New Developments in Schrödinger Bridge, Stochastic Control and Stochastic Learning

Abhishek Halder

Department of Aerospace Engineering
Iowa State University

Joint work with students and collaborators

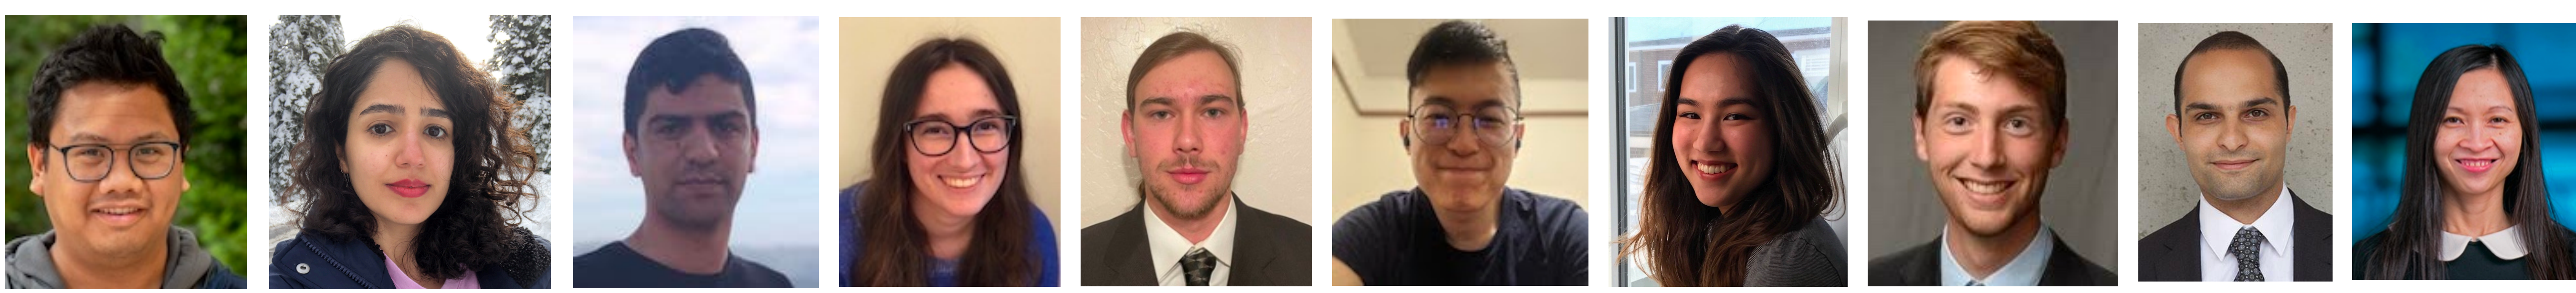

Analysis and Probability Seminar, Department of Mathematics, Iowa State University
March 06, 2024

Theme of this talk

Control and learning of measures/densities

Density Control: Generalized Schrödinger Bridge

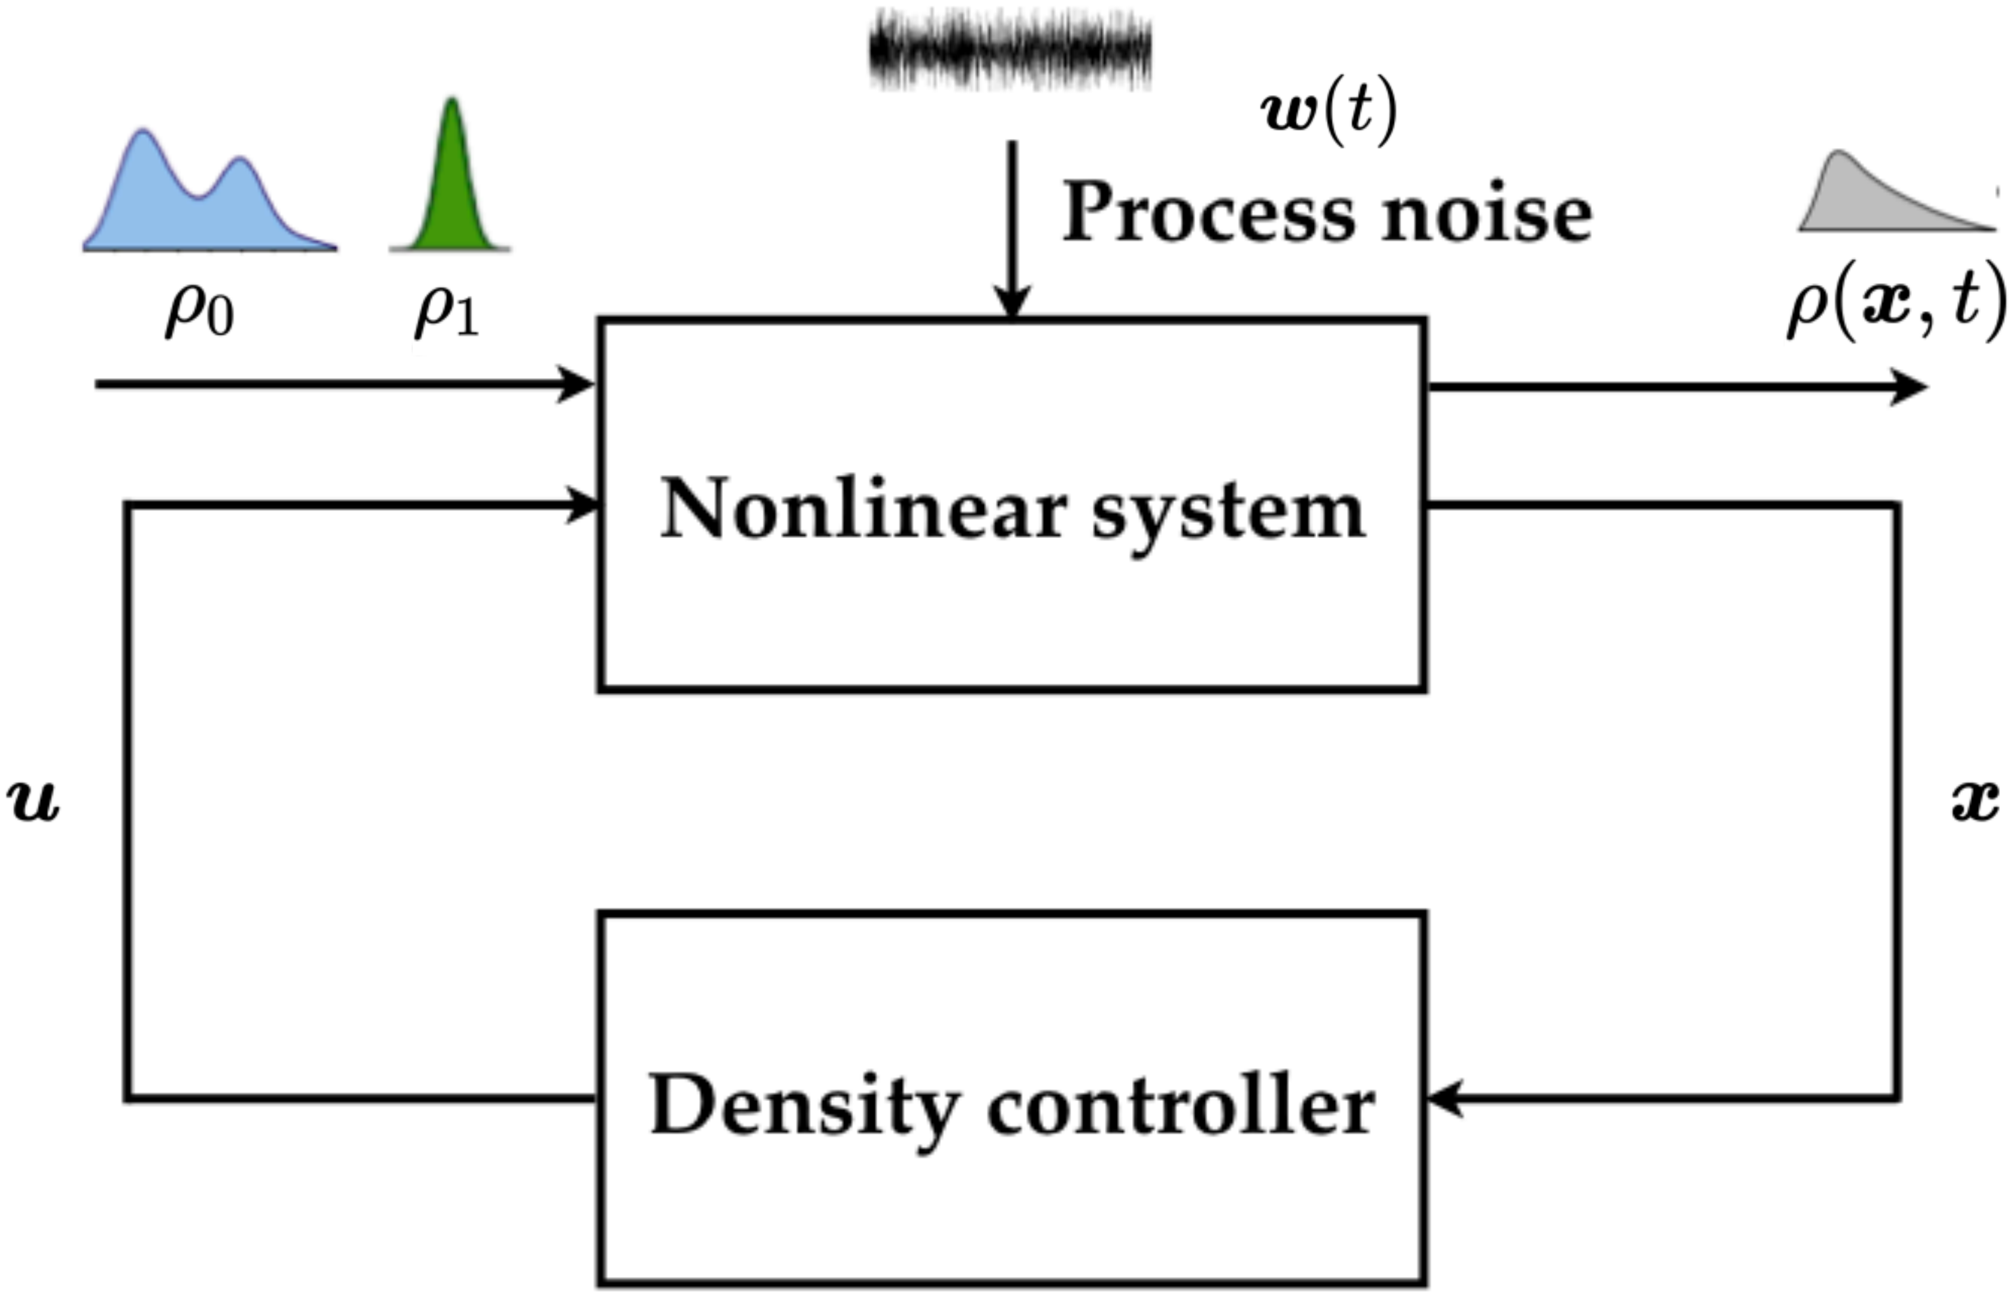

Motivating Applications

Distribution ~ Probability

Spacecraft landing with desired statistical accuracy

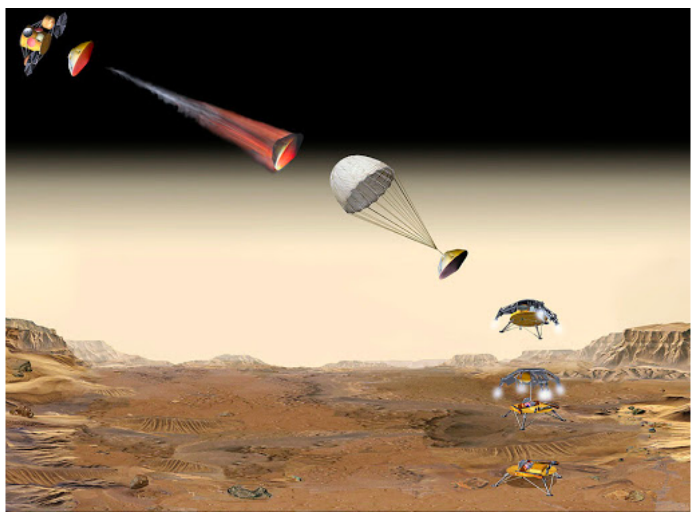

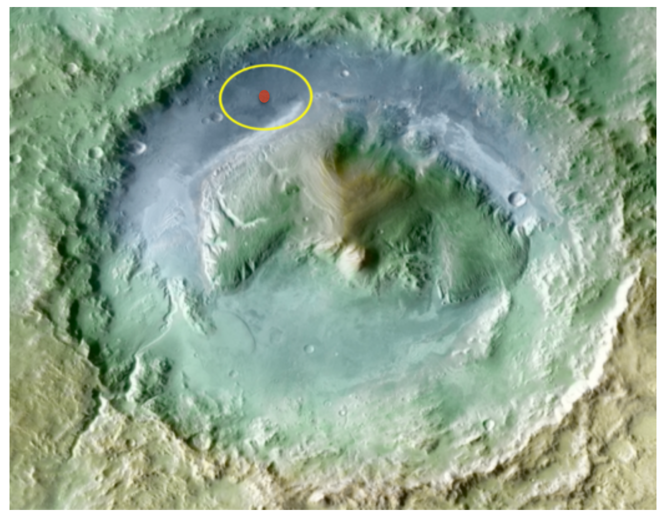

Gale Crater (4.49S, 137.42E)

Risk management for automated driving in multi-lane highways

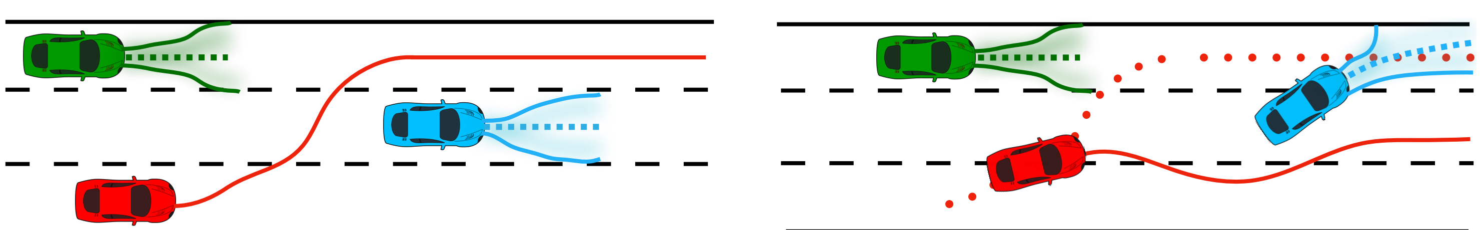

Control of uncertainties

Distribution ~ Population

Dynamic shaping of swarms

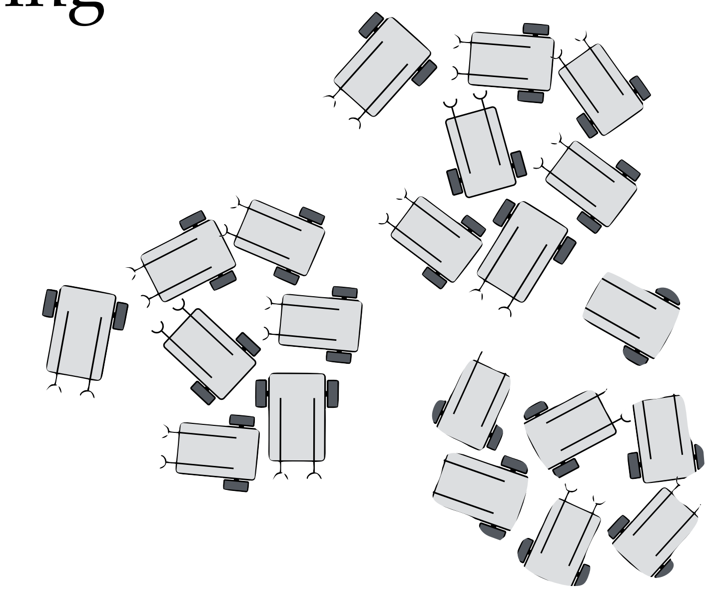

Feedback sync. and desync. of neuronal population

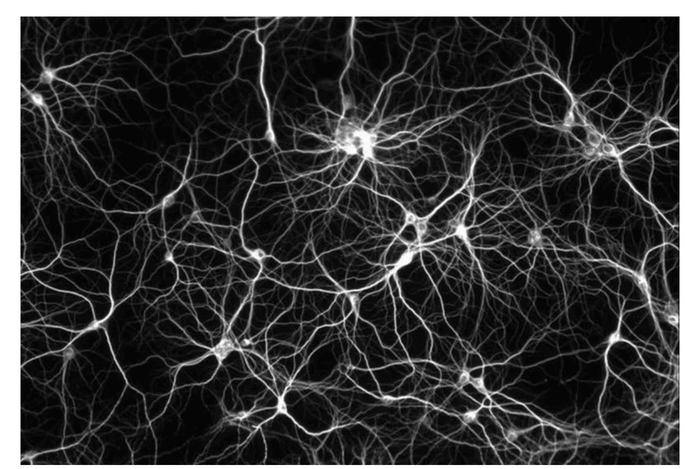

Control of ensemble

State Feedback Density Steering

Steer joint state PDF via feedback control over finite time horizon

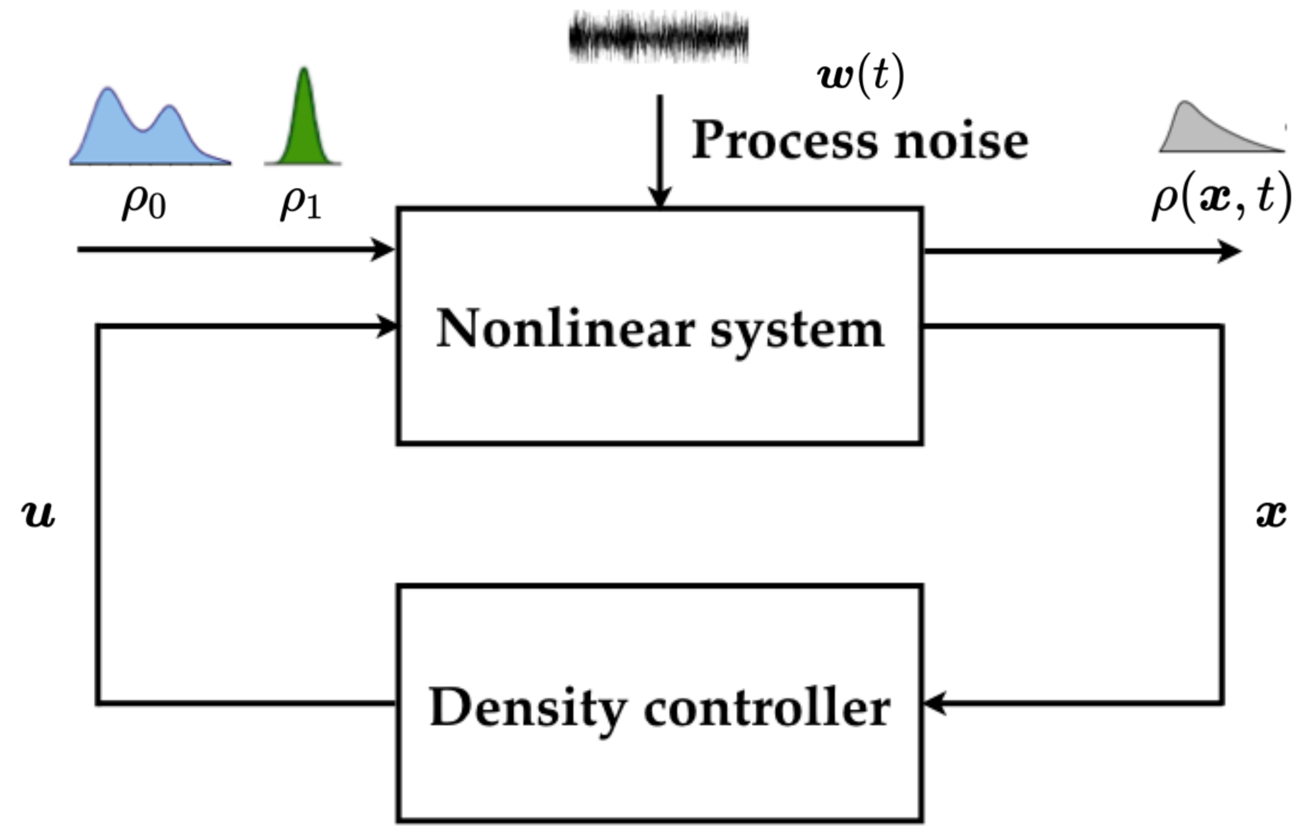

Common scenario: $\boldsymbol{G} \equiv \boldsymbol{B}$

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

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

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 ⇝ Optimal Mass Transport

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

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

subject to

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

0

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

Controlled Liouville PDE

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

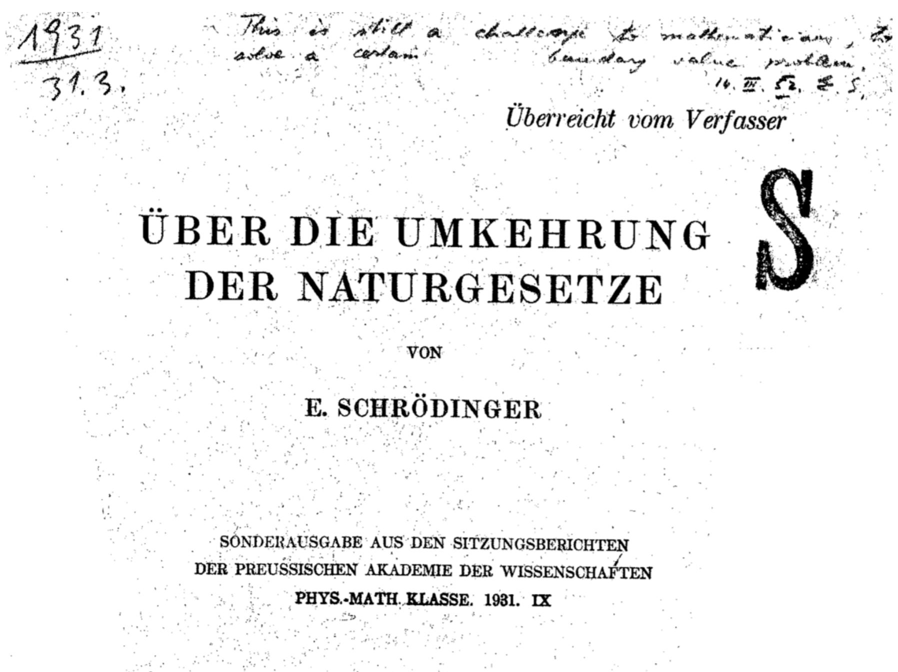

1931
31.3.

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

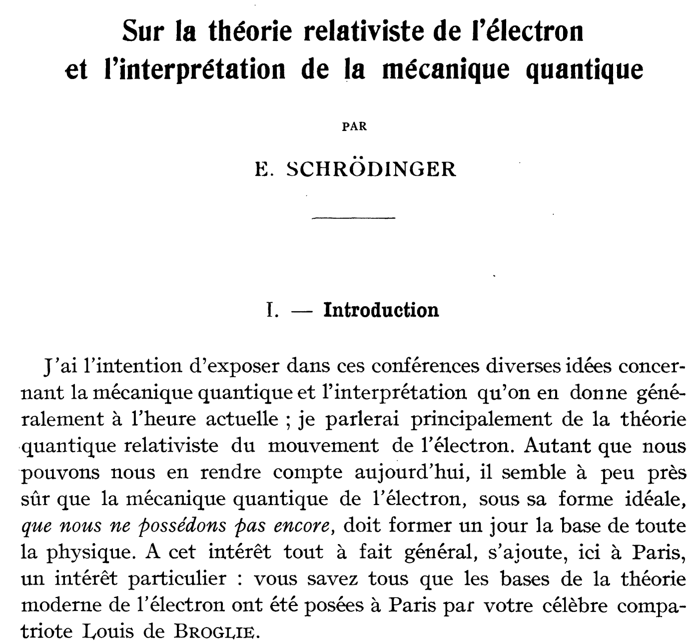

Sur la théorie relativiste de l'électron
et l'interprétation de la mécanique quantique

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:

$$\left(\rho^{\text{opt}}, \psi\right) \mapsto \left(\widehat{\varphi}, \varphi\right)$$ — **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: $\boldsymbol{u}^{\text{opt}}(\boldsymbol{x}, t) = 2\boldsymbol{B}^{\top}\nabla_{\boldsymbol{x}} \log \varphi(\boldsymbol{x}, t)$

What Exactly are Schrödinger Factors?

Classical: $$\rho^{\mathrm{opt}}(\boldsymbol{x},t) = \varphi(\boldsymbol{x},t)\widehat{\varphi}(\boldsymbol{x},t)$$

$$\left(\frac{\partial}{\partial t} + \frac{1}{2}\Delta - q\right)\varphi = 0$$ [Backward reaction-diffusion PDE]

$$\left(\frac{\partial}{\partial t} - \frac{1}{2}\Delta + q\right)\widehat{\varphi} = 0$$ [Forward reaction-diffusion PDE]

Quantum: $$\rho^{\mathrm{opt}}(\boldsymbol{x},t) = \Psi(\boldsymbol{x},t)\widehat{\Psi}(\boldsymbol{x},t)$$ [Born's relation]

wave function

$$\left(\sqrt{-1}\frac{\partial}{\partial t} + \frac{1}{2}\Delta - q\right)\Psi = 0$$ [Schrödinger PDE]

$$\left(-\sqrt{-1}\frac{\partial}{\partial t} - \frac{1}{2}\Delta + q\right)\widehat{\Psi} = 0$$ [Adjoint Schrödinger PDE]

Feedback Synthesis via the Schrödinger System

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

Uncontrolled forward-backward Kolmogorov PDEs:

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

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

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 $\left(\varphi_1, \hat{\varphi}_0\right)$

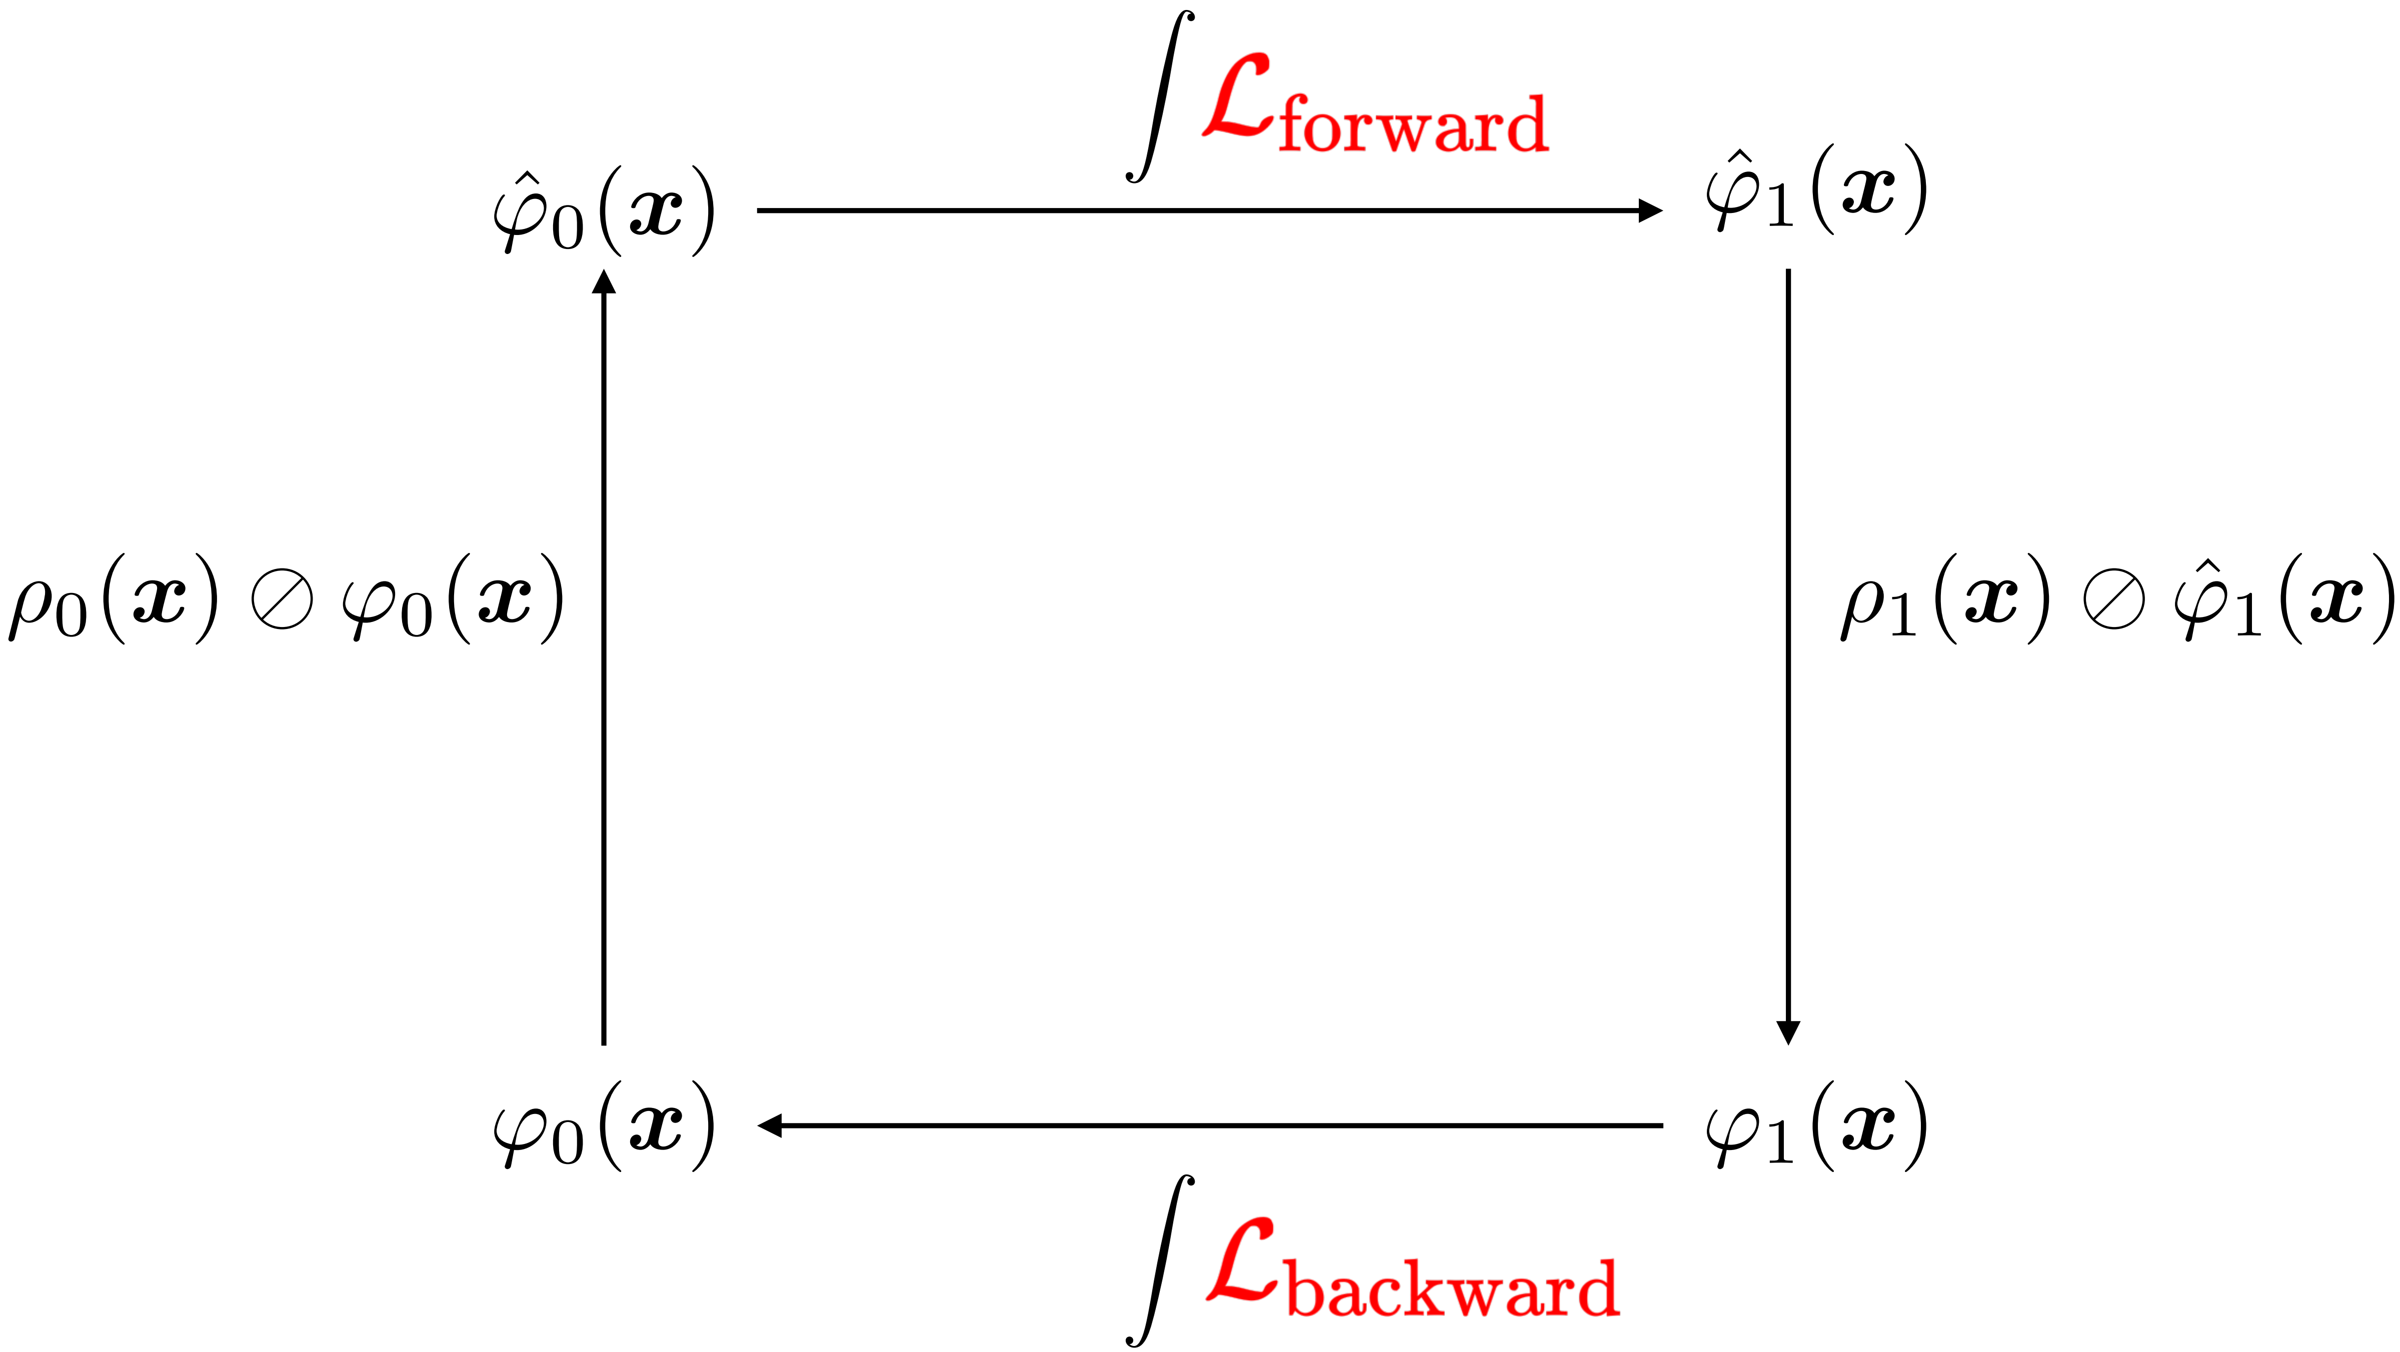

This recursion is contractive in the Hilbert's projective metric!!

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

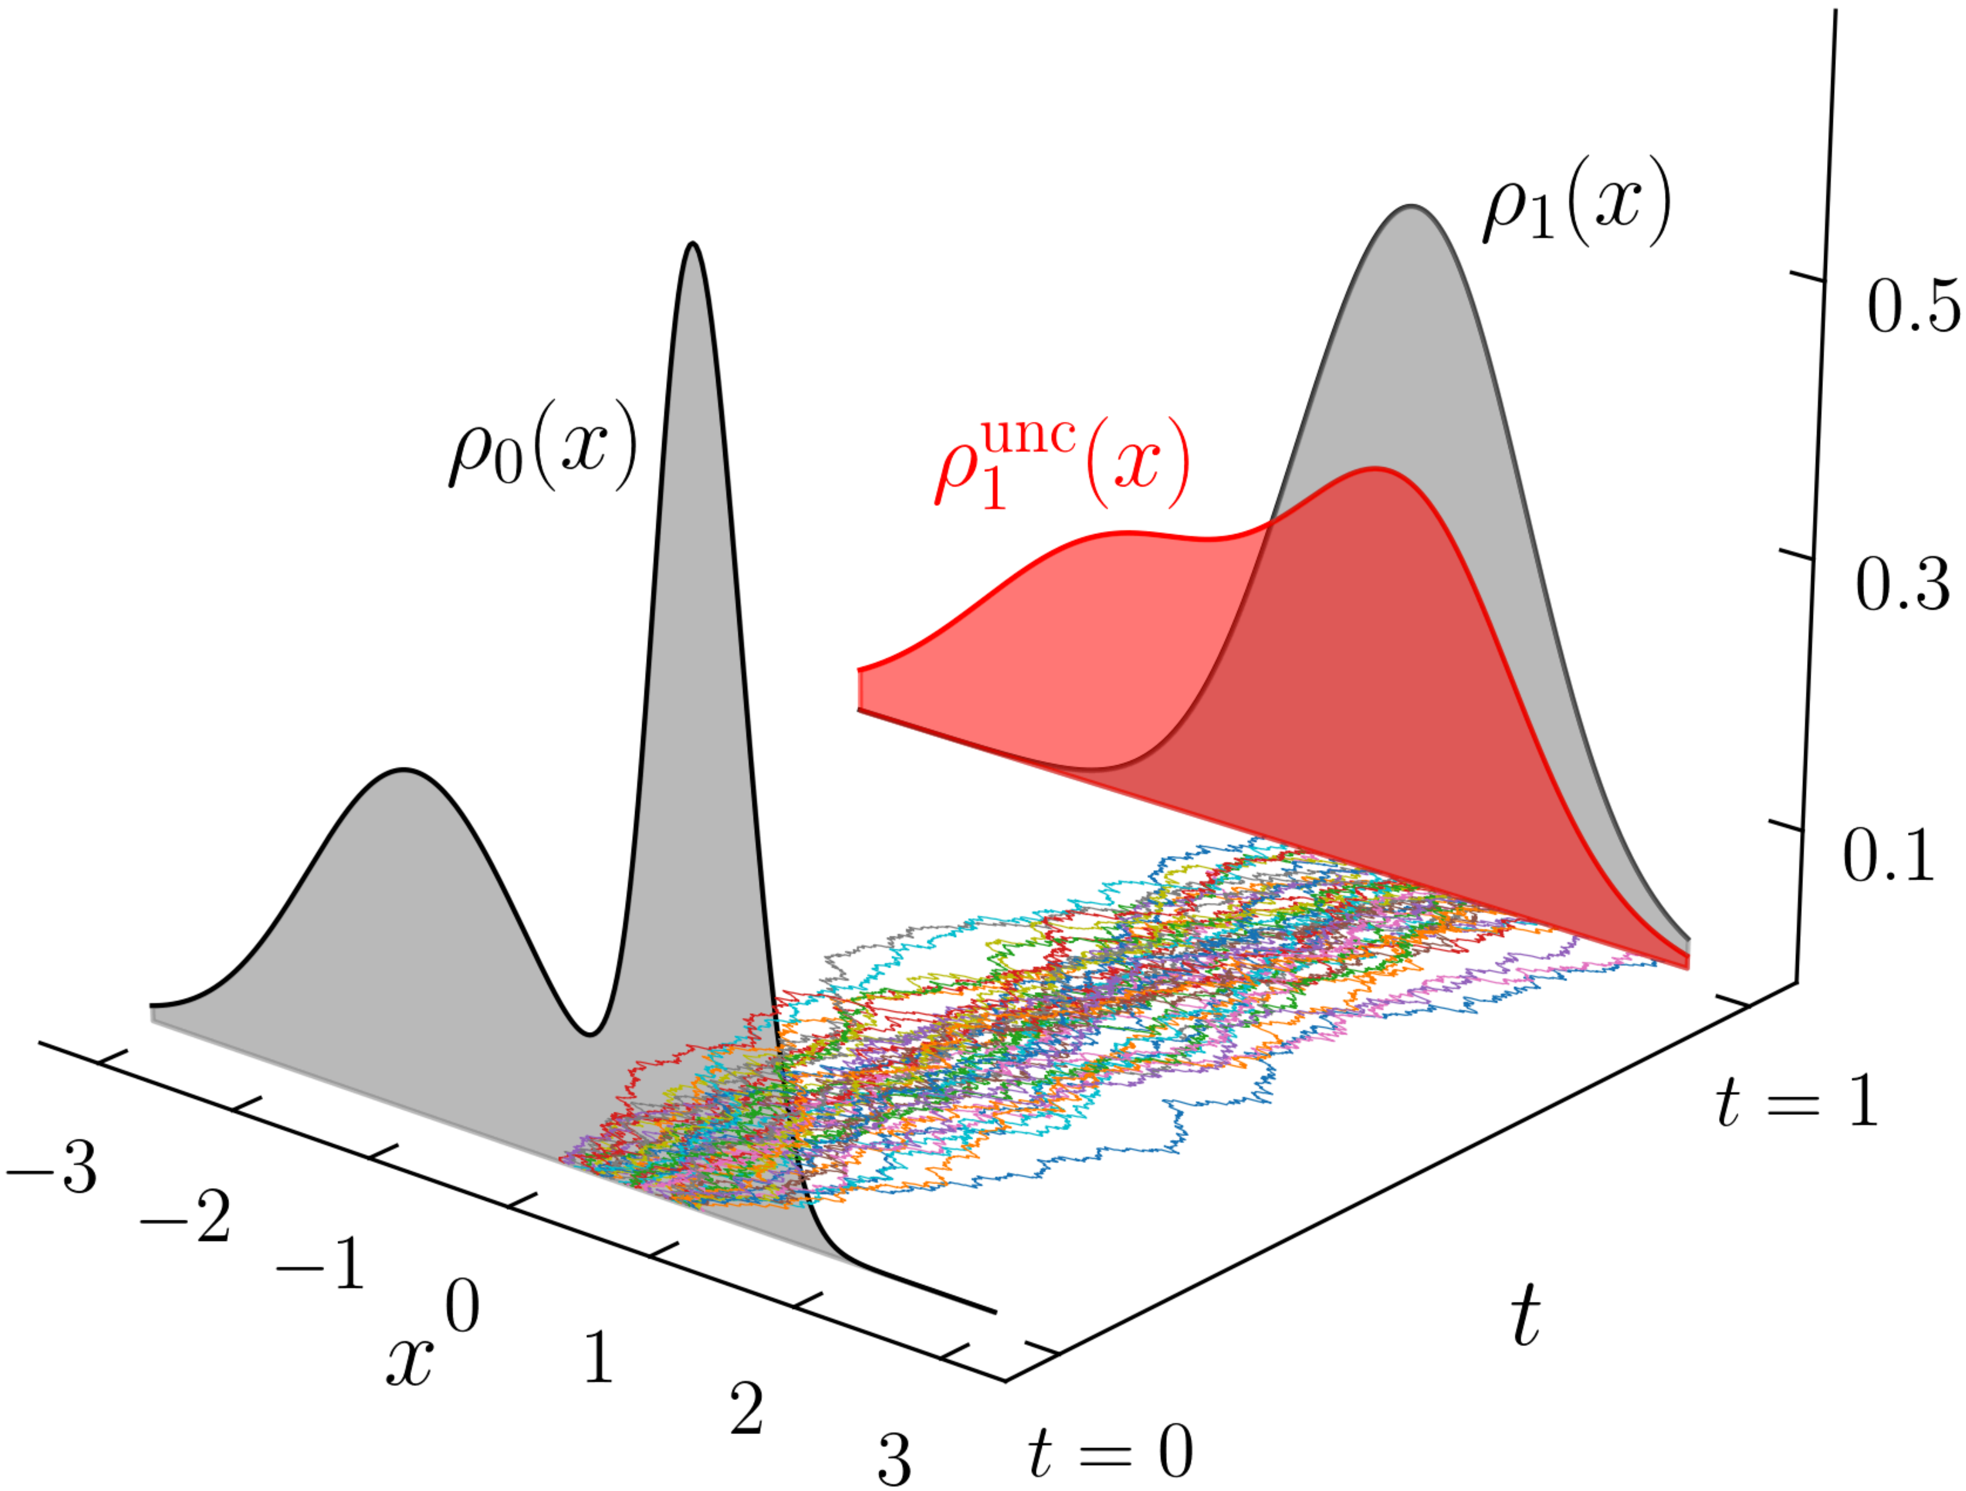

Zero prior dynamics

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

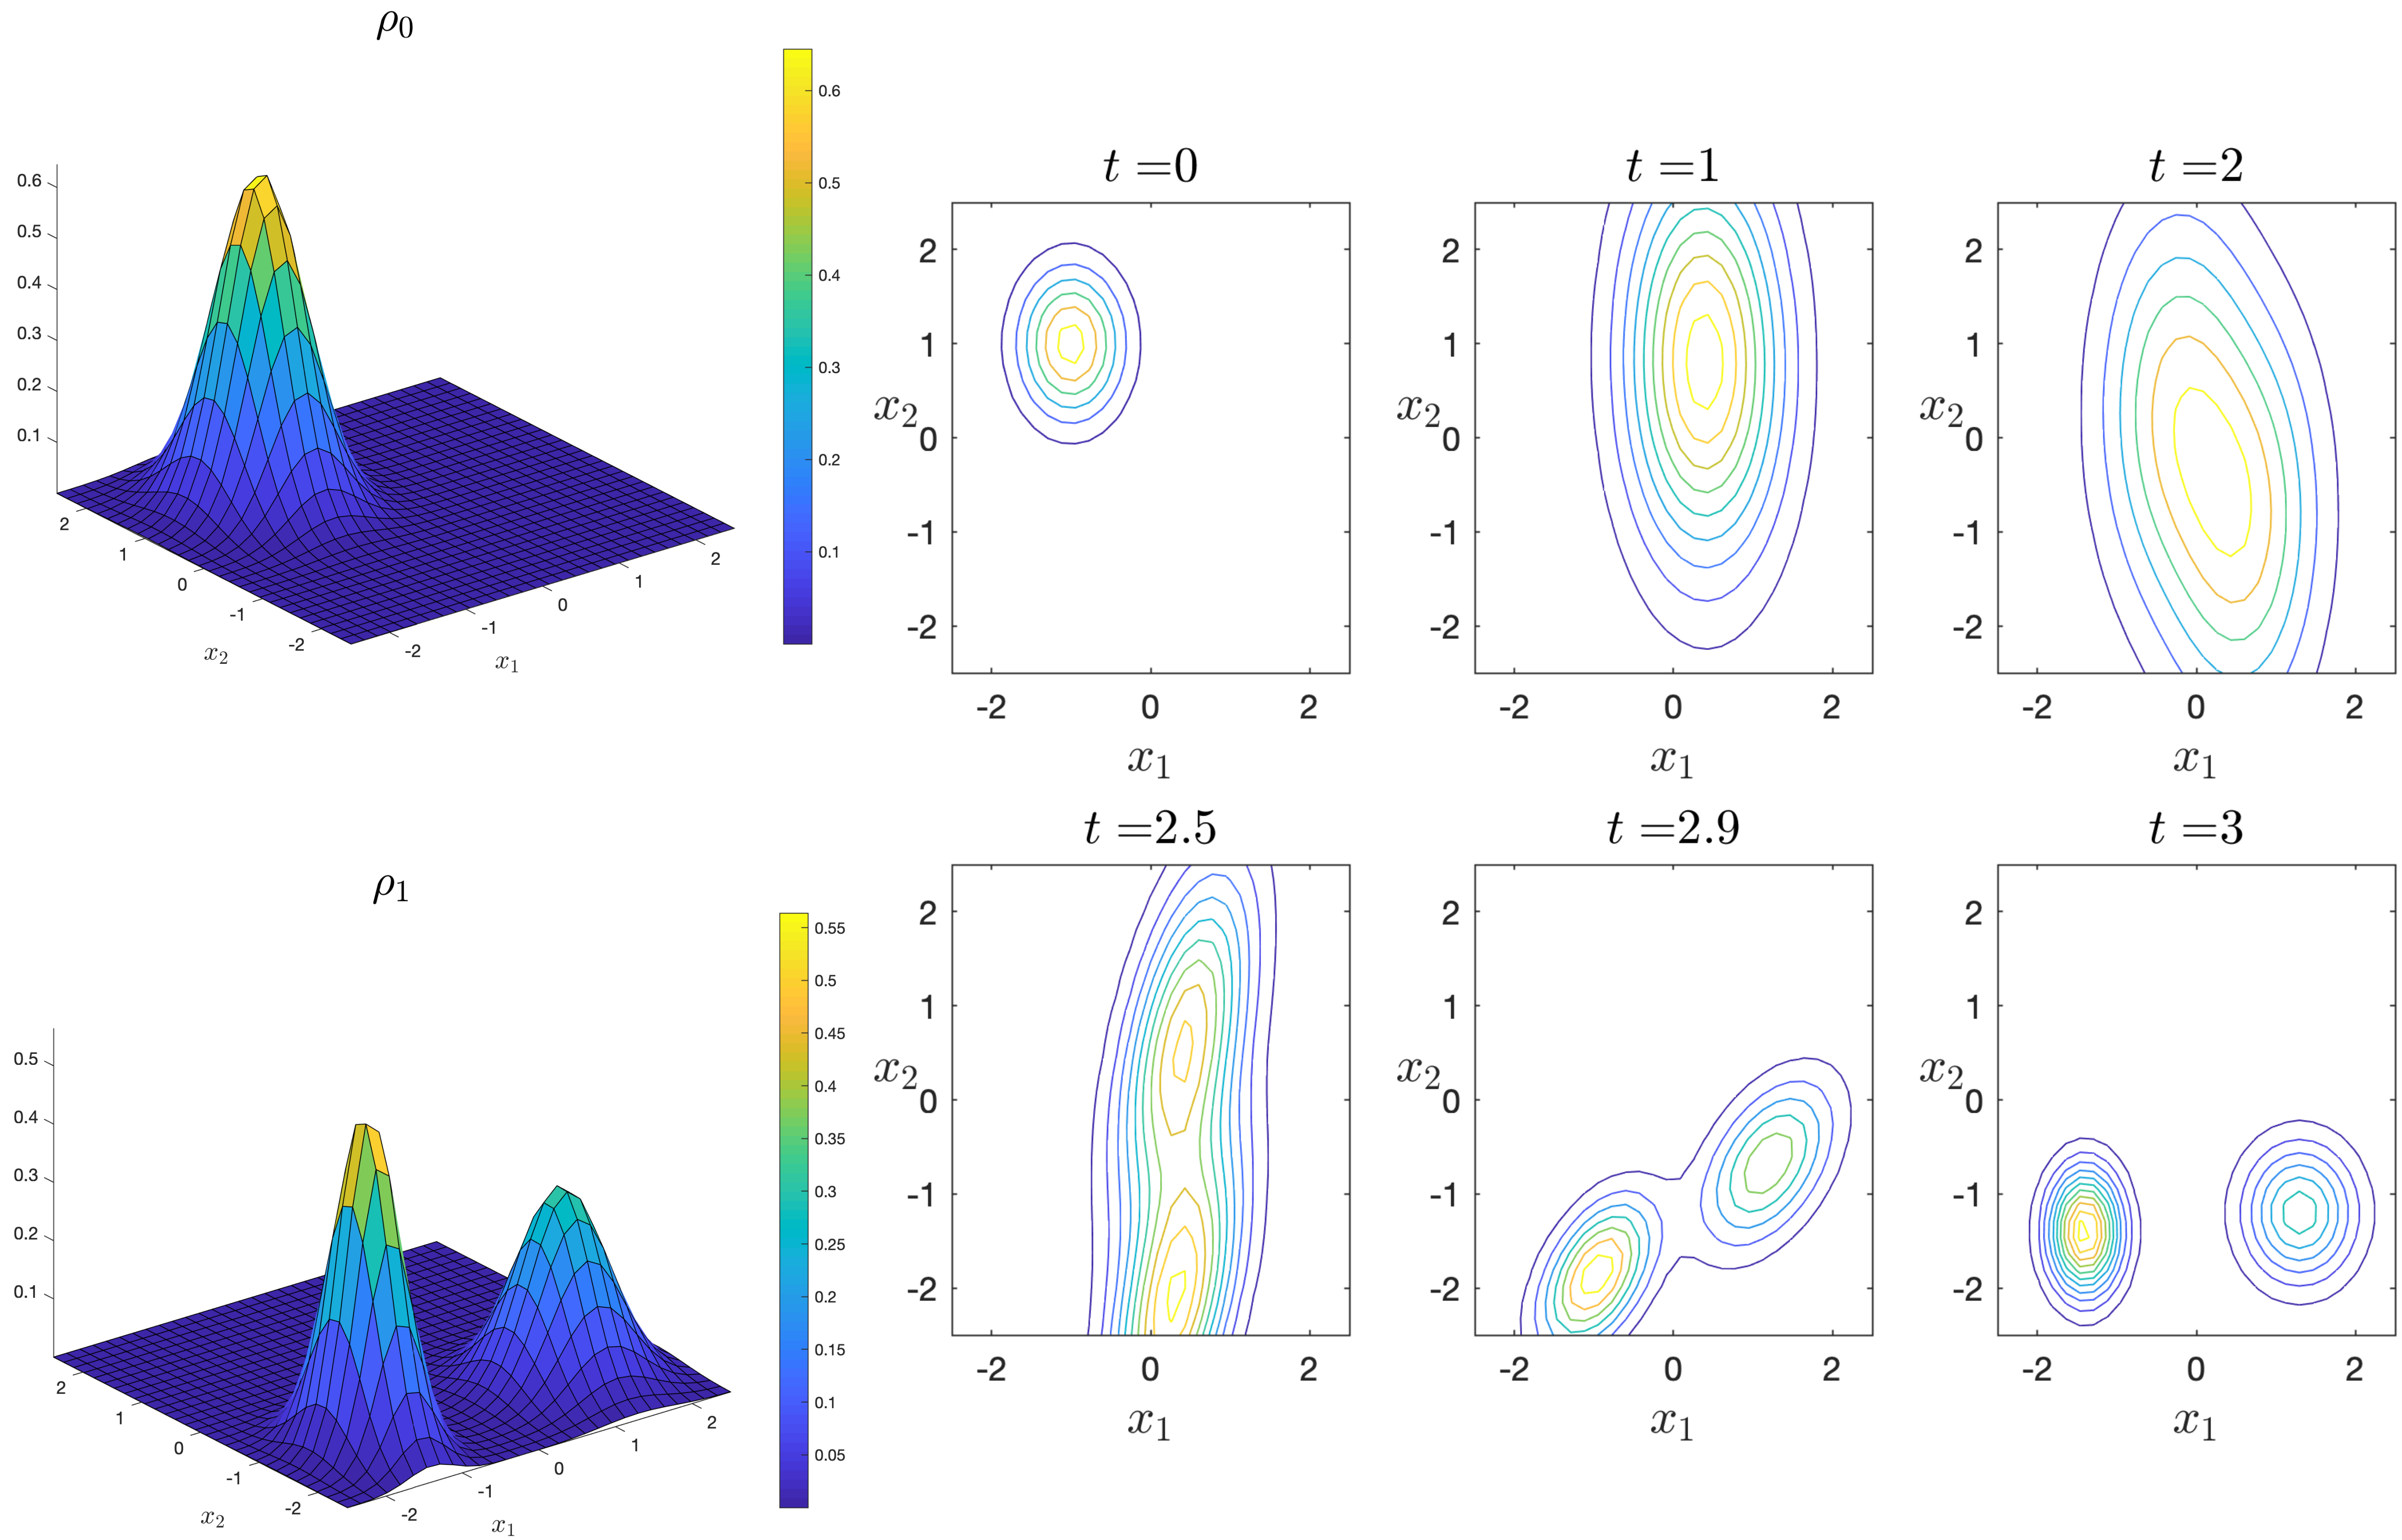

Linear prior dynamics

Feedback Density Control: Nonlinear Grad. Drift

Uncontrolled joint PDF evolution:

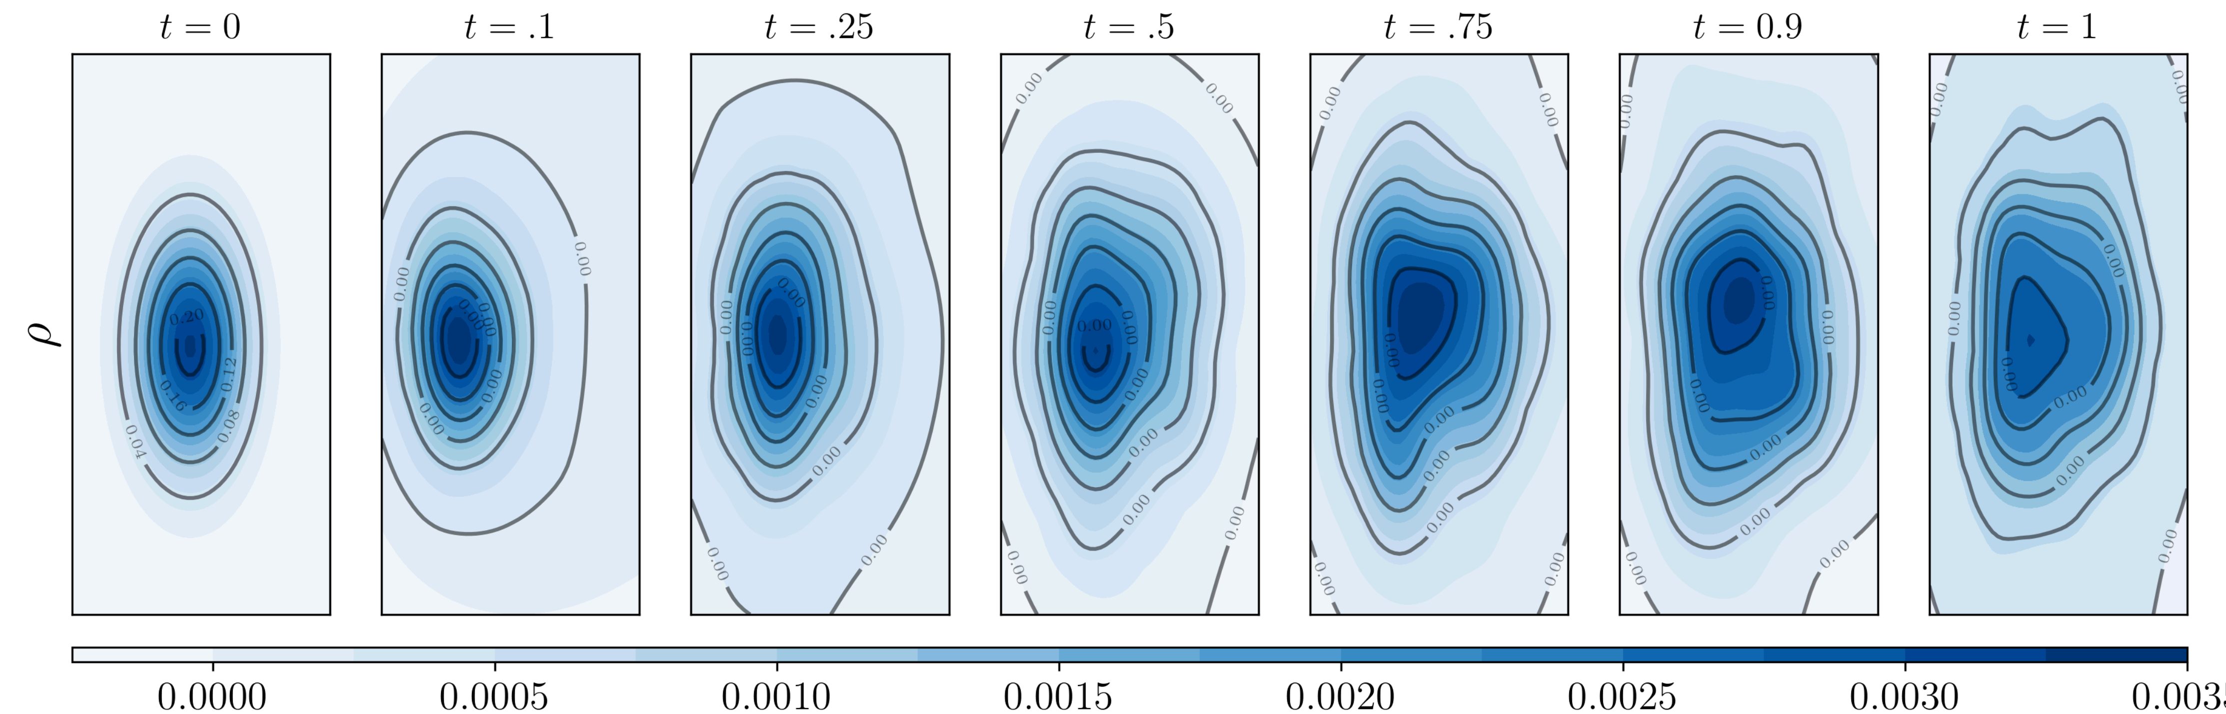

Optimal controlled joint PDF evolution:

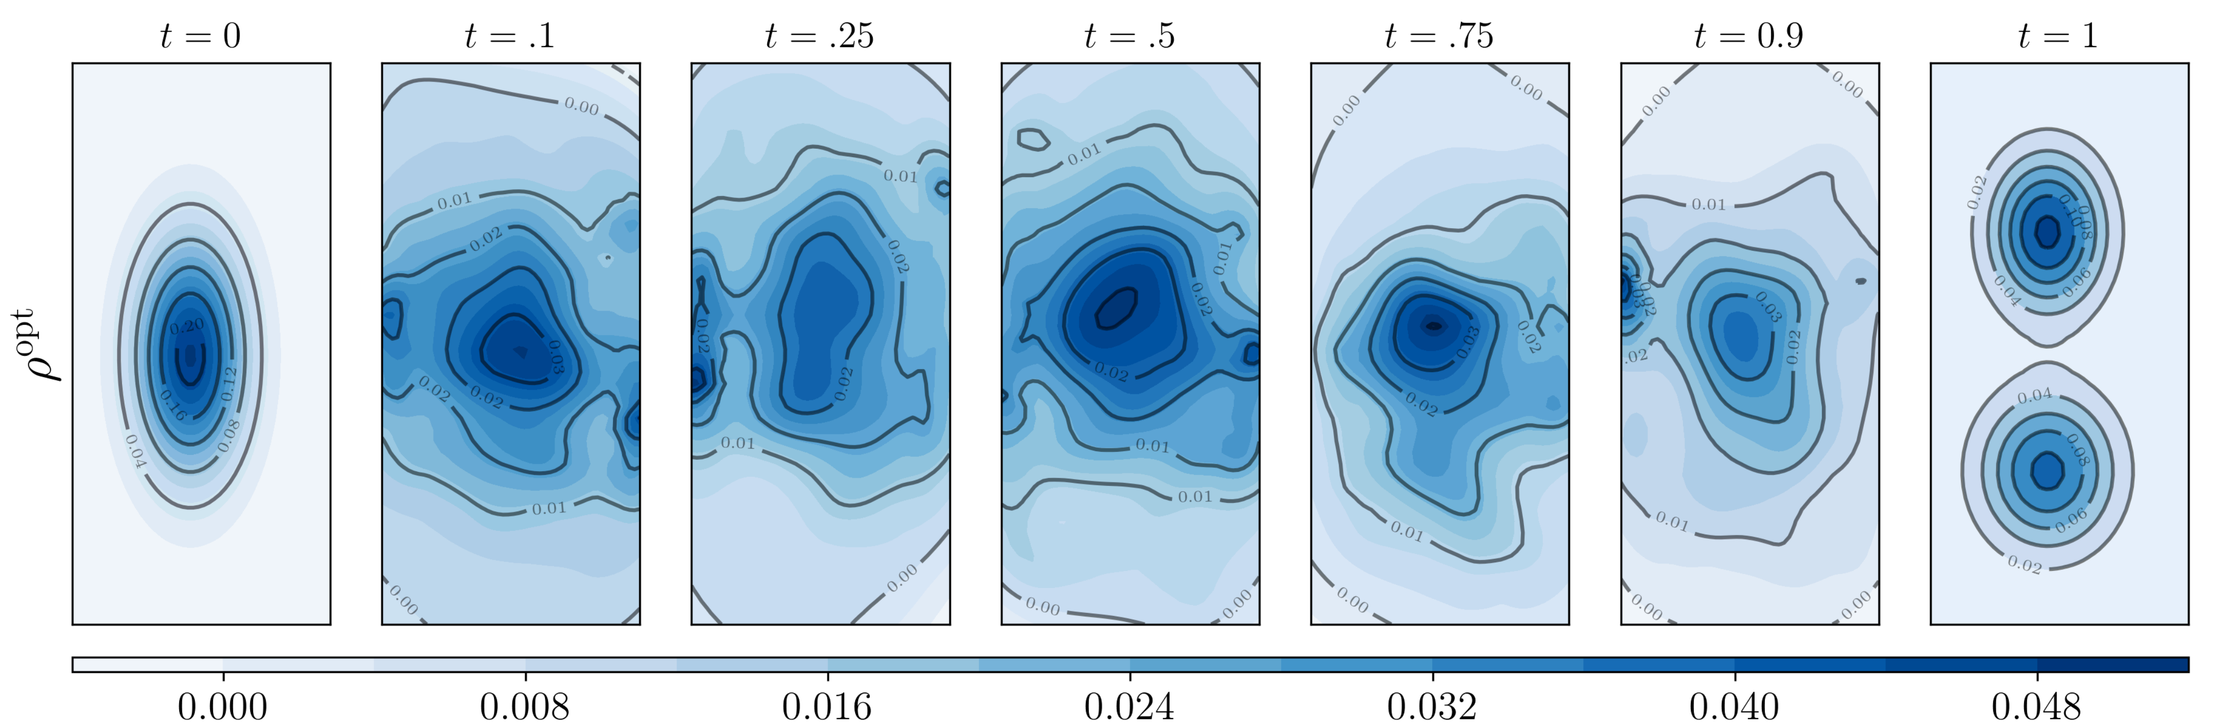

Feedback Density Control: Mixed Conservative-Dissipative Drift

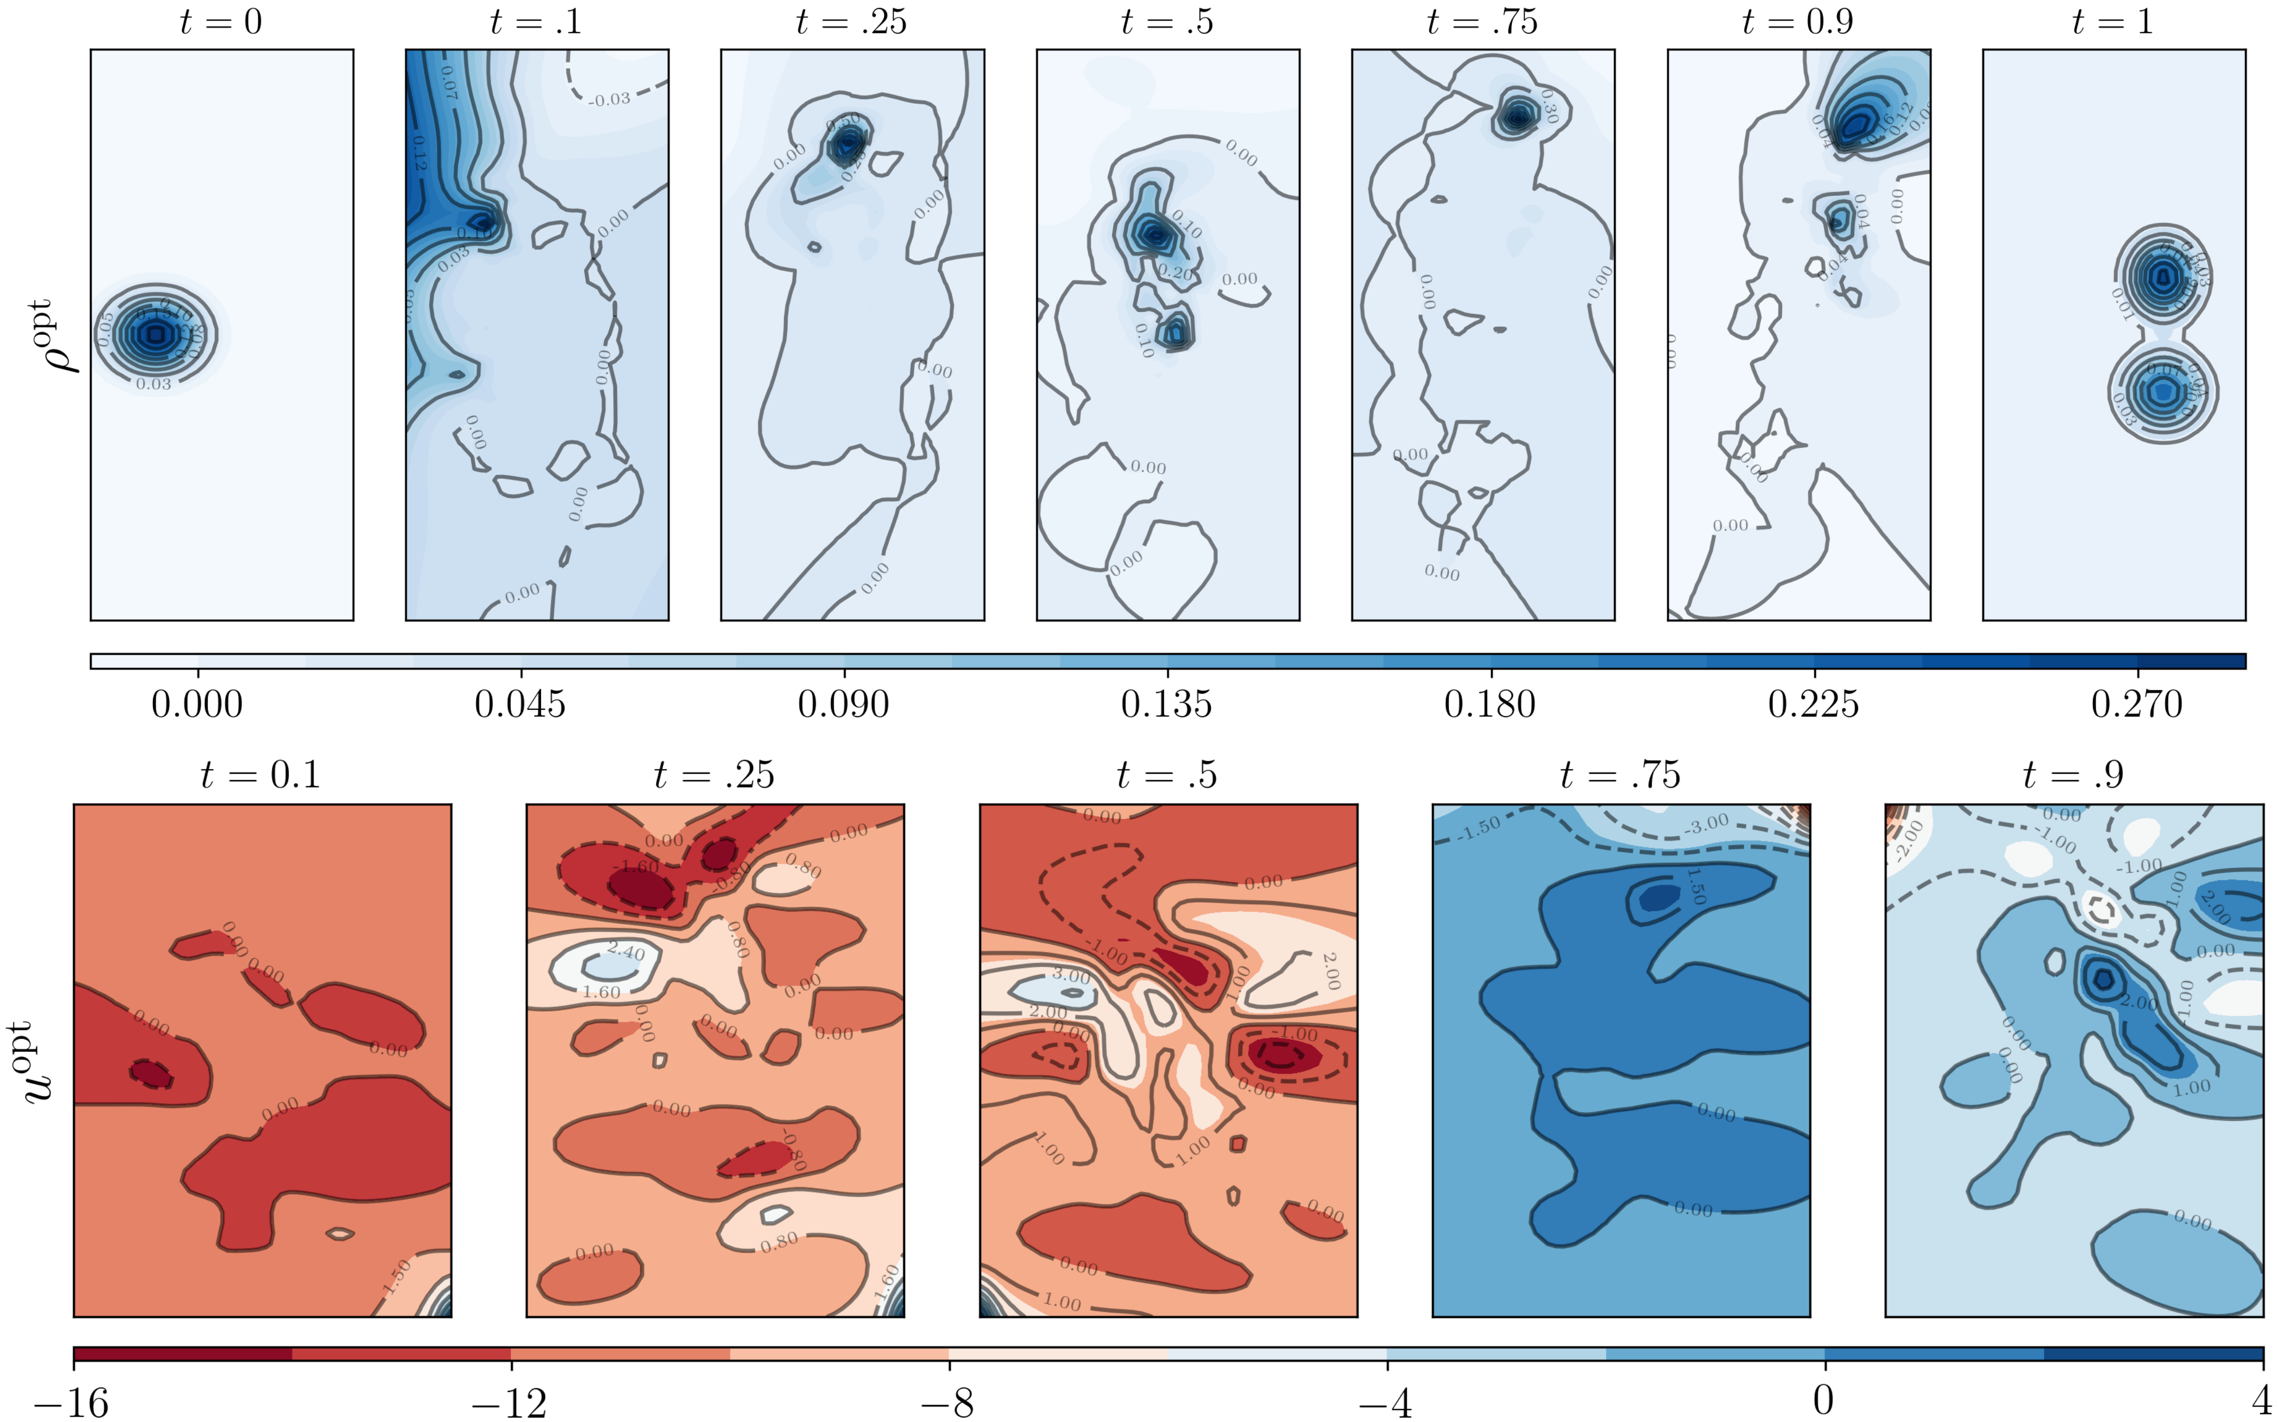

K.F. Caluya and A.H., Wasserstein proximal algorithms for the Schrödinger bridge problem: density control with nonlinear drift, *IEEE TAC* 2021.

Application: Multi-lane Automated Driving

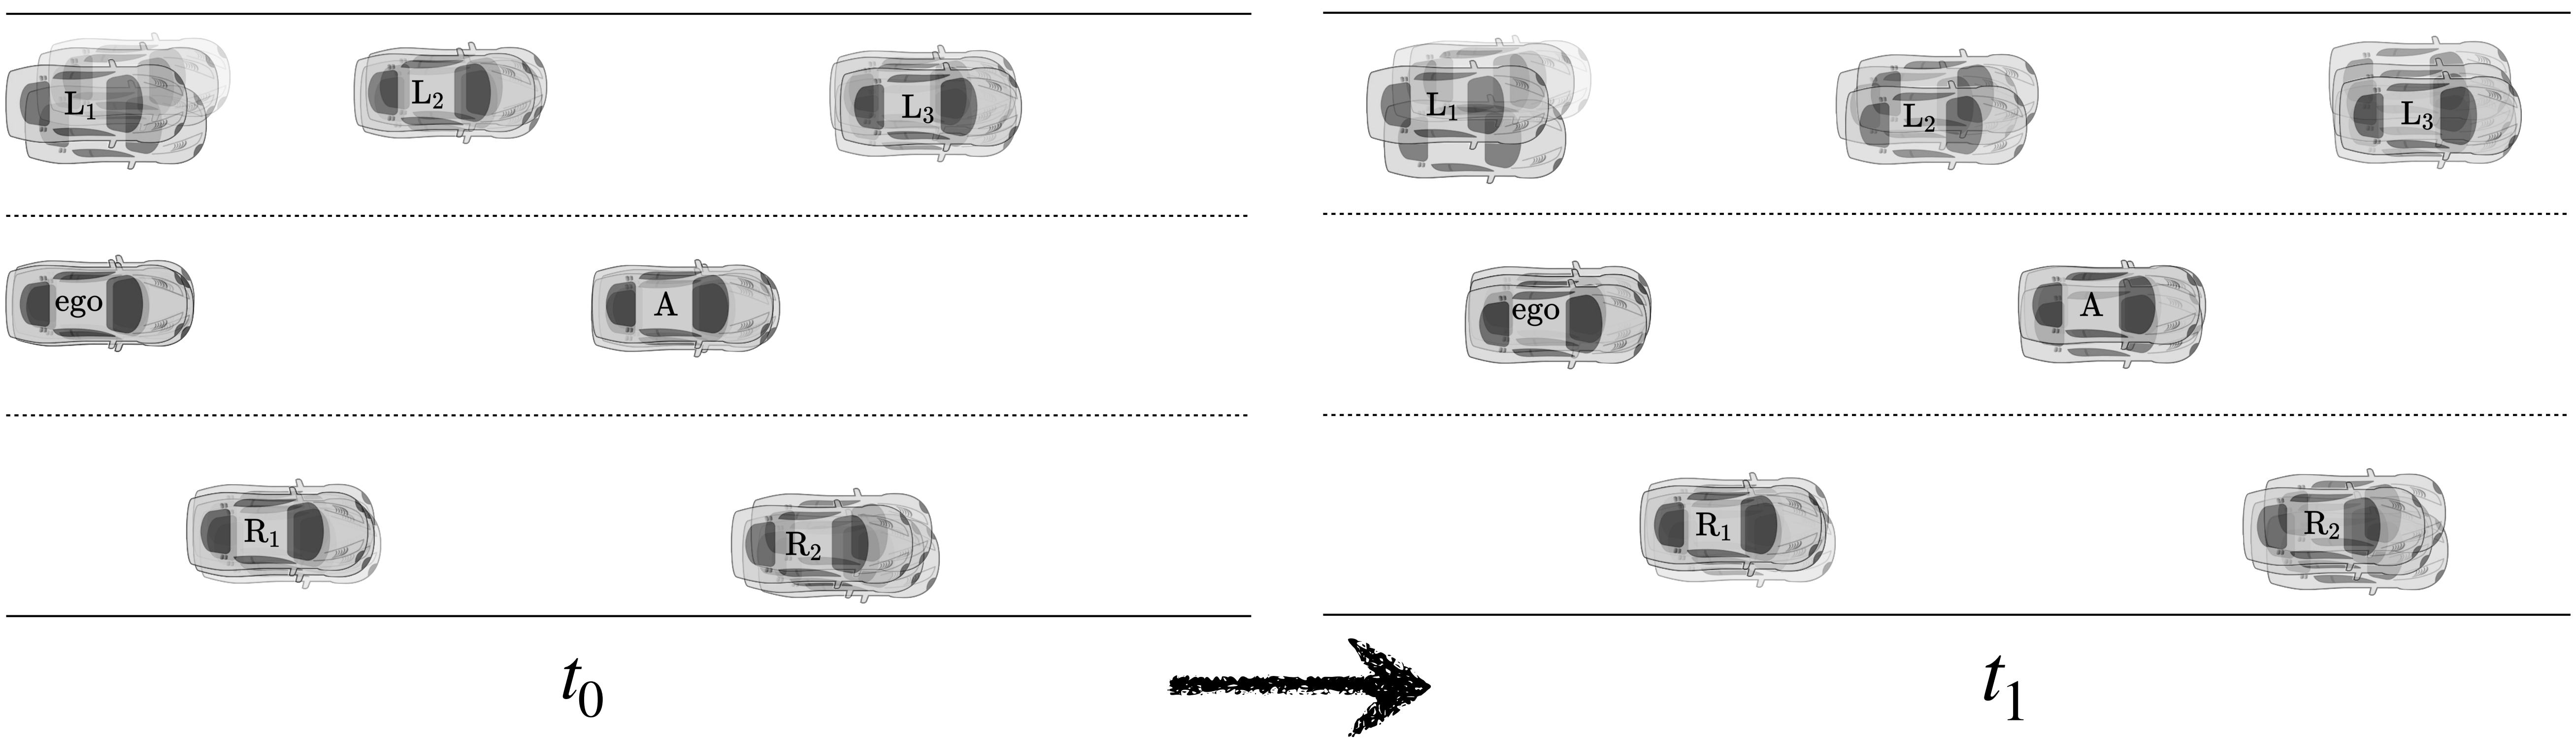

x marginals

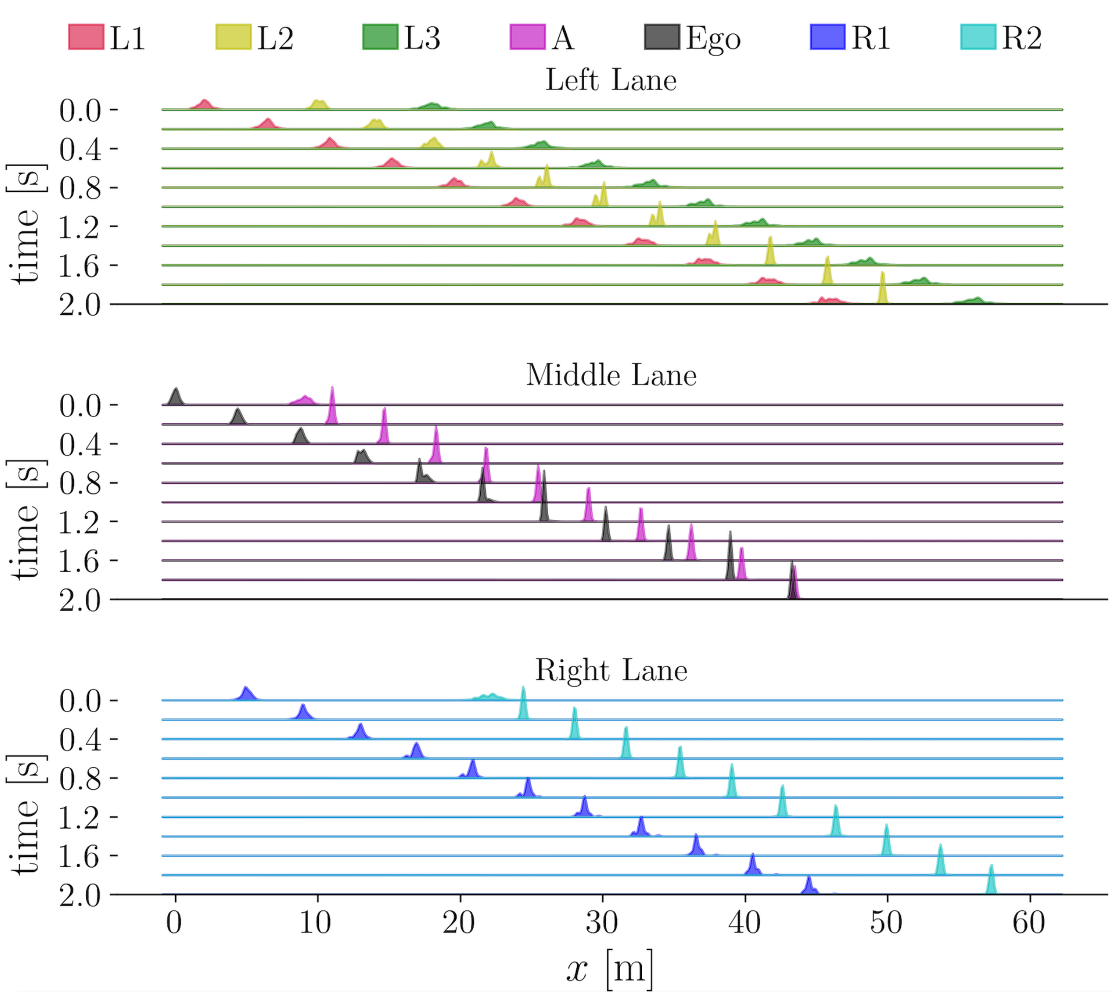

y marginals

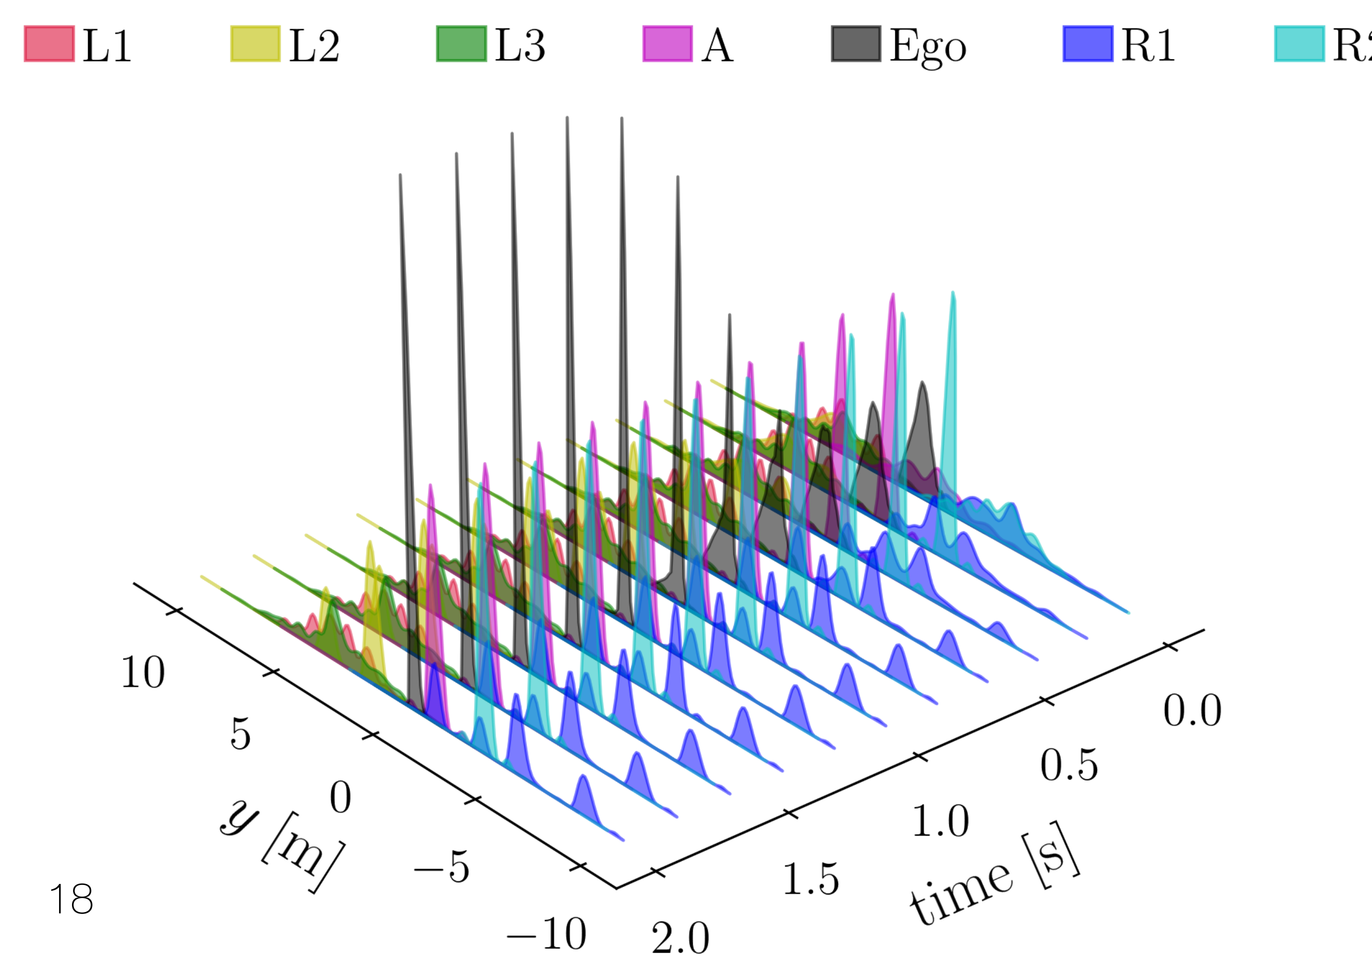

Application: Multi-lane Automated Driving

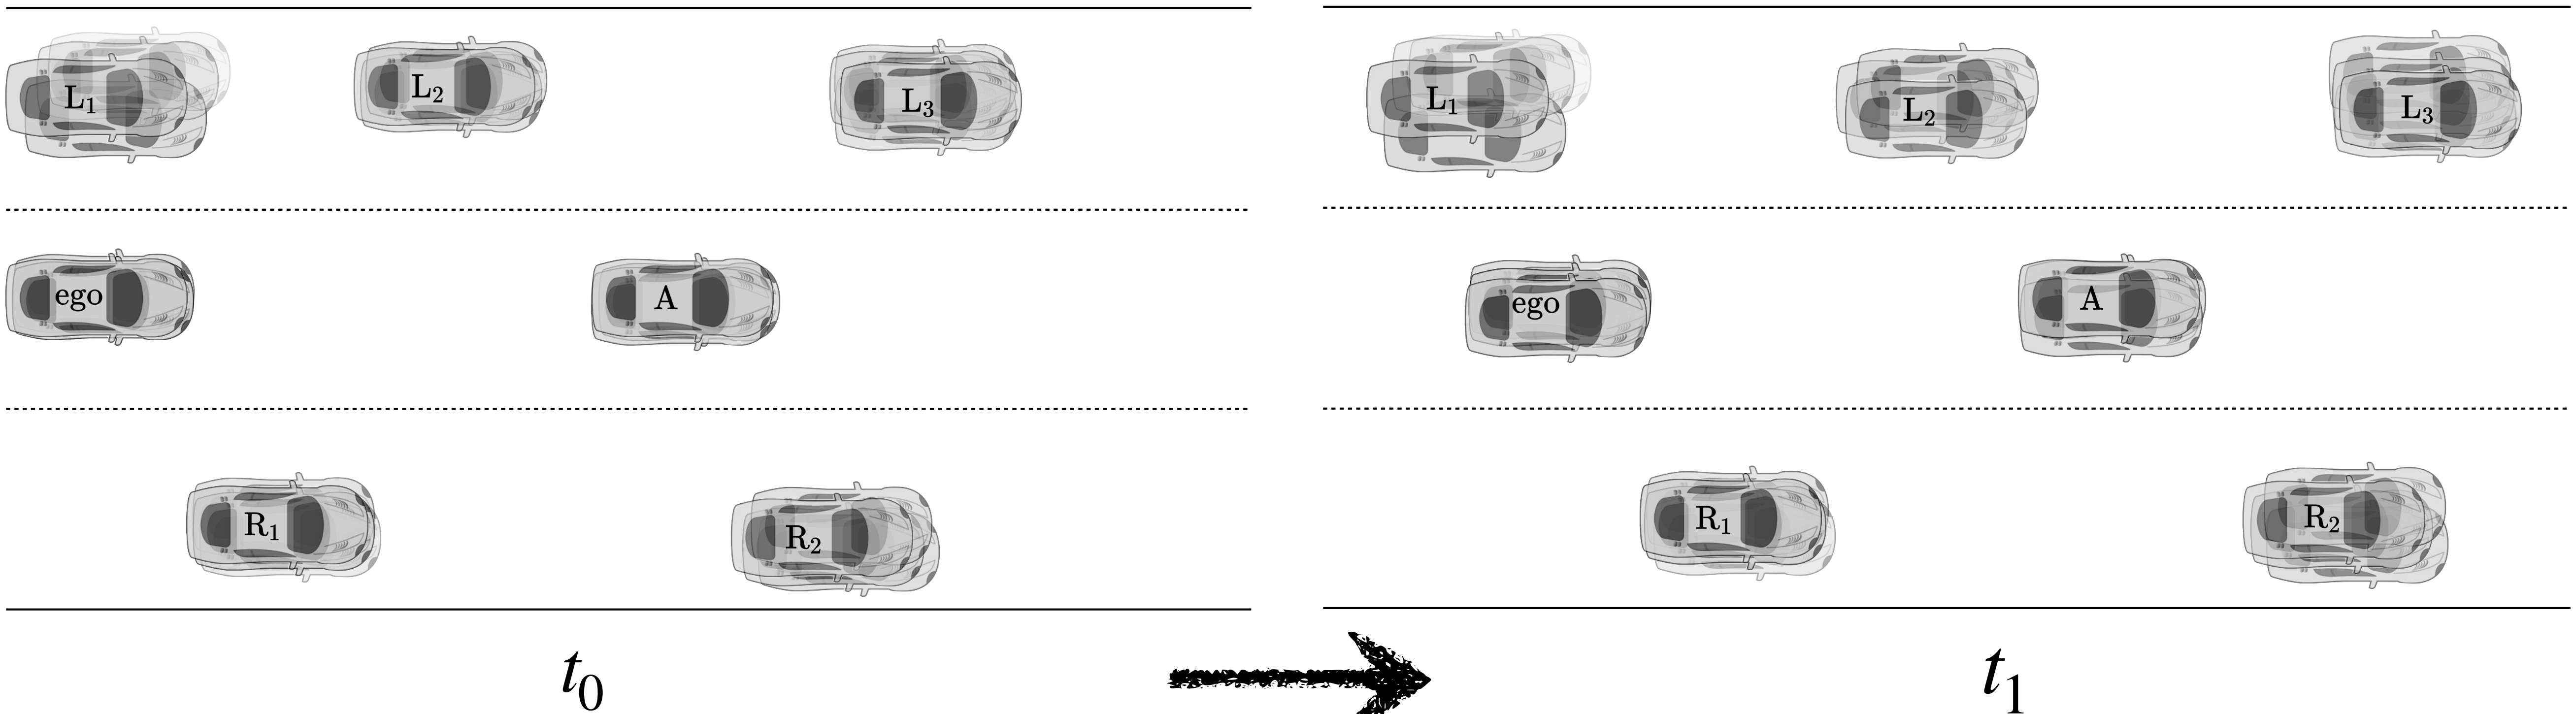

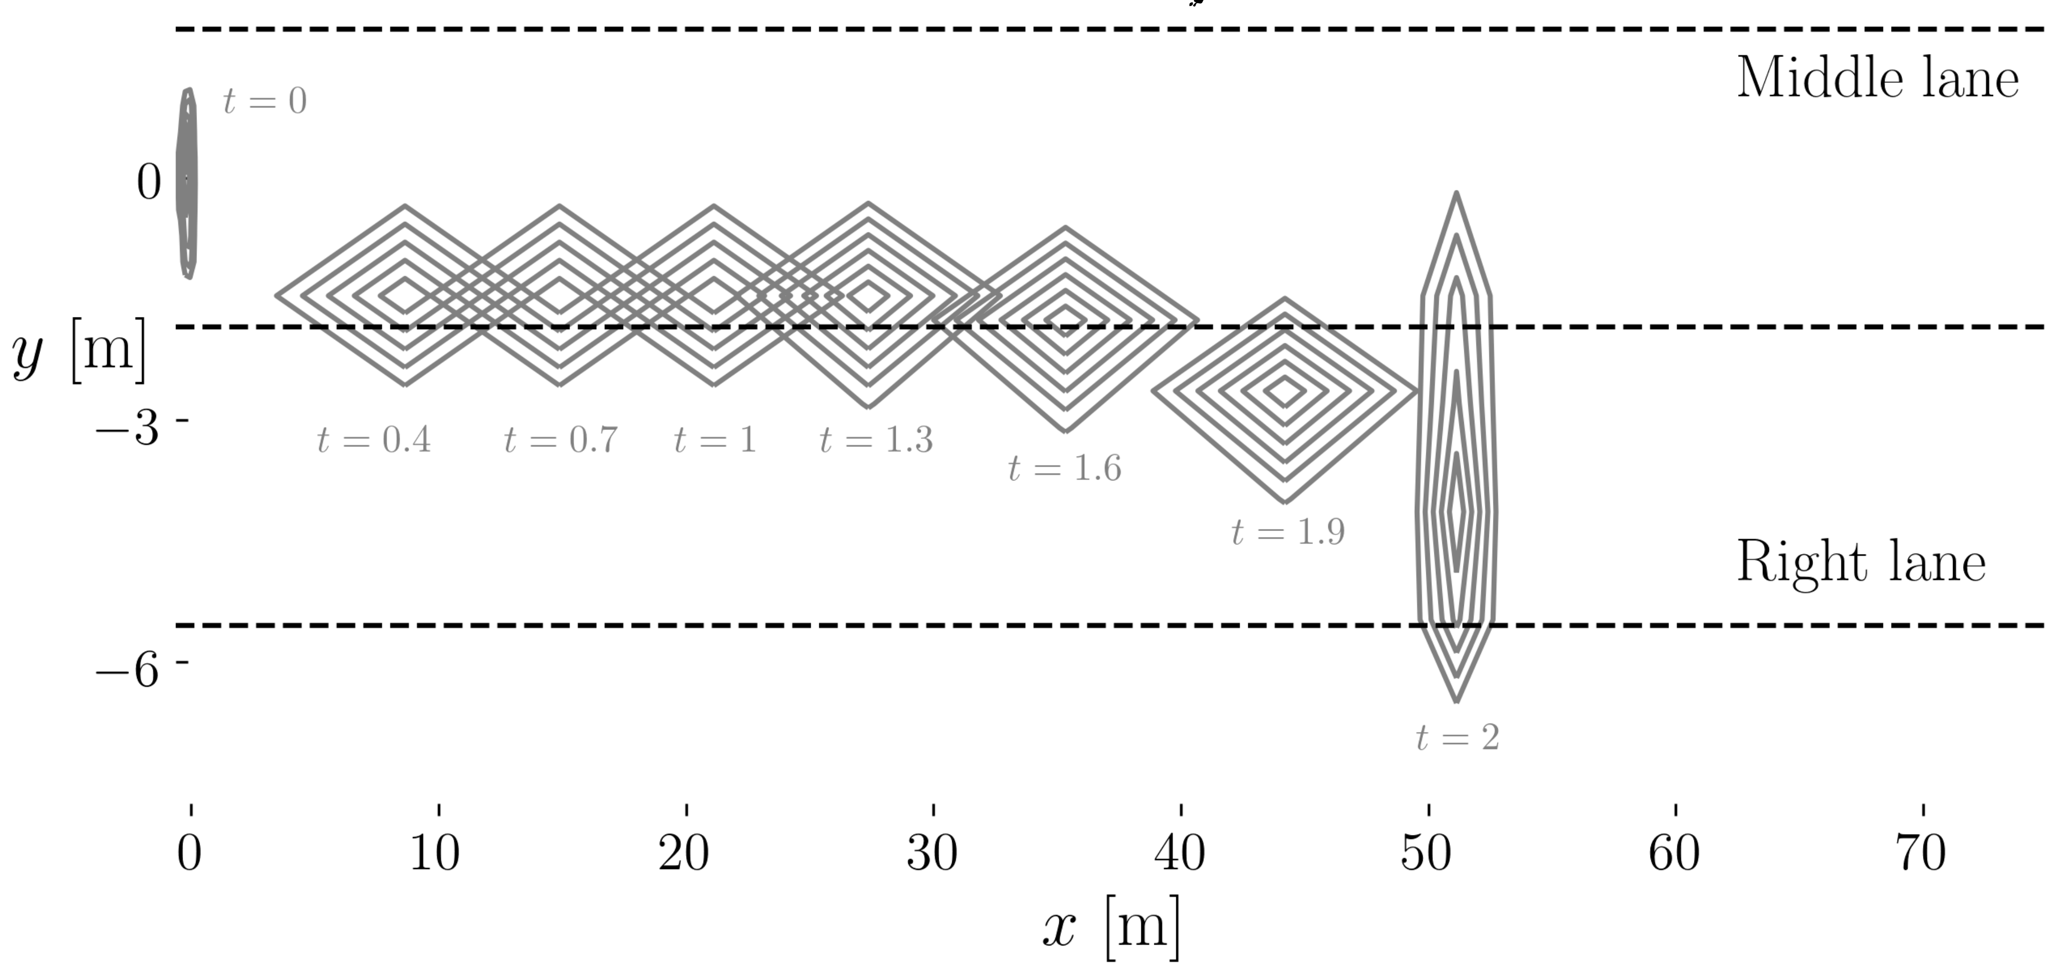

Hard Path Constraints: Reflected SBP

Main idea: path constraints $\sim$ reflected Itô SDEs
modify the controlled sample path dynamics to

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

$\boldsymbol{x}_t^{\boldsymbol{u}} \in \overline{\mathcal{X}} := \mathcal{X} \cup \partial\mathcal{X}$, closure of connected smooth $\mathcal{X}$

$\boldsymbol{n}$ is inward unit normal to the boundary $\partial\mathcal{X}$

γ_t is minimal local time stochastic process

Reflected SBP: Schrödinger Factor Recursion

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

$$\hat{\varphi}_0(\boldsymbol{x}) \longrightarrow \hat{\varphi}_1(\boldsymbol{x})$$

$$\rho_0(\boldsymbol{x}) \oslash \varphi_0(\boldsymbol{x}) \qquad\qquad \rho_1(\boldsymbol{x}) \oslash \hat{\varphi}_1(\boldsymbol{x})$$

$$\varphi_0(\boldsymbol{x}) \longleftarrow \varphi_1(\boldsymbol{x})$$

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

Reflected SBP: Numerics with Gradient Drift

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

Optimal controlled state PDFs:

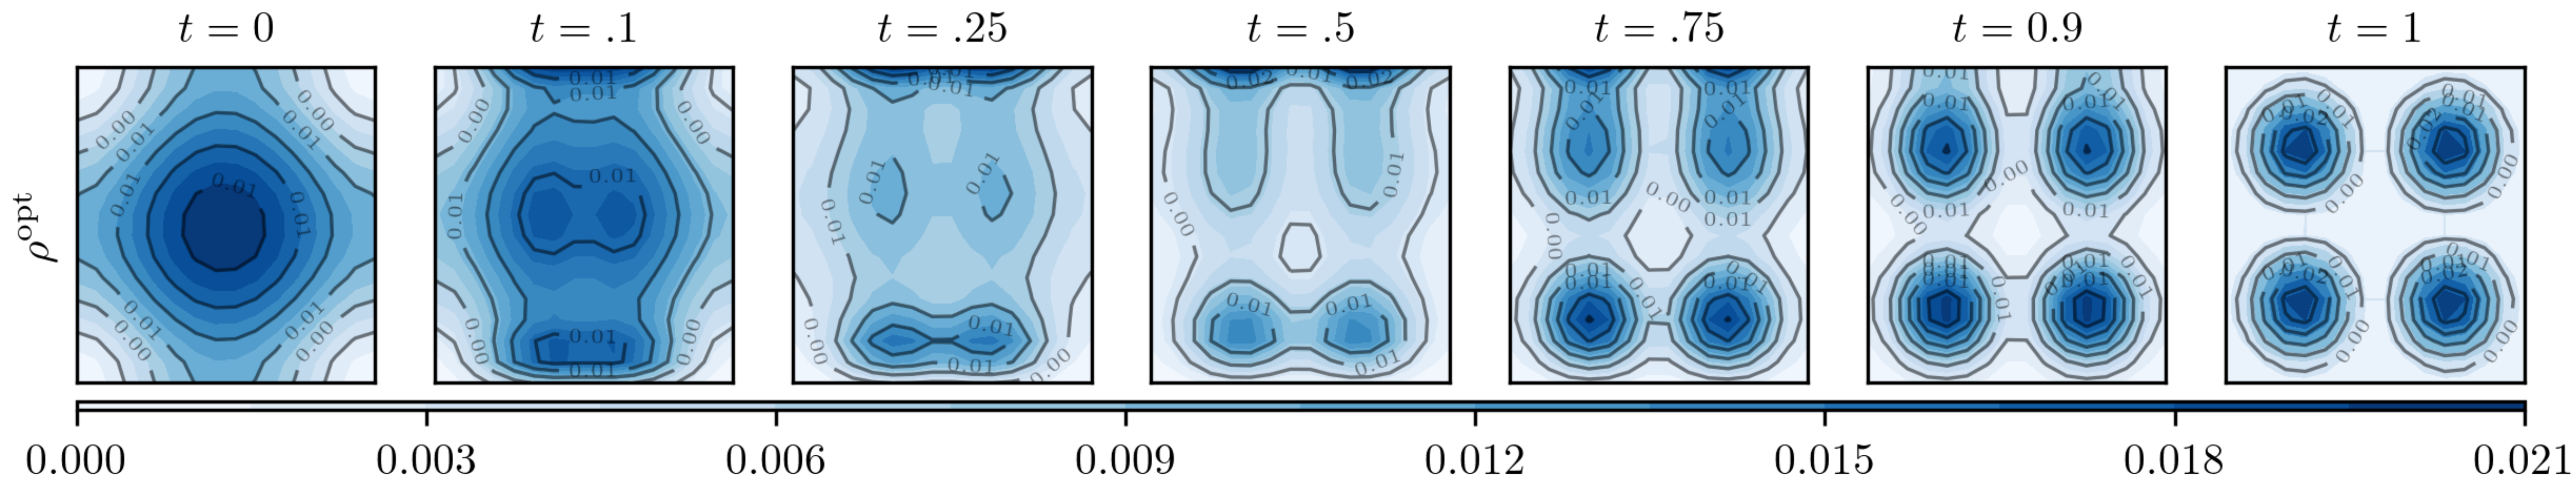

Uncontrolled state PDFs:

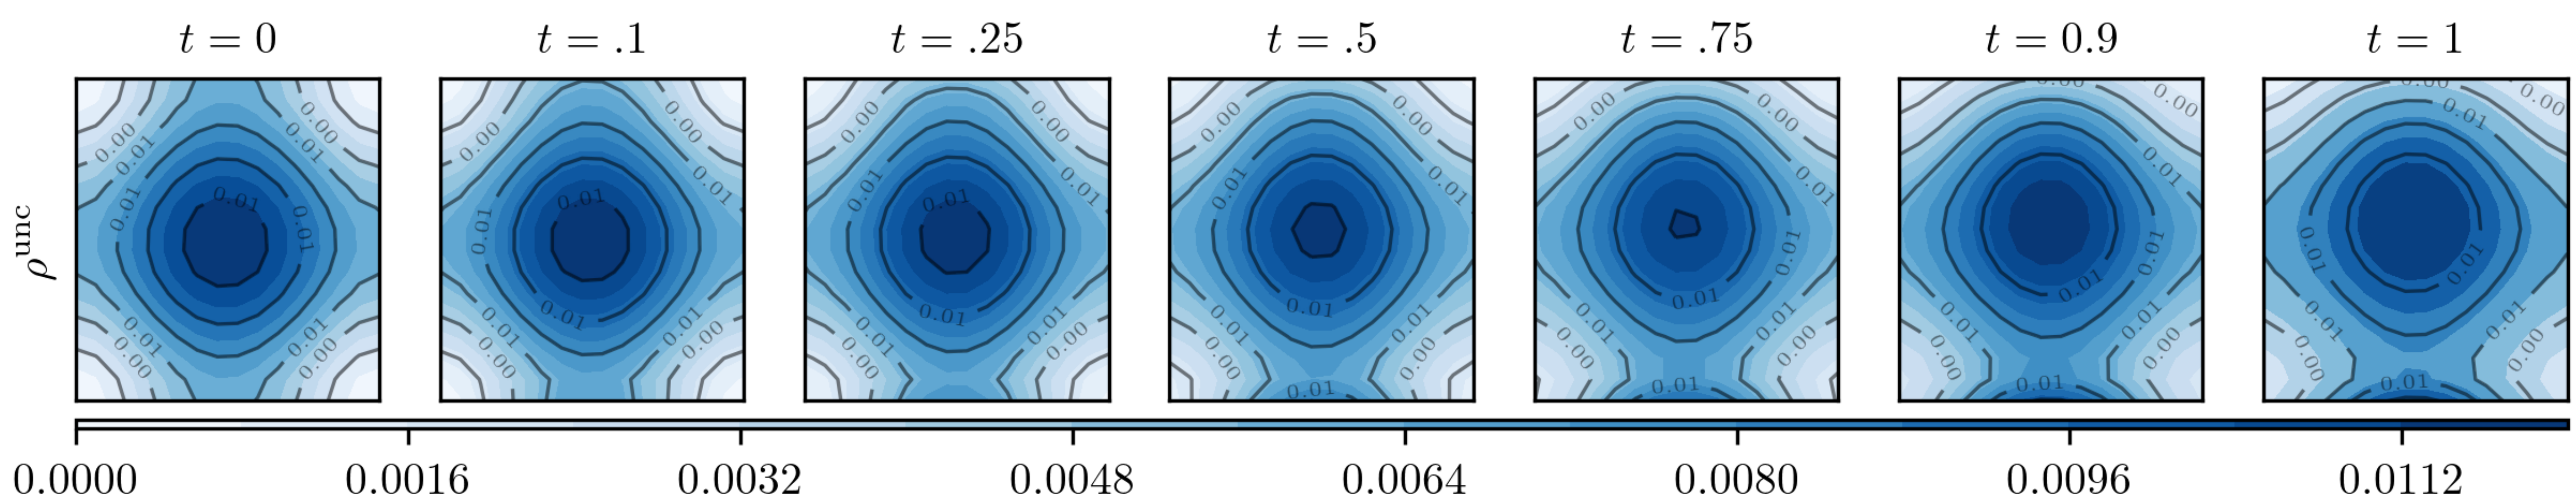

Control Non-affine SBP: Optimality Conditions

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

Drift coefficient   Diffusion tensor

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

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

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

$$\rho_{\rm opt}^{\boldsymbol{u}}(0,\boldsymbol{x}) = \rho_0, \qquad \rho_{\rm opt}^{\boldsymbol{u}}(T,\boldsymbol{x}) = \rho_T,$$

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

Stochastic Learning: Generalized SBP

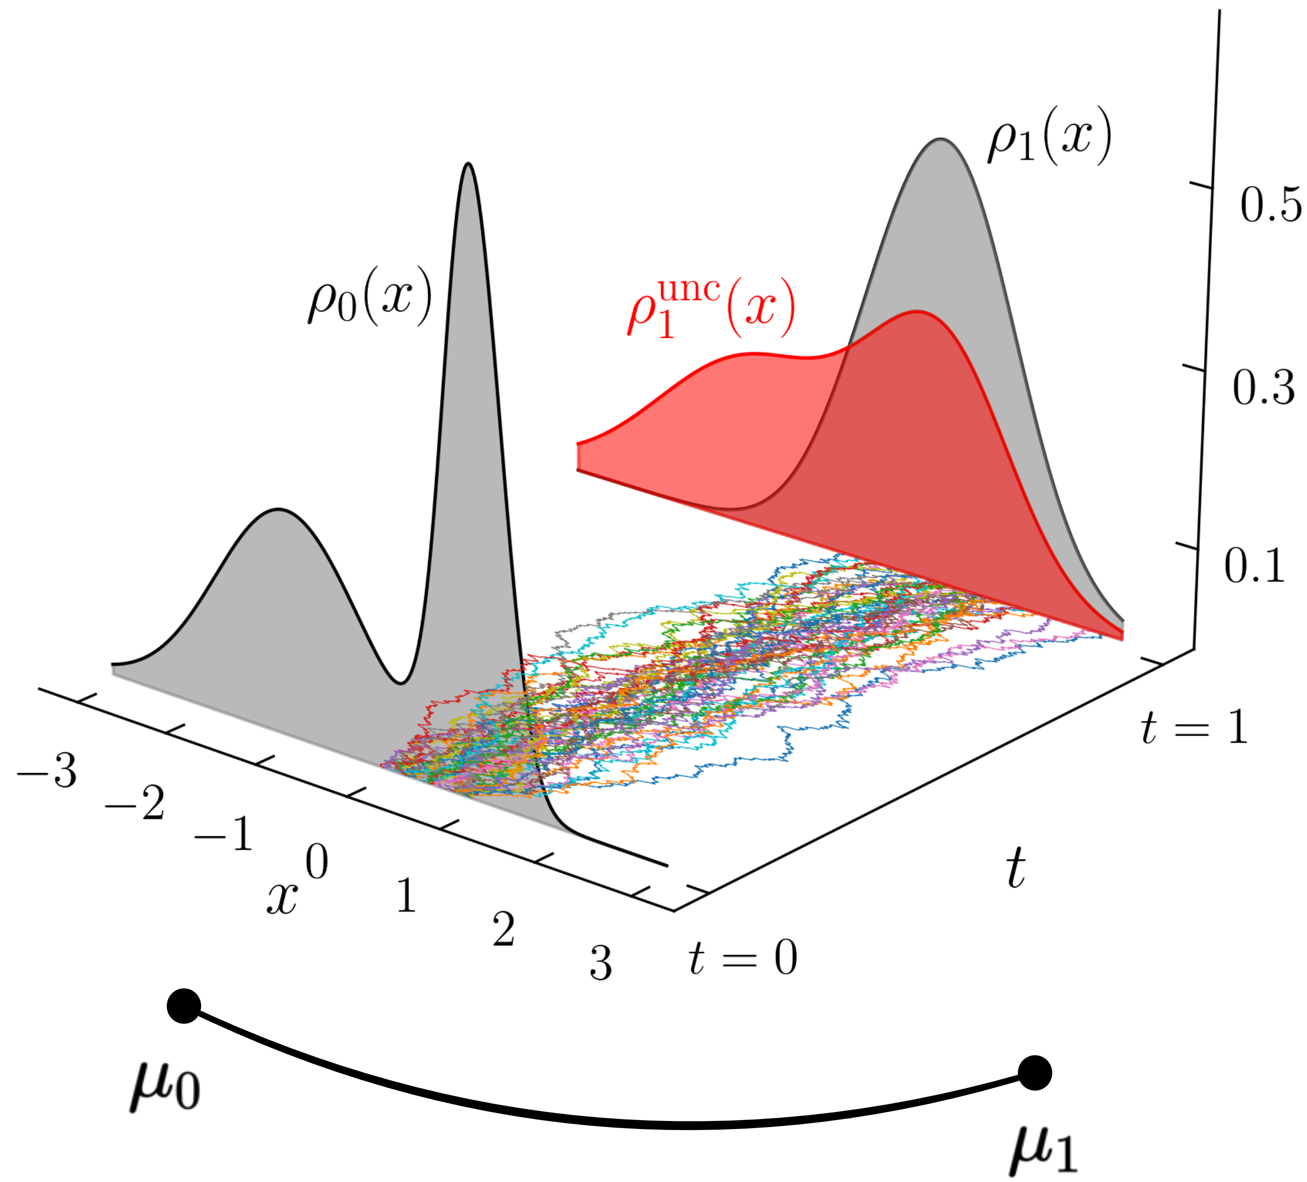

$$\mathcal{P}(\mathrm{AC}([0,1];\mathcal{P}_2(\mathcal{X})))$$

Large deviation principle on path measure

Motivating Applications

Learn most likely progression of medical condition

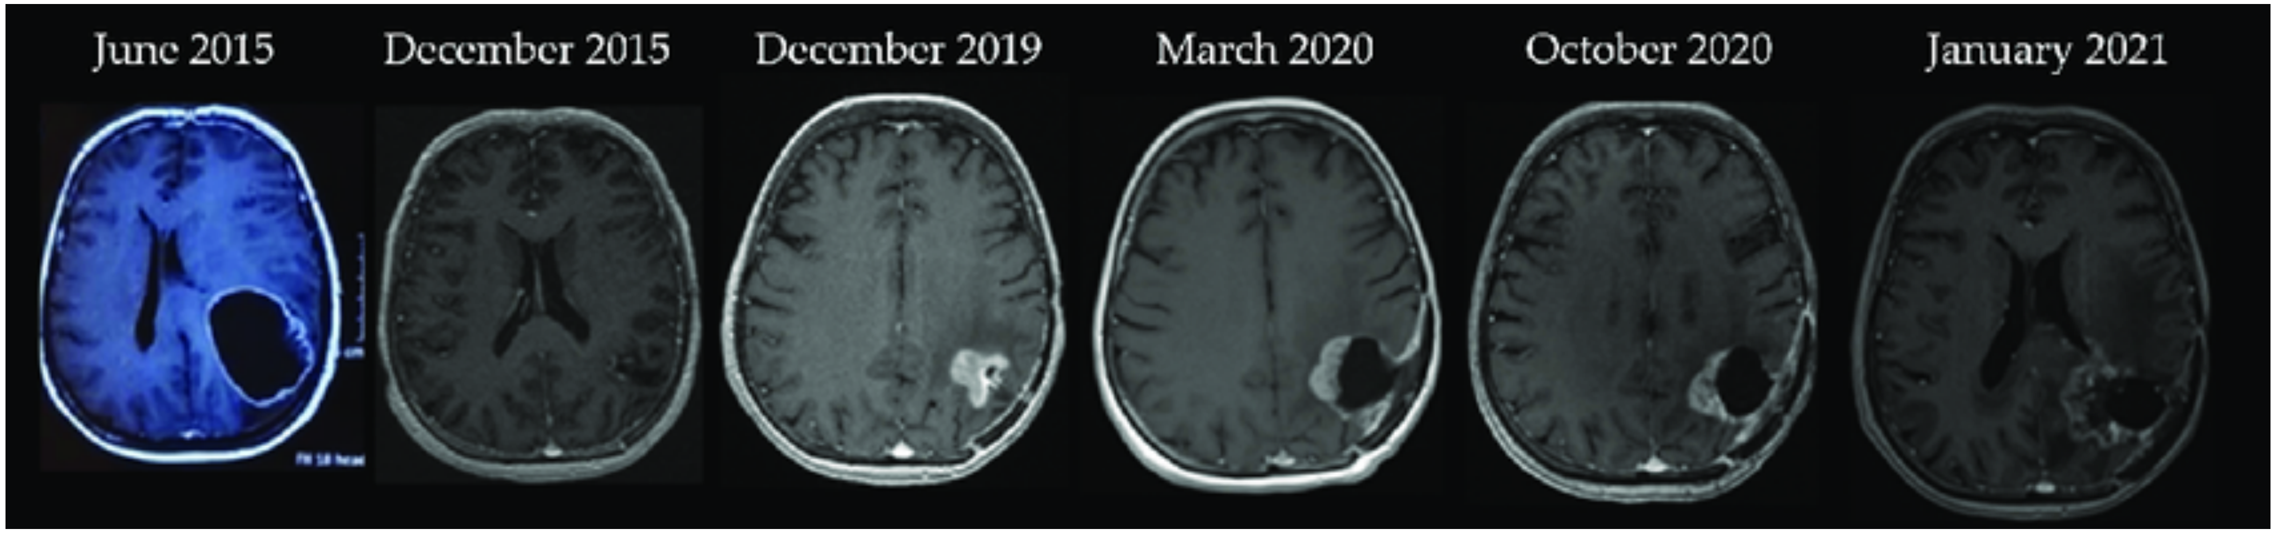

Learn joint stochastic time-varying hardware resource availability

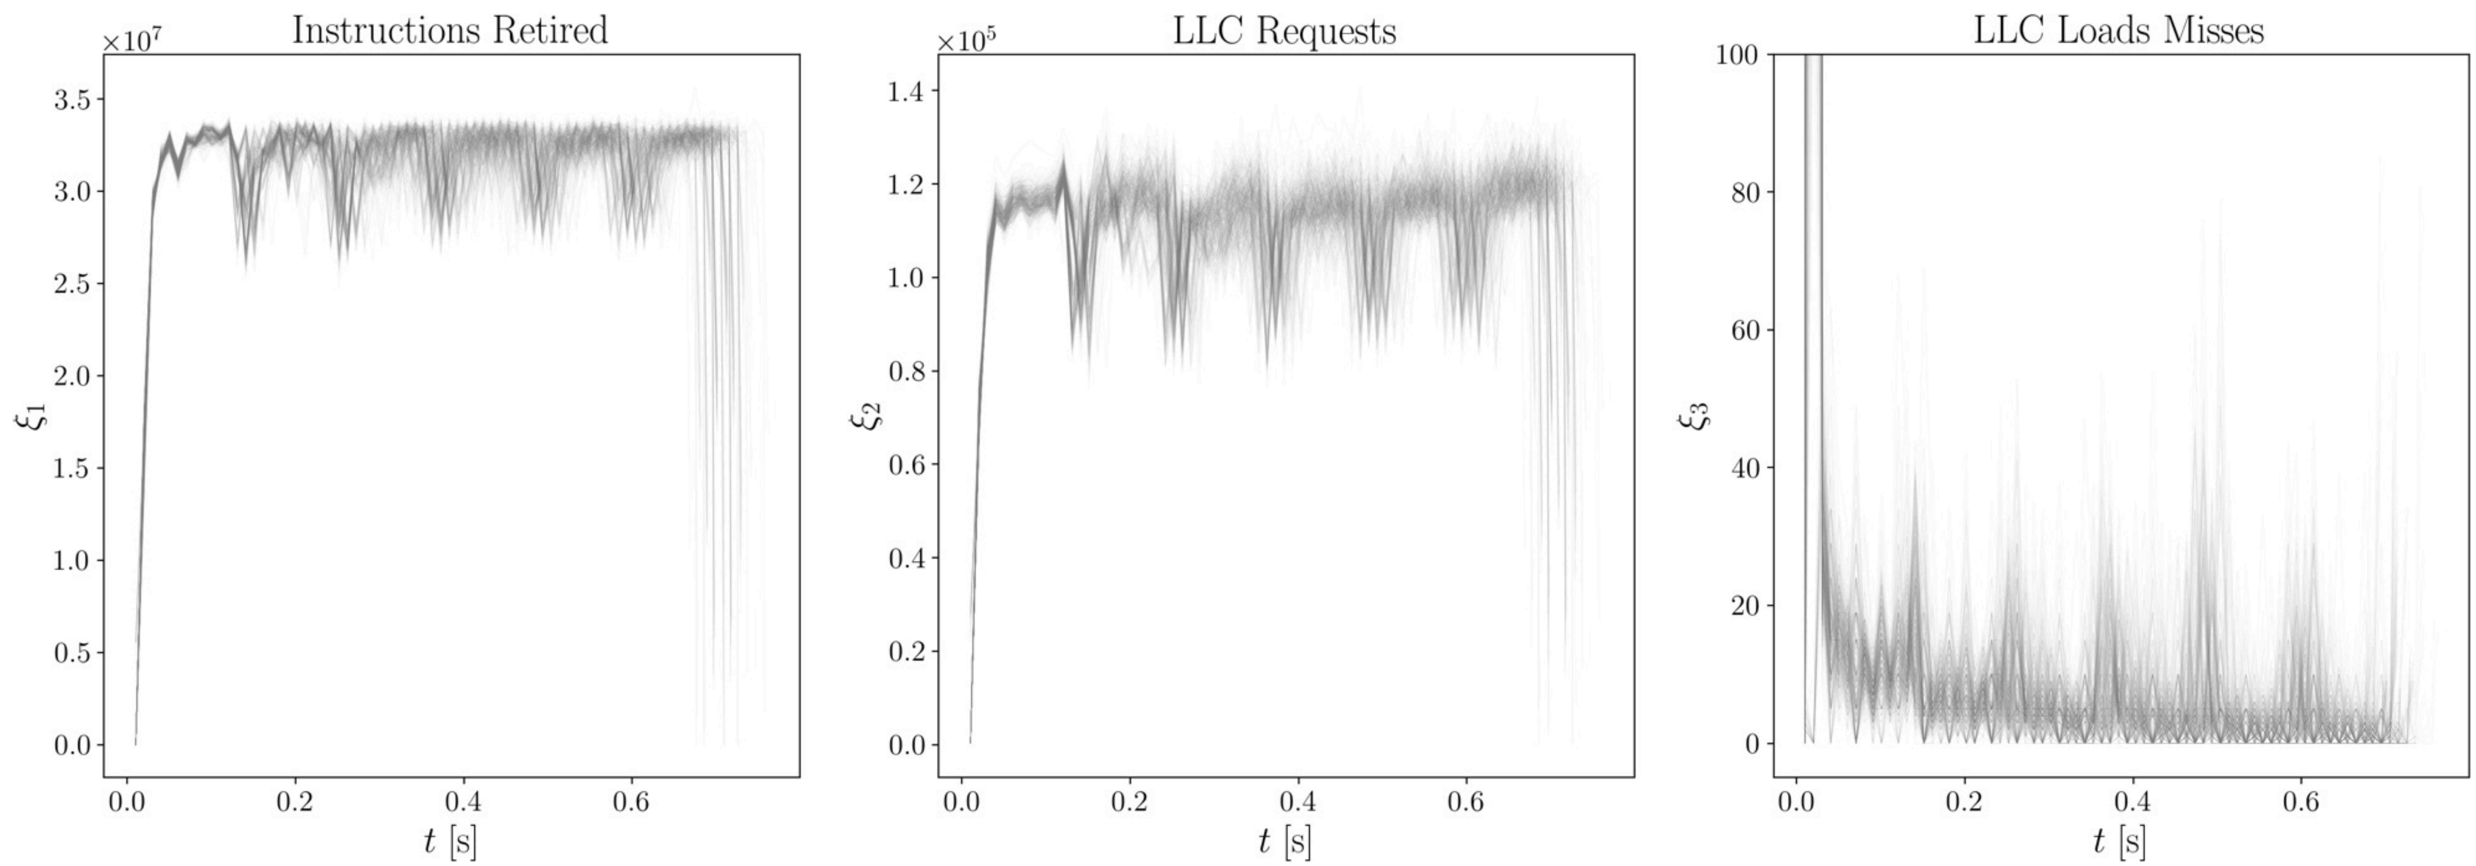

Learning-Control Duality

SBP as stochastic control problem:

$$\underset{(\rho,\boldsymbol{u})\in\mathcal{P}_{01}\times\mathcal{U}}{\arg\inf}\int_{t_0}^{t_1}\int_{\mathbb{R}^n}\left(\frac{1}{2}|\boldsymbol{u}|^2+q(\boldsymbol{x})\right)\rho(\boldsymbol{x},t)\mathrm{d}\boldsymbol{x}\mathrm{d}t$$

$$\frac{\partial\rho}{\partial t}+\nabla_{\boldsymbol{x}}\cdot(\rho\boldsymbol{u})=\varepsilon\Delta_{\boldsymbol{x}}\rho,$$

$$\boldsymbol{x}(t=t_0)\sim\rho_0\ (\text{given}),\quad \boldsymbol{x}(t=t_1)\sim\rho_1\ (\text{given})$$

SBP as large deviation principle on path space:

$$\underset{\mathbb{P}\in\Pi_{01}}{\arg\inf}\, D_{\mathrm{KL}}\left(\mathbb{P}\|\frac{\exp\left(-\frac{1}{2\varepsilon}\int_{t_0}^{t_1}q(\boldsymbol{x})\mathrm{d}t\right)\mathbb{W}}{Z}\right)$$

$$\Pi_{01}:=\{\mathbb{M}\in\mathcal{M}(\Omega)\mid\mathbb{M}\text{ has marginal }\rho_i\,\mathrm{d}\boldsymbol{x}\text{ at time }t_i\forall i\in\{0,1\},\rho_0,\rho_1\in\mathcal{P}_2(\mathbb{R}^n)\}$$

Generalization: Multimarginal SBP (MSBP)

Multi-marginal version: MSBP formulation

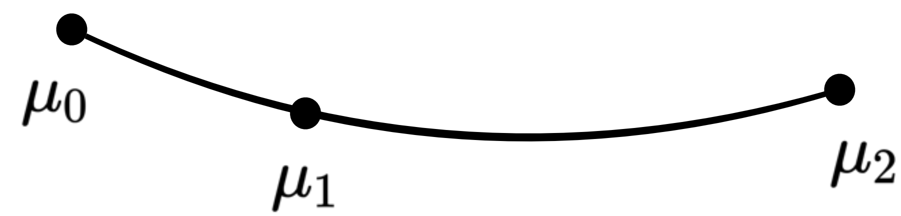

$\mathcal{X}_\sigma := \text{support}\,(\mu_\sigma) \subseteq \mathbb{R}^d \;\forall \sigma \in [\![s]\!],\qquad \mathcal{X}_1 \times \mathcal{X}_2 \times \ldots \times \mathcal{X}_s =: \boldsymbol{\mathcal{X}} \subseteq \left(\mathbb{R}^d\right)^{\otimes s}$

$\mathcal{M}(\mathcal{X}_\sigma)$ and $\mathcal{M}(\boldsymbol{\mathcal{X}})$ denote manifold of prob. measures on $\mathcal{X}_\sigma$ and $\boldsymbol{\mathcal{X}}$

Ground cost $\boldsymbol{C} : \boldsymbol{\mathcal{X}} \mapsto \mathbb{R}_{\geq 0}$

Let

$$\mathrm{d}\boldsymbol{\xi}_{-\sigma} := \mathrm{d}\boldsymbol{\xi}(\tau_1) \times \ldots \times \mathrm{d}\boldsymbol{\xi}(\tau_{\sigma-1}) \times \mathrm{d}\boldsymbol{\xi}(\tau_{\sigma+1}) \times \ldots \times \mathrm{d}\boldsymbol{\xi}(\tau_s)$$
$$\boldsymbol{\mathcal{X}}_{-\sigma} := \mathcal{X}_1 \times \ldots \times \mathcal{X}_{\sigma-1} \times \mathcal{X}_{\sigma+1} \times \ldots \times \mathcal{X}_s$$

MSBP:

$$\min_{\boldsymbol{M} \in \mathcal{M}(\boldsymbol{\mathcal{X}})} \int_{\boldsymbol{\mathcal{X}}} \{\boldsymbol{C}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s)) + \varepsilon \log \boldsymbol{M}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s))\} \boldsymbol{M}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s)) \mathrm{d}\boldsymbol{\xi}(\tau_1) \ldots \mathrm{d}\boldsymbol{\xi}(\tau_s)$$
$$\text{subject to} \int_{\boldsymbol{\mathcal{X}}_{-\sigma}} \boldsymbol{M}(\boldsymbol{\xi}(\tau_1), \ldots, \boldsymbol{\xi}(\tau_s)) \mathrm{d}\boldsymbol{\xi}_{-\sigma} = \mu_\sigma \forall \sigma \in [\![s]\!].$$

Large Deviation Interpretation for MSBP

Multimarginal Gibbs kernel $\boldsymbol{K}(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s))\mu_1 \otimes \ldots \otimes \mu_s$

$$\boldsymbol{K}(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s)) := \exp\left(-\frac{\boldsymbol{C}(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s))}{\varepsilon}\right)$$

Then MSBP is the same as

$$\min_{\pi\in\Pi(\mu_1,\ldots,\mu_s)} \varepsilon D_{\mathrm{KL}}\left(\pi \| \boldsymbol{K}\left(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s)\right)\mu_1 \otimes \ldots \otimes \mu_s\right)$$

Set of all path measures on $\mathcal{C}([\tau_1,\tau_s],\mathbb{R}^d)$ whose time τ_σ marginal is $\mu_\sigma\ \forall\sigma \in [\![s]\!]$

Discrete Formulation of MSBP

Ground cost is order s tensor $\boldsymbol{C} \in (\mathbb{R}^n)^{\otimes s}_{\geq 0}$, with components $[\boldsymbol{C}_{i_1,\ldots,i_s}] = \boldsymbol{C}(\boldsymbol{\xi}_{i_1}, \ldots, \boldsymbol{\xi}_{i_s})$.

Ditto for the discrete mass tensor $\boldsymbol{M} \in (\mathbb{R}^n)^{\otimes s}_{\geq 0}$

Define (marginalized) projection from nonneg tensor to nonneg vector:

$$\left[\mathrm{proj}_{\sigma}(\boldsymbol{M})_j\right] = \sum_{i_1,\ldots,i_{\sigma-1},i_{\sigma+1},\ldots,i_s} \boldsymbol{M}_{i_1,\ldots,i_{\sigma-1},j,i_{\sigma+1},\ldots,i_s}.$$

Discrete MSBP on scattered data:

$$\begin{aligned} &\min_{\boldsymbol{M}\in(\mathbb{R}^n)^{\otimes s}_{\geq 0}} \quad \langle \boldsymbol{C} + \varepsilon \log \boldsymbol{M}, \boldsymbol{M} \rangle \\ &\text{subject to} \quad \mathrm{proj}_{\sigma}(\boldsymbol{M}) = \boldsymbol{\mu}_{\sigma} \quad \forall \sigma \in [\![s]\!]. \end{aligned}$$

Strictly convex program in n^s decision variables

MSBP with Sequential Information Structure

Snapshot observation is a path tree: 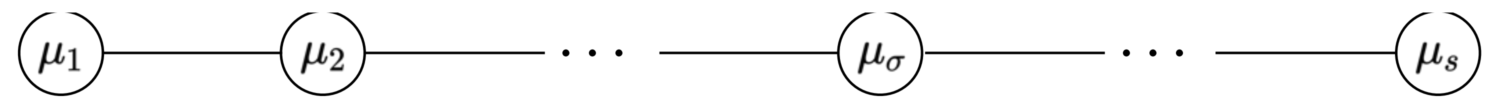

Ground cost admits path structure: $$\boldsymbol{C}(\boldsymbol{\xi}(\tau_1),\ldots,\boldsymbol{\xi}(\tau_s)) = \sum_{\sigma=1}^{s-1} c_\sigma\left(\boldsymbol{\xi}(\tau_\sigma),\boldsymbol{\xi}(\tau_{\sigma+1})\right).$$

KKT: $\boldsymbol{M}_{\text{opt}} = \boldsymbol{K} \odot \boldsymbol{U}$ where $\boldsymbol{K} := \exp(-\boldsymbol{C}/\varepsilon) \in (\mathbb{R}^n)^{\otimes s}_{>0}$, $\boldsymbol{U} := \otimes_{\sigma=1}^{s}\boldsymbol{u}_\sigma \in (\mathbb{R}^n)^{\otimes s}_{>0}$, $\boldsymbol{u}_\sigma := \exp(\boldsymbol{\lambda}_\sigma/\varepsilon)$

where $\boldsymbol{u}_\sigma$ solves multi marginal Sinkhorn contractive fixed point recursions:

$$\boldsymbol{u}_\sigma \leftarrow \boldsymbol{u}_\sigma \odot \boldsymbol{\mu}_\sigma \oslash \operatorname{proj}_\sigma(\boldsymbol{K}\odot\boldsymbol{U}) \quad \forall\sigma\in[\![s]\!]$$

But computing $\boldsymbol{K}\odot\boldsymbol{U}$ requires $\mathcal{O}\left(n^s\right)$ operations

From Exp to Linear Complexity

Thm.

$$\mathrm{proj}_\sigma(\boldsymbol{K} \odot \boldsymbol{U}) = \left(\boldsymbol{u}_1^\top K^{1\to 2} \prod_{j=2}^{\sigma-1} \mathrm{diag}(\boldsymbol{u}_j) K^{j\to j+1}\right)^\top \odot \boldsymbol{u}_\sigma \odot$$

$$\left(\left(\prod_{j=\sigma+1}^{s-1} K^{j-1\to j}\mathrm{diag}(\boldsymbol{u}_j)\right) K^{s-1\to s}\boldsymbol{u}_s\right) \quad \forall \sigma \in [\![s]\!],$$

Recursions become

$$\boldsymbol{u}_\sigma \leftarrow \boldsymbol{\mu}_\sigma \oslash \left(\left(\boldsymbol{u}_1^\top K^{1\to 2} \prod_{j=2}^{\sigma-1} \mathrm{diag}(\boldsymbol{u}_j) K^{j\to j+1}\right)^\top \right.$$

$$\left. \odot \left(\left(\prod_{j=\sigma+1}^{s-1} K^{j-1\to j}\mathrm{diag}(\boldsymbol{u}_j)\right) K^{s-1\to s}\boldsymbol{u}_s\right)\right) \quad \forall \sigma \in [\![s]\!].$$

Only $s-1$ matrix-vector multiplications: complexity $\mathcal{O}\left((s-1)n^2\right)$

Numerical Case Study: Convergence

$n = 500, s = 26:$ solving for ~ 1.49×10^{70} decision variables in ~ 10 s in MATLAB

Linear convergence of multimarginal Sinkhorn iterates in Hilbert's projective metric

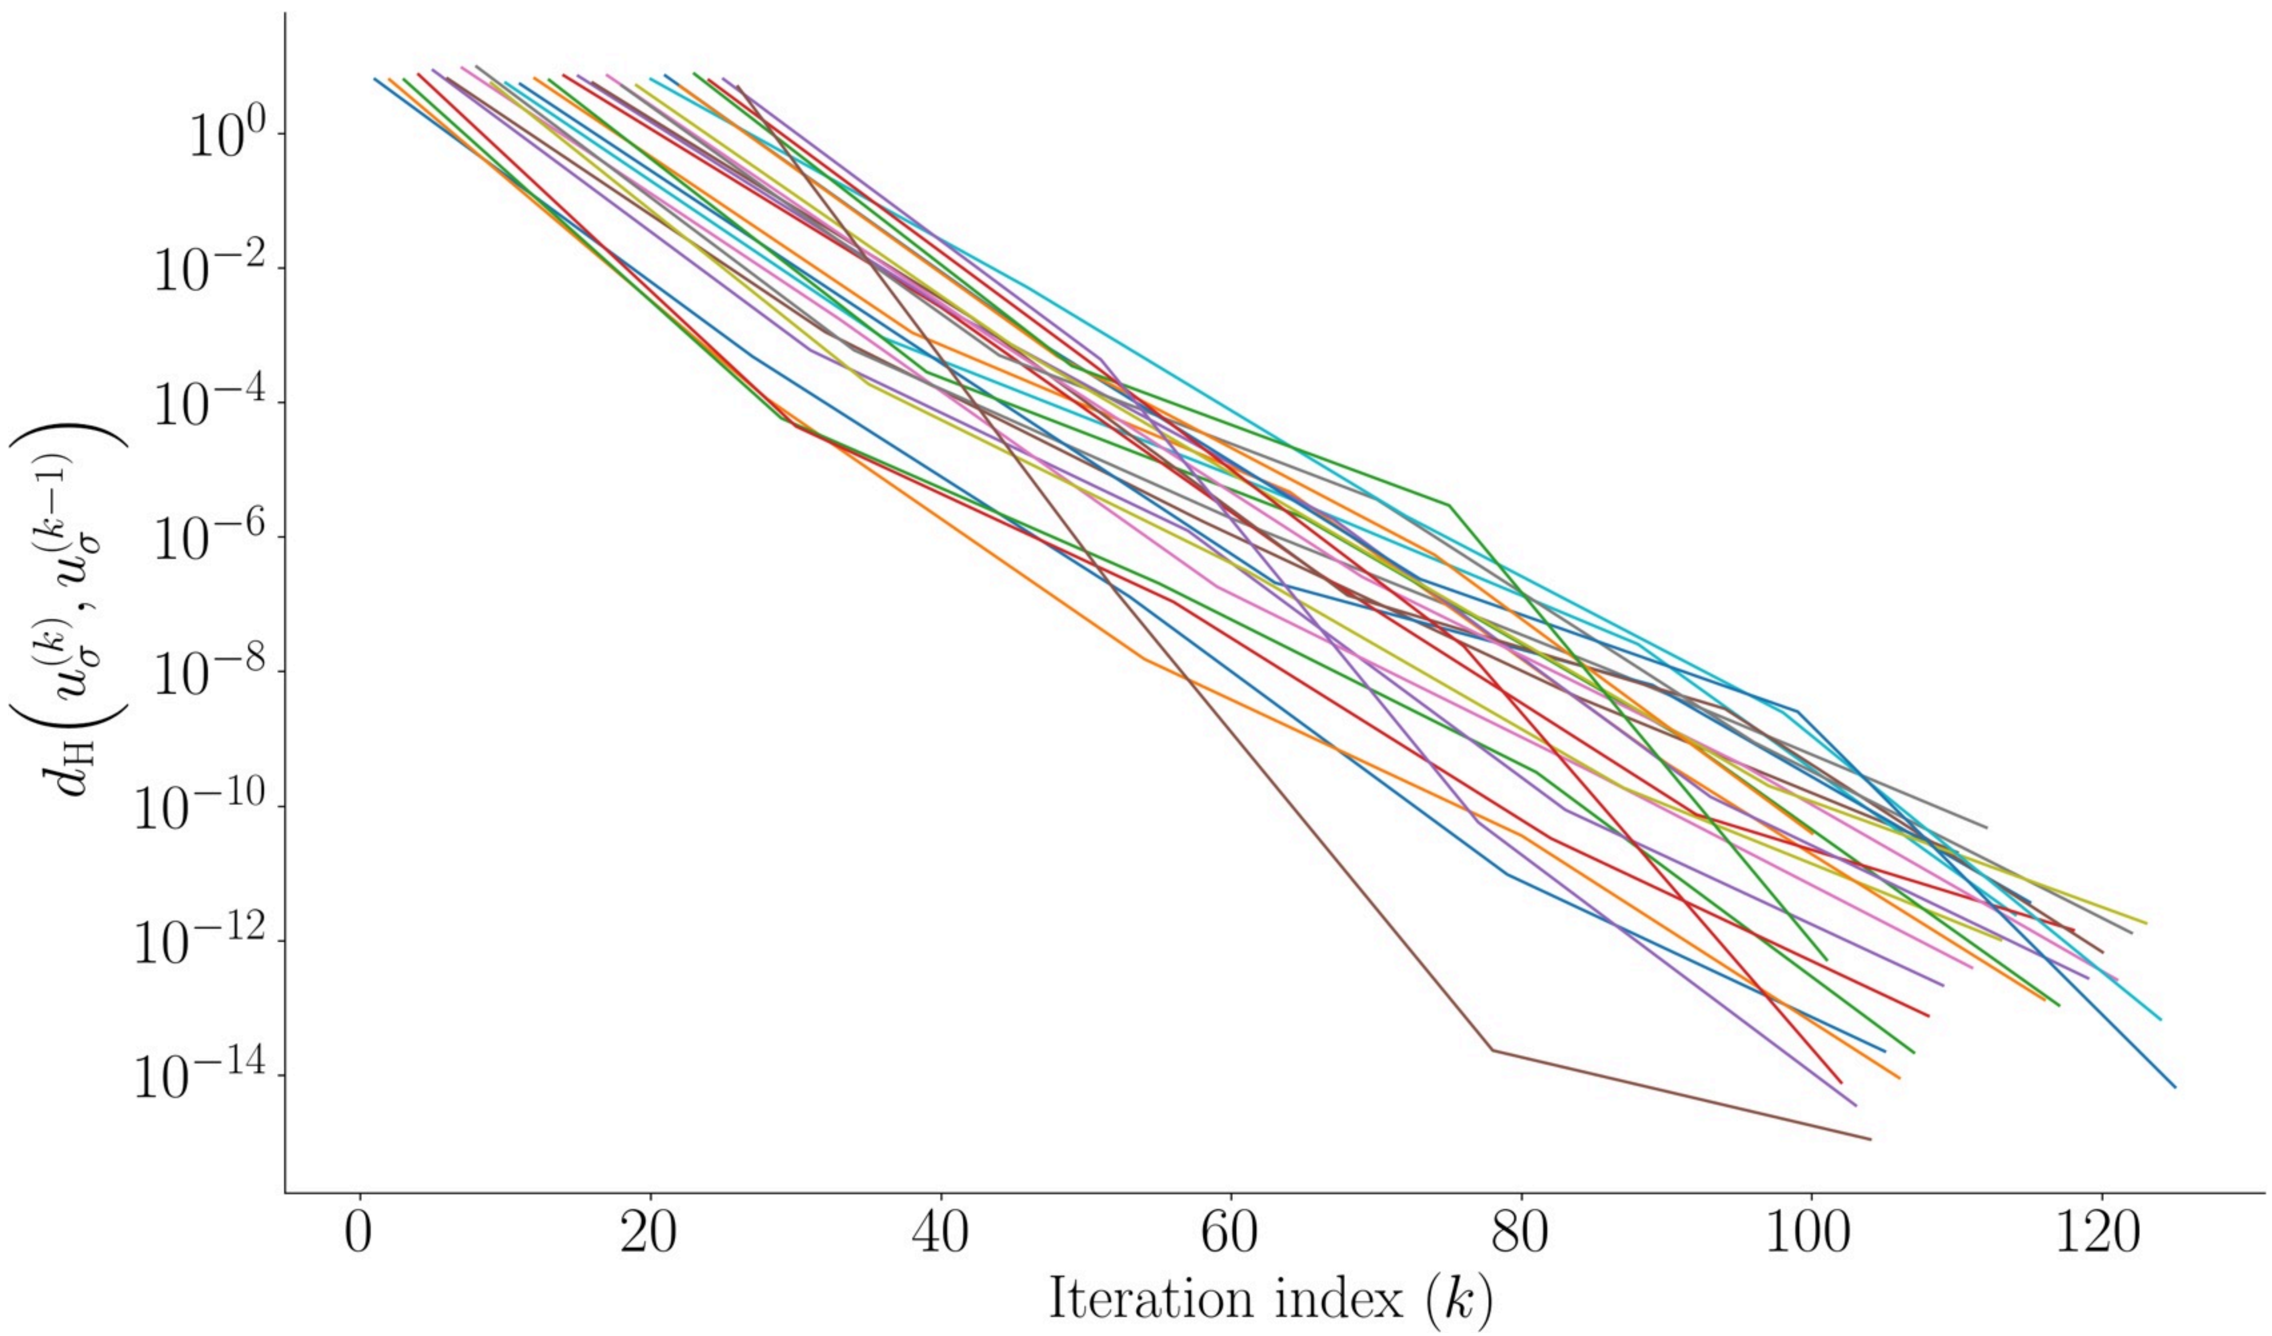

Numerical Case Study: Predicted vs Measured

Blue: predicted, red: measured, black: measured at control cycle boundaries

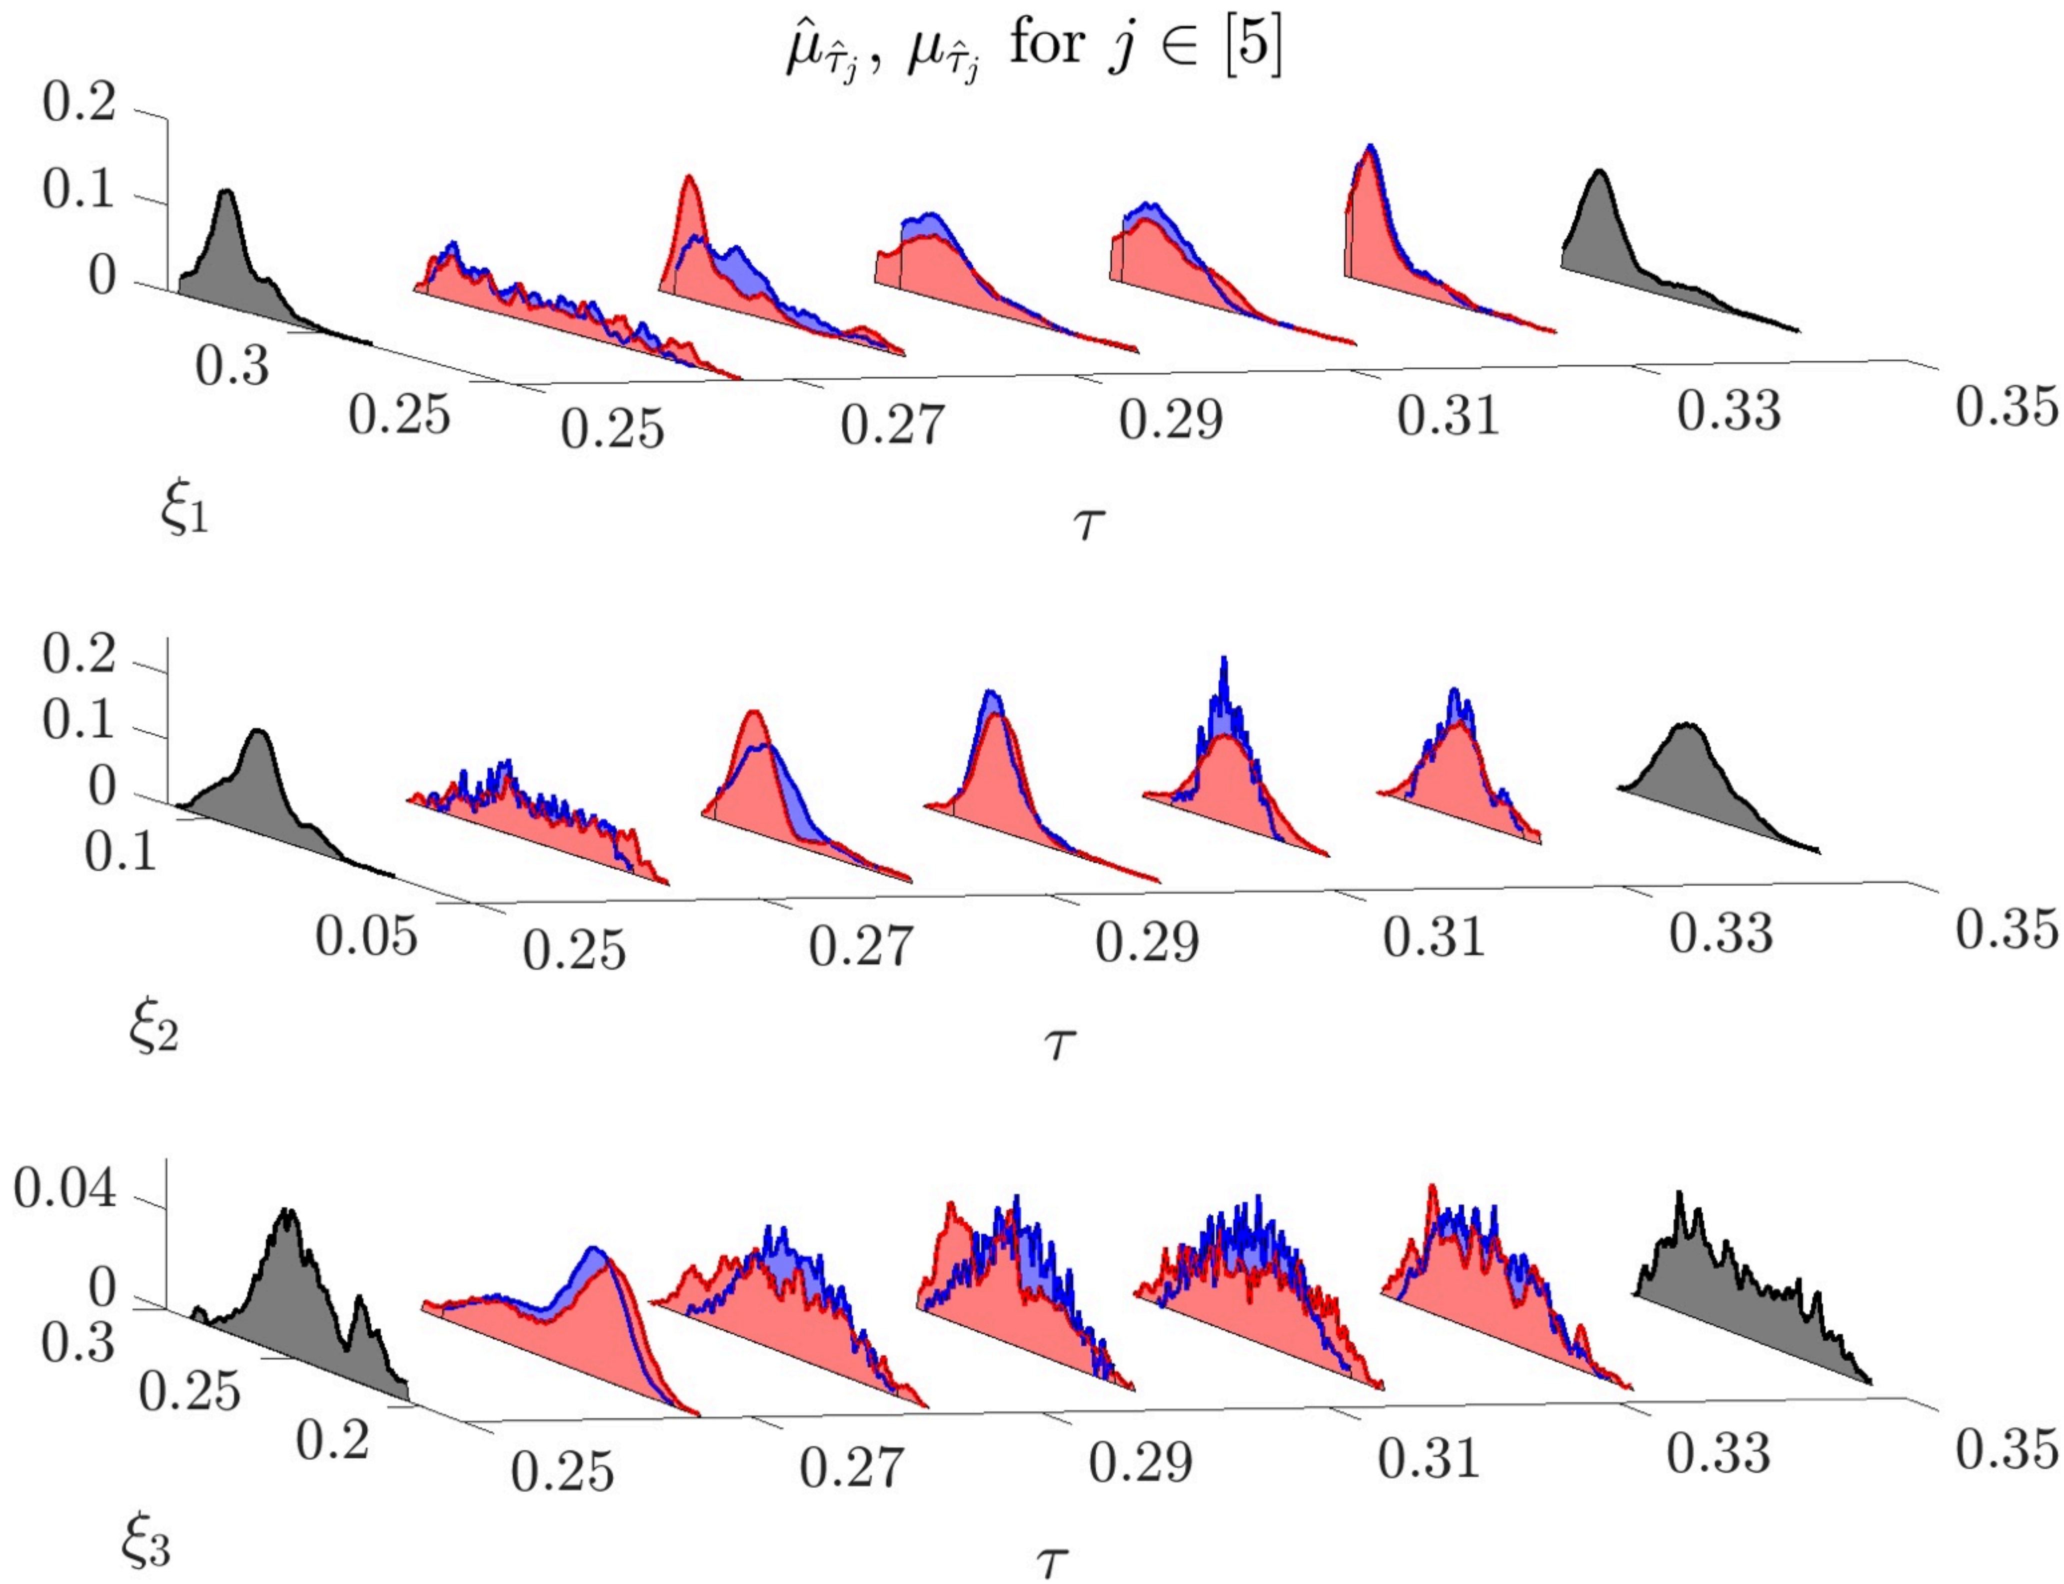

Outlook

- Density control and learning: undergoing rapid developments

- Lots of mathematics, algorithms and applications to be done

- Growing community in systems-control

- Strong intersections with many areas: probability, analysis, geometry, optimization, AI/ML, statistics, information theory, robotics, systems biology

We are hiring for Grad students and Postdocs

- Need strong mathematical background
- Interdisciplinary team
- Exciting projects at the intersection of control and ML

Thank You

Support: NSF, CITRIS PEOPLE AND ROBOTS, Ford