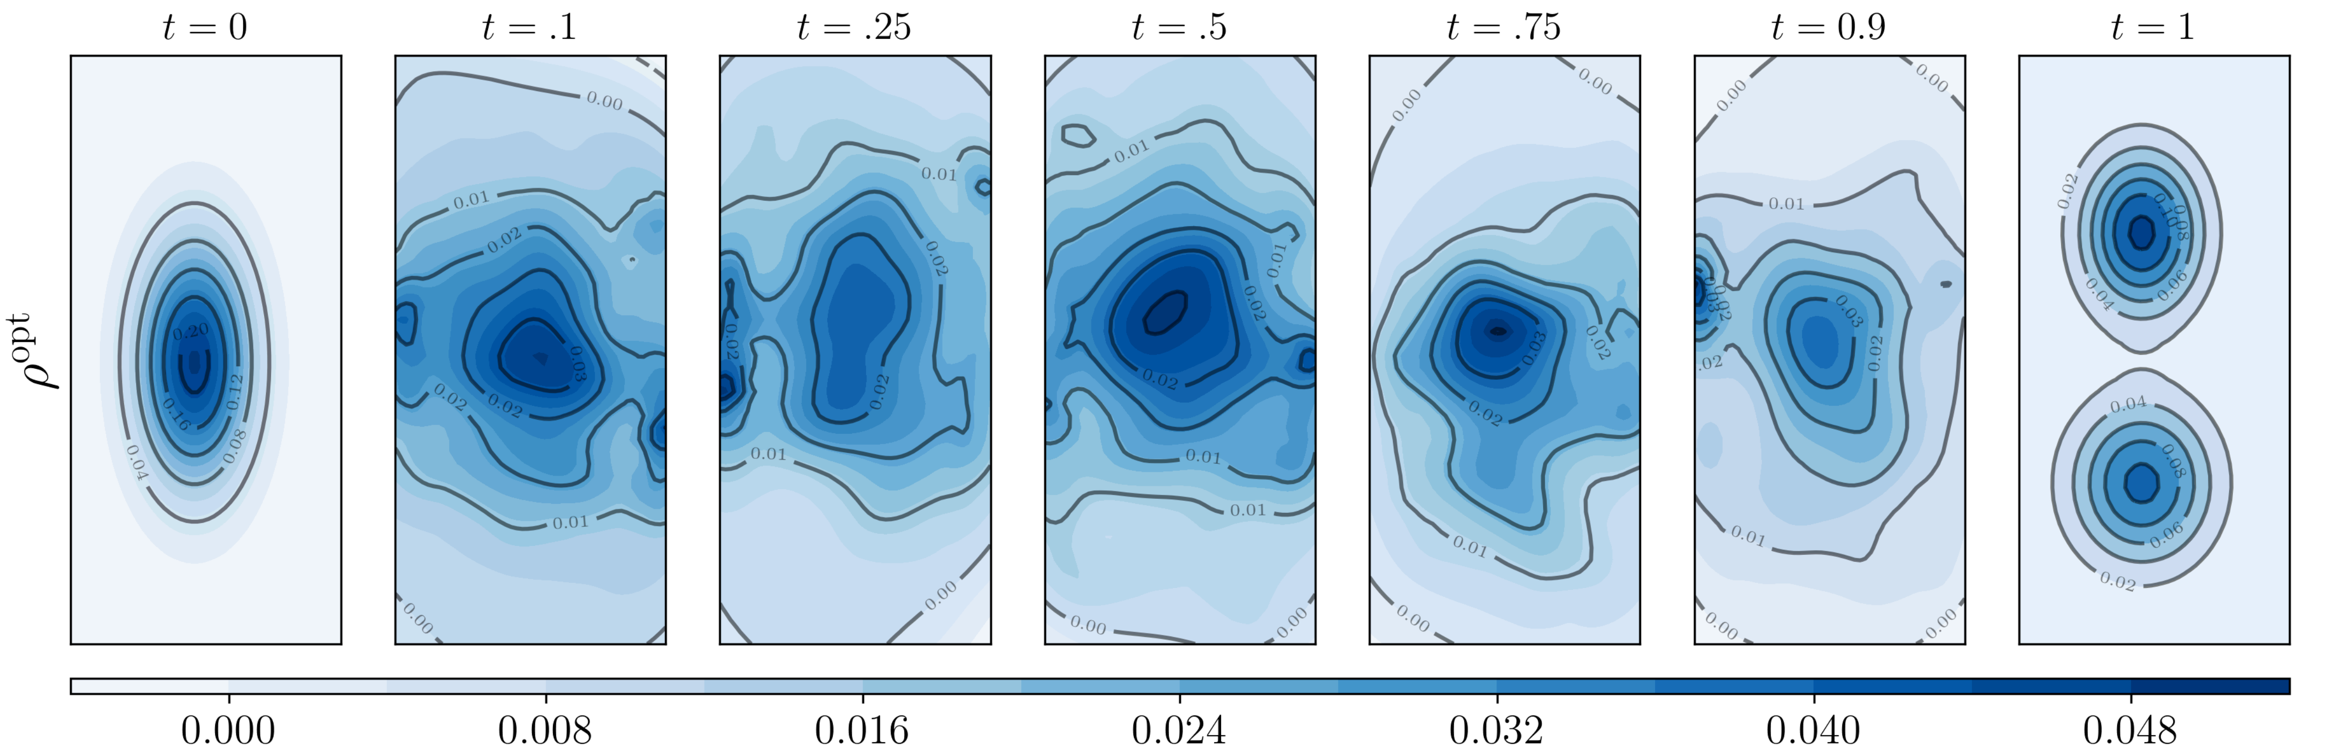

Feedback Density Control: Mixed Conservative-Dissipative Drift

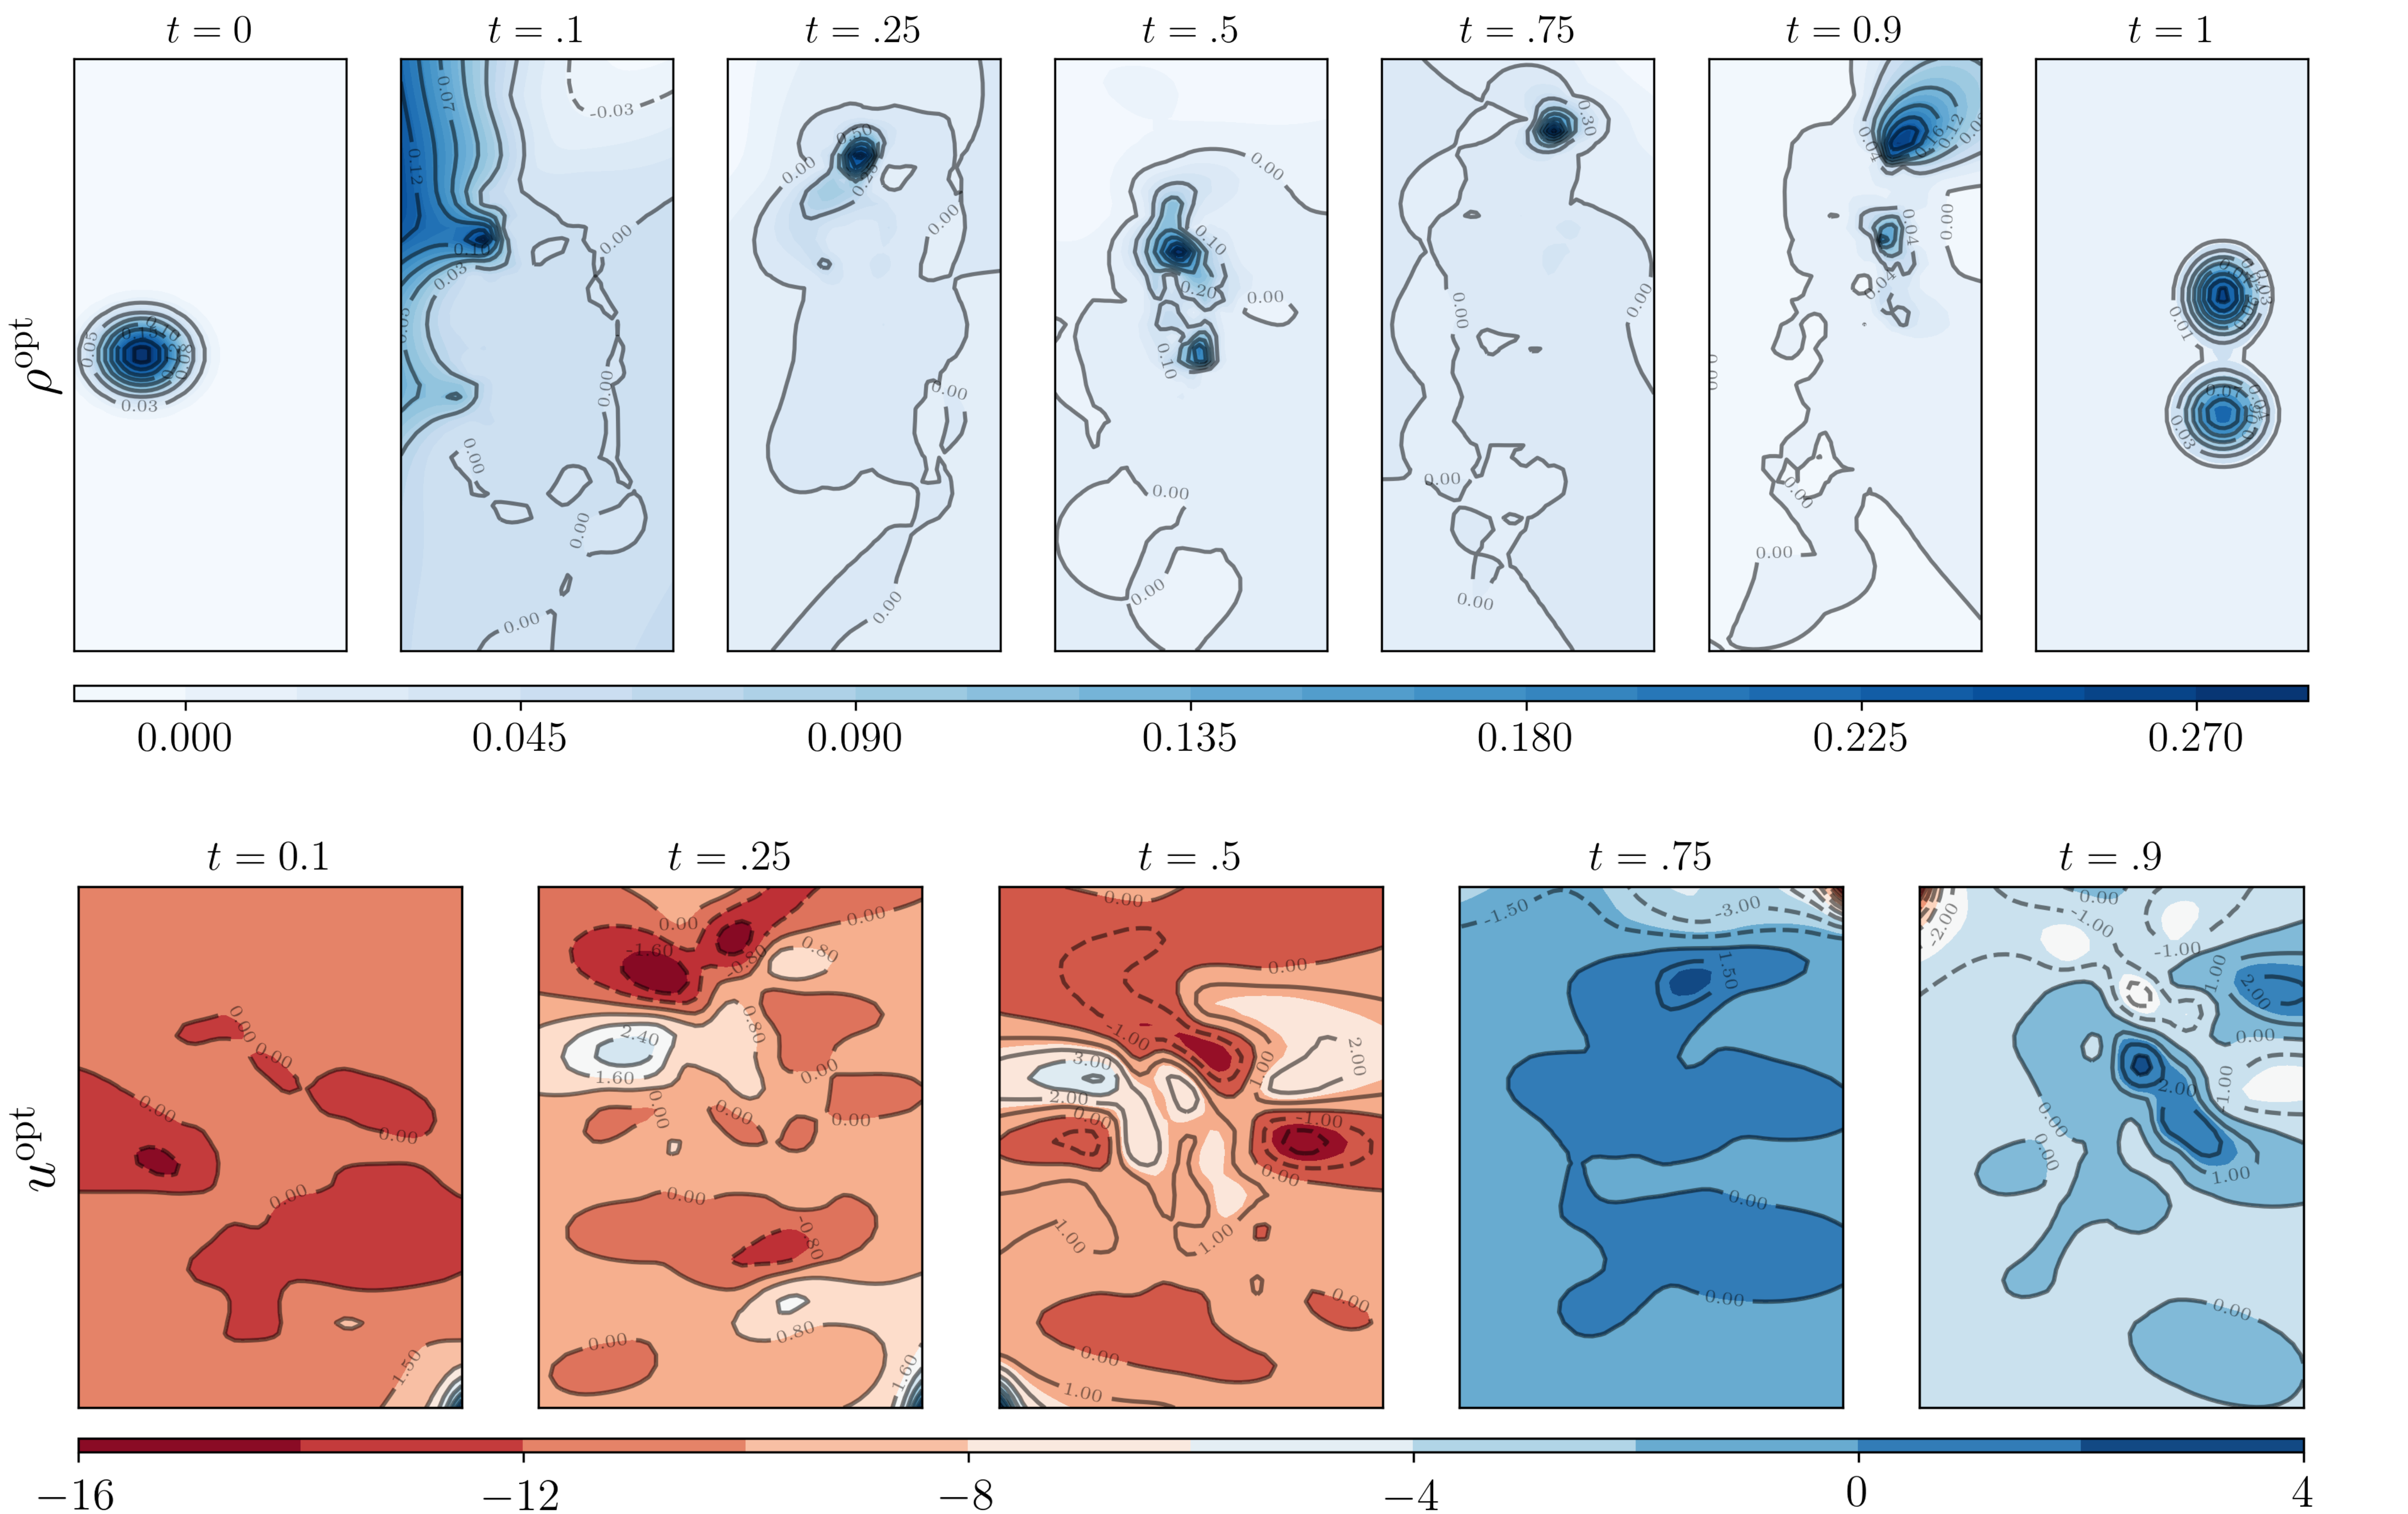

Nonlinear Density Steering with Deterministic Path Constraints

Optimal controlled state PDFs:

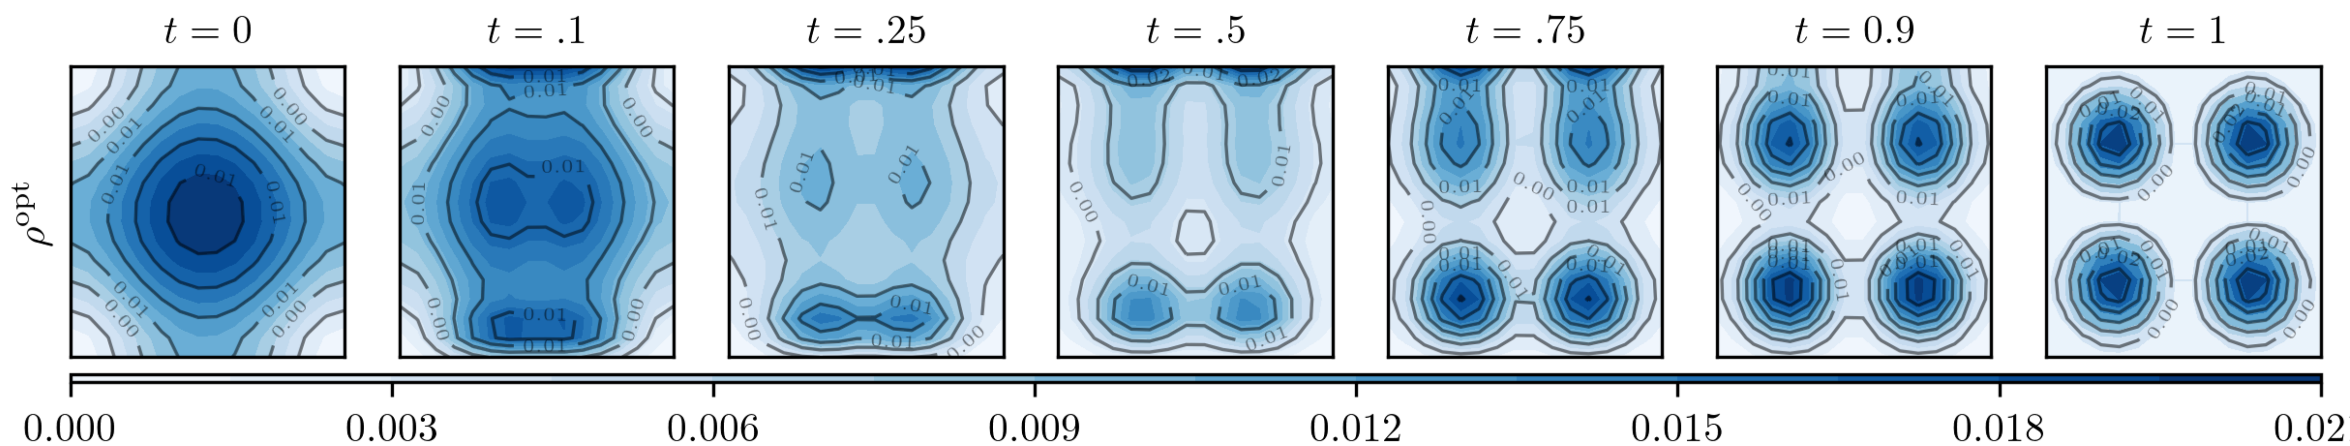

Uncontrolled state PDFs:

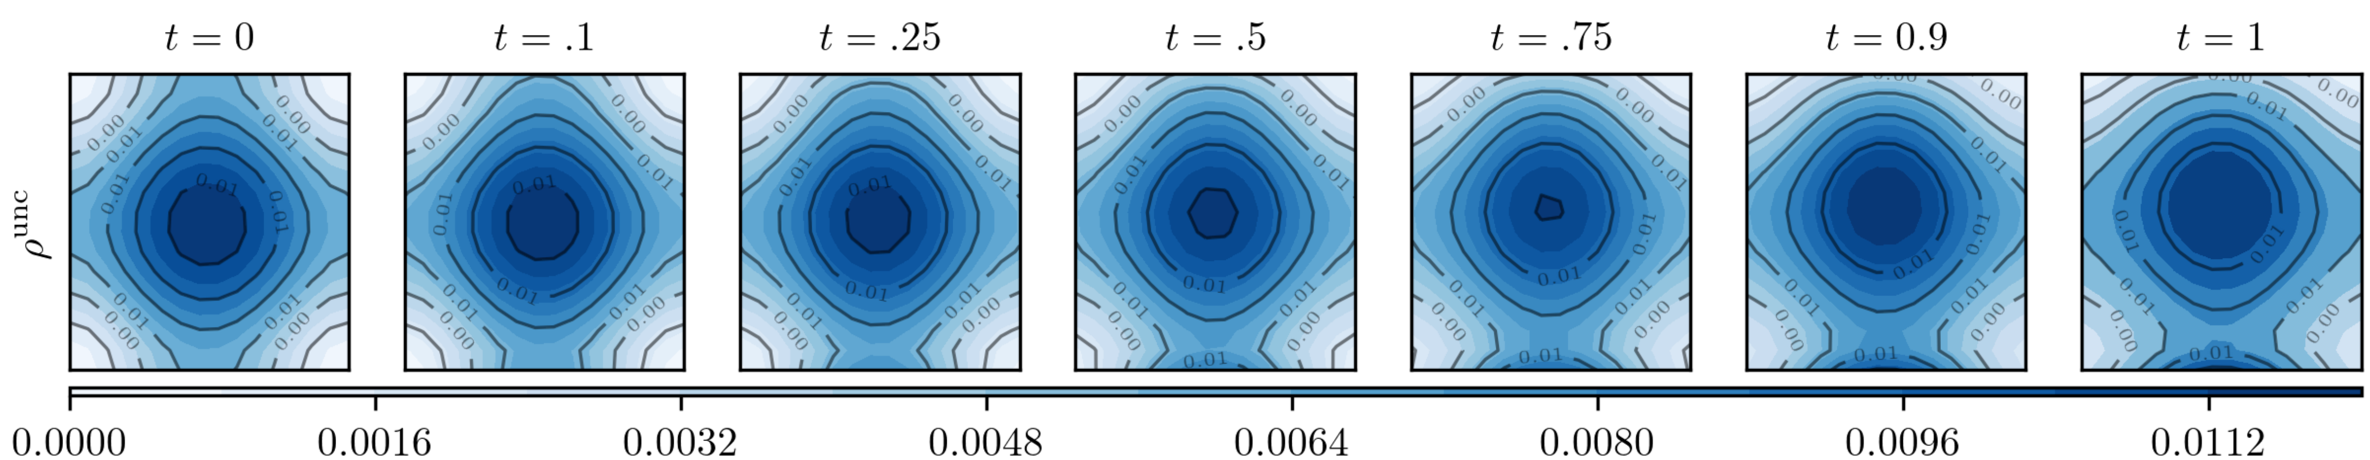

This ACC: Session TuA22

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

Application: Multi-lane Automated Driving

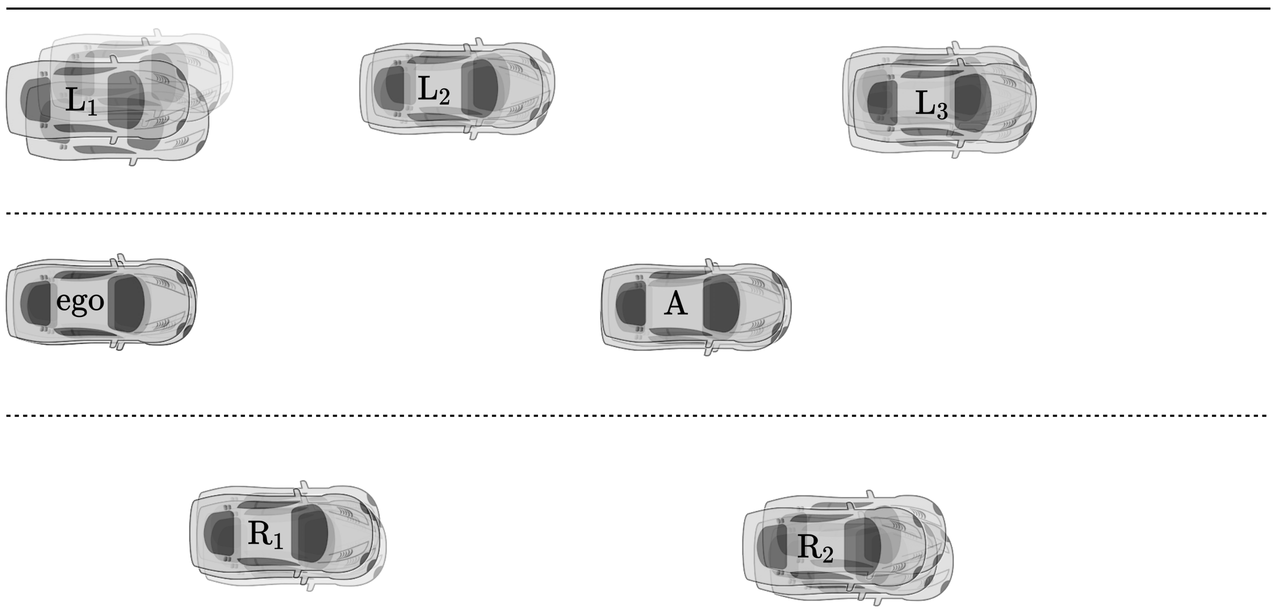

t_0

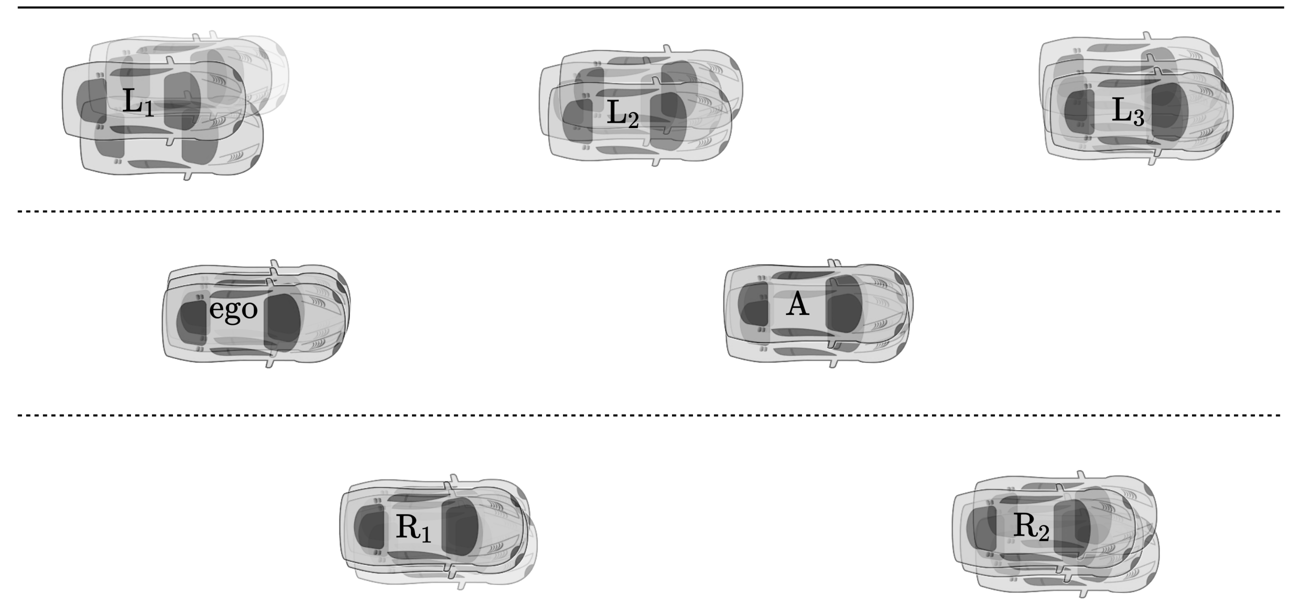

t_1

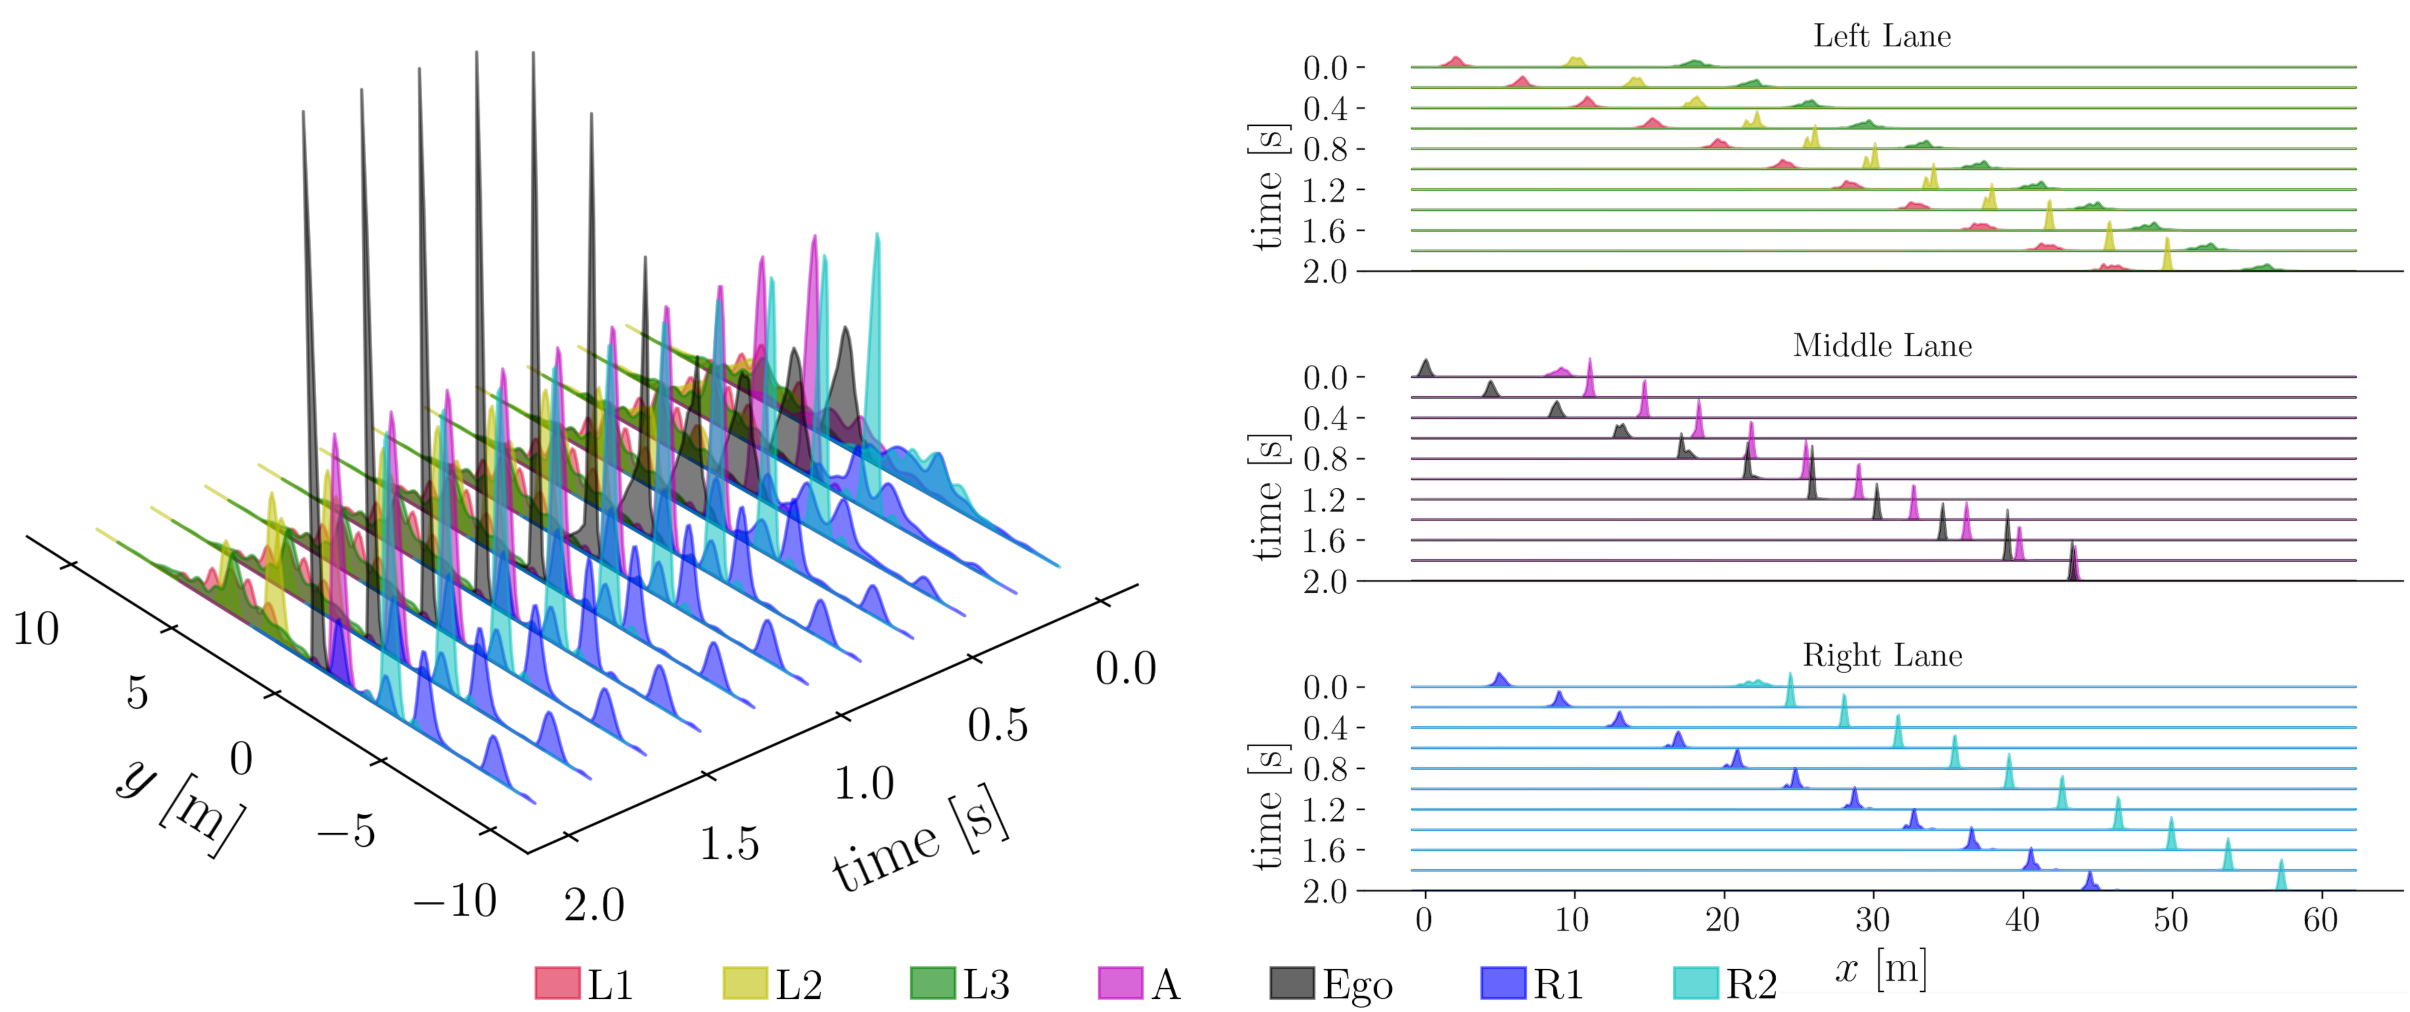

Application: Multi-lane Automated Driving

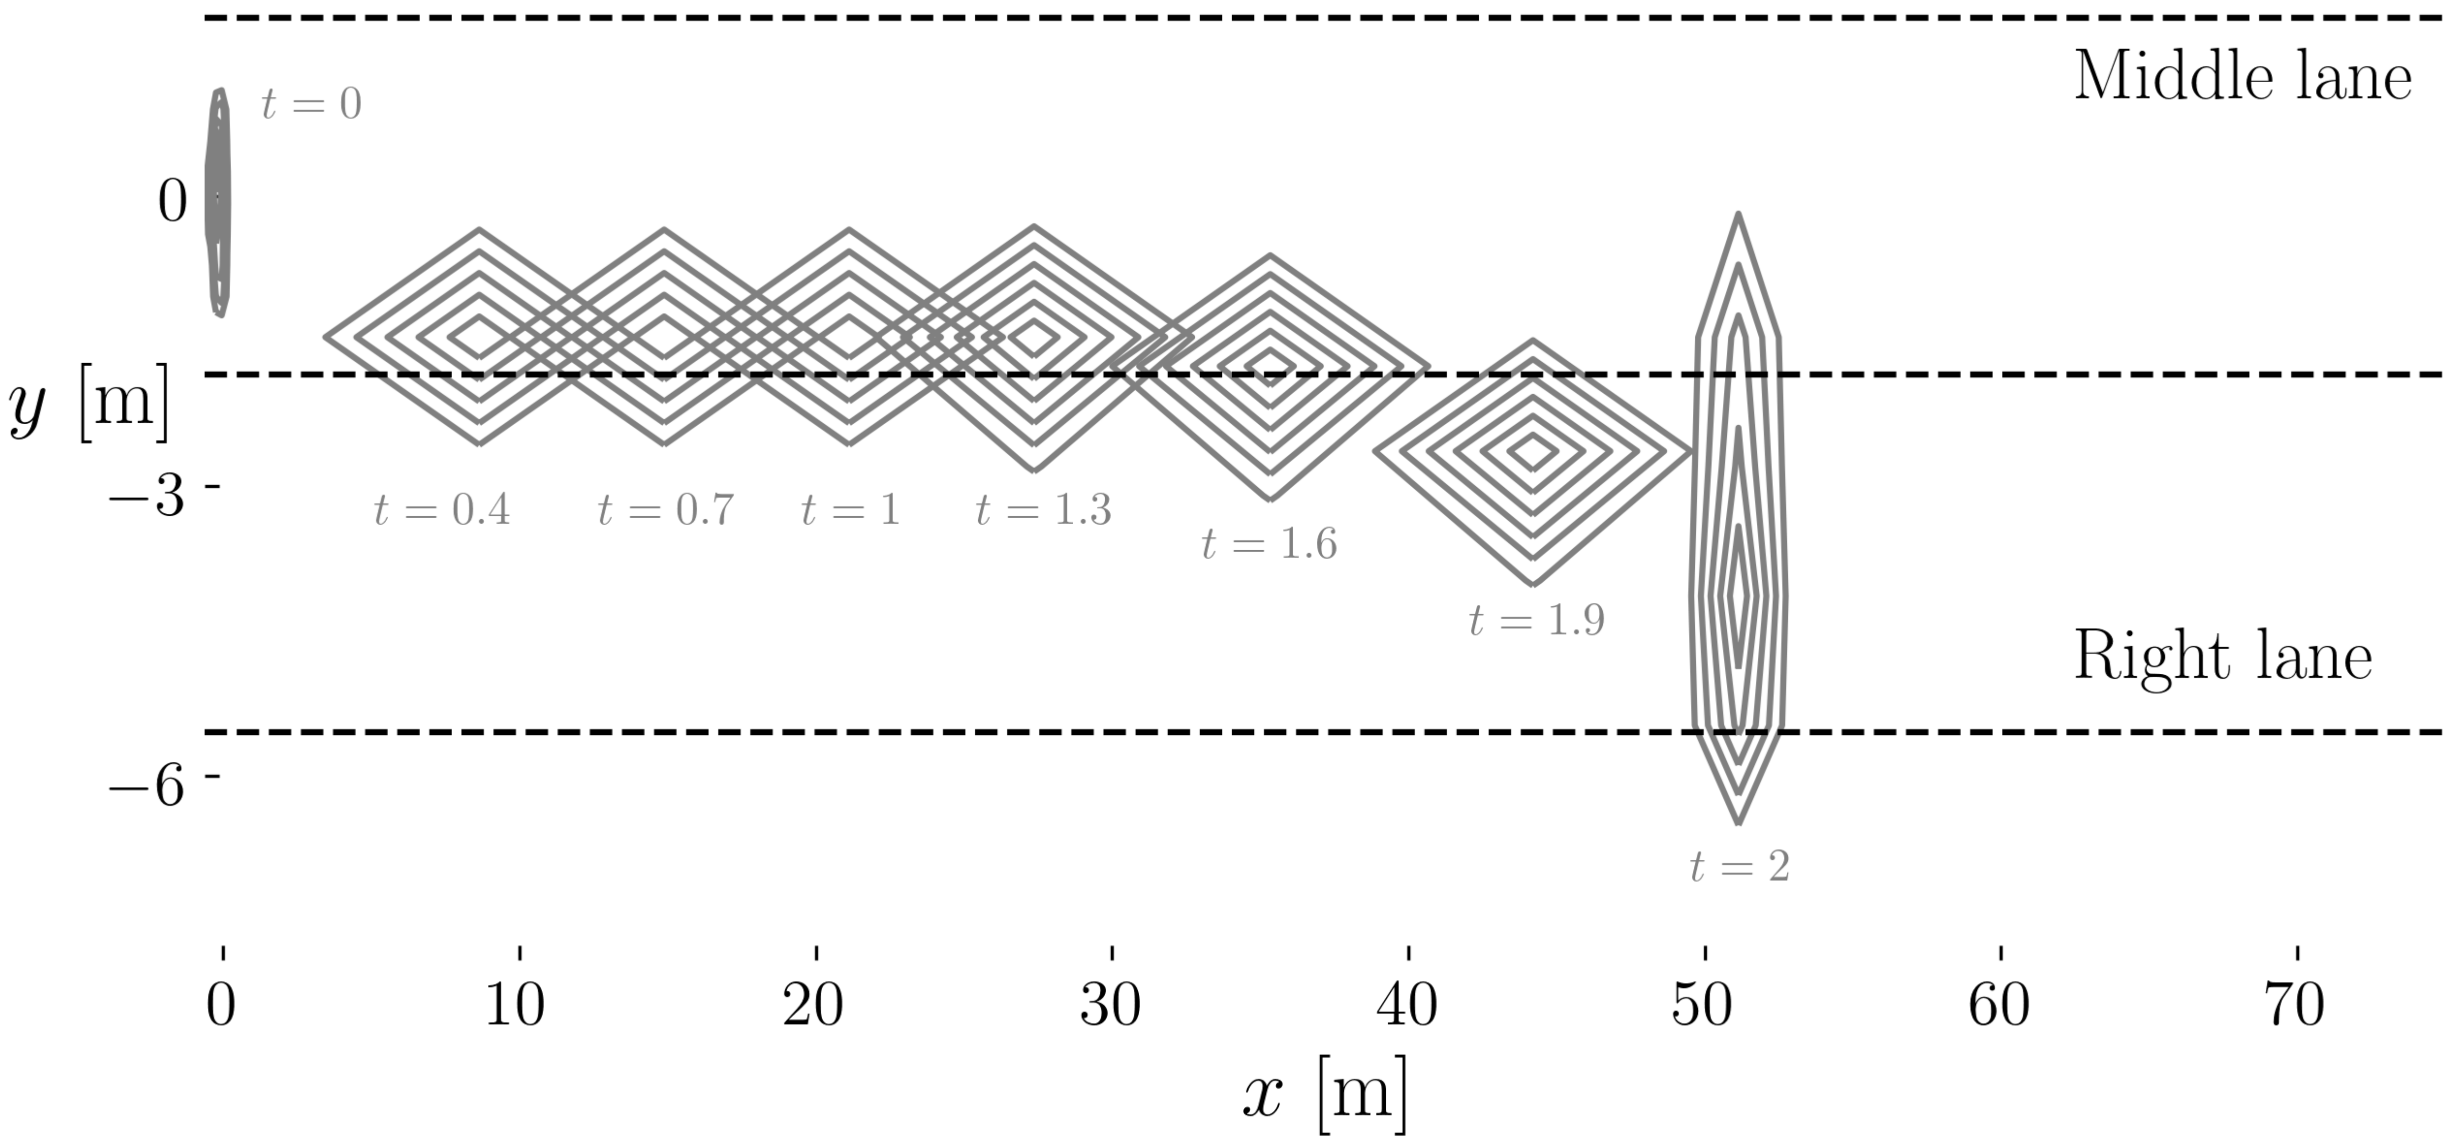

Details

— 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: 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*.

Outlook

- Distribution control undergoing rapid developments
- Lots of theory, algorithms and applications to be done
- Growing community in systems-control
- Excellent intersections with related communities: learning, information theory, robotics, systems biology

Thank You

Support: NSF CITRIS PEOPLE AND ROBOTS Ford