Generalized Gradient Flows for Stochastic Prediction, Estimation, Learning and Control

Abhishek Halder

Department of Aerospace Engineering
Iowa State University
Ames, IA 50011

Joint work with students and collaborators

       

Translational AI Center Seminar, Iowa State University
September 22, 2023

Topic of this talk

Control theory and algorithms for measures/distributions and densities

Why bother?

Prediction Problem

Compute joint state PDF $\rho\left(x,t\right)$

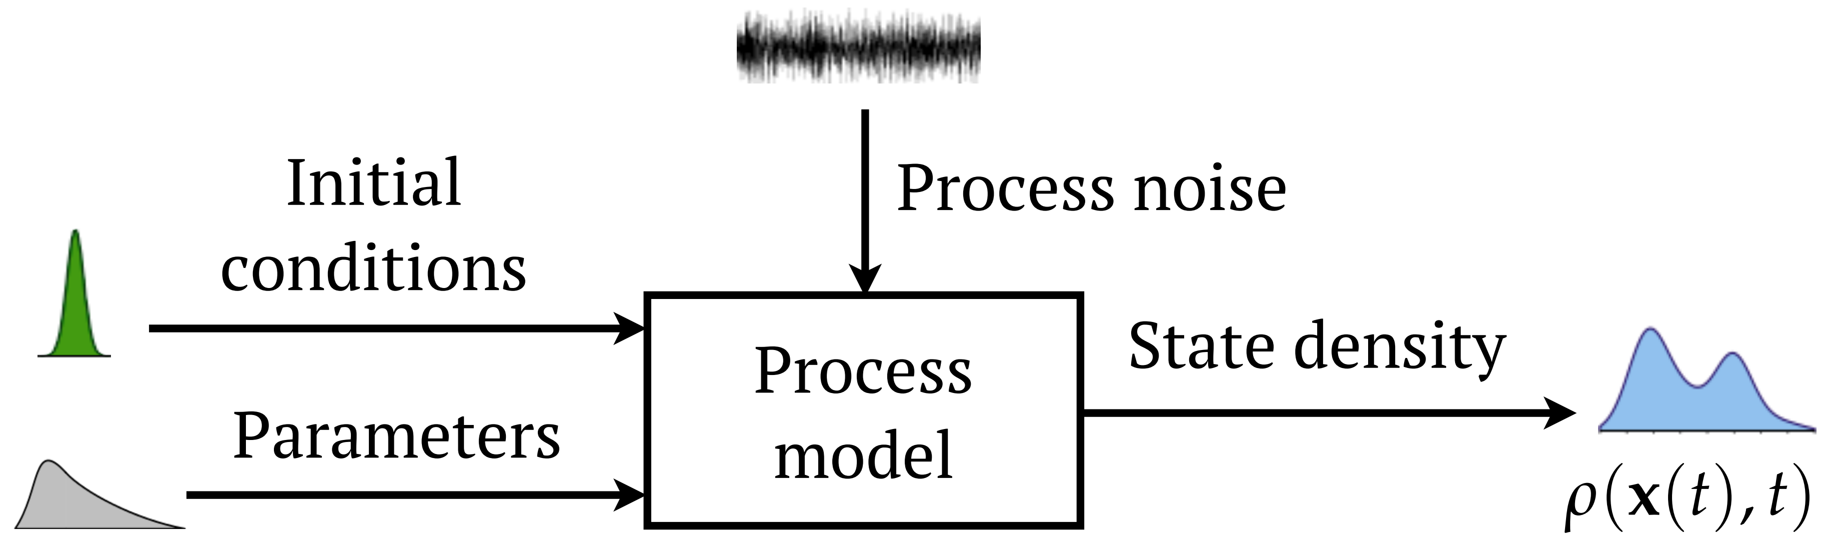

Trajectory flow:

$$\mathrm{d}\mathbf{X}(t) = \mathbf{f}(\mathbf{X},t)\,\mathrm{d}t + \mathbf{g}(\mathbf{X},t)\,\mathrm{d}\mathbf{w}(t), \quad \mathrm{d}\mathbf{w}(t) \sim \mathcal{N}(0,\mathbf{Q}\mathrm{d}t)$$

Density flow:

$$\frac{\partial \rho}{\partial t} = \mathcal{L}_{\mathrm{FP}}(\rho) := -\nabla\cdot(\rho\mathbf{f}) + \frac{1}{2}\sum_{i,j=1}^{n}\frac{\partial^2}{\partial x_i \partial x_j}\left(\left(\mathbf{g}\mathbf{Q}\mathbf{g}^\top\right)_{ij}\rho\right)$$

Filtering Problem

Compute conditional joint state PDF

$$\rho^{+} := \rho\left(x, t \mid z(s), 0 \le s \le t\right)$$

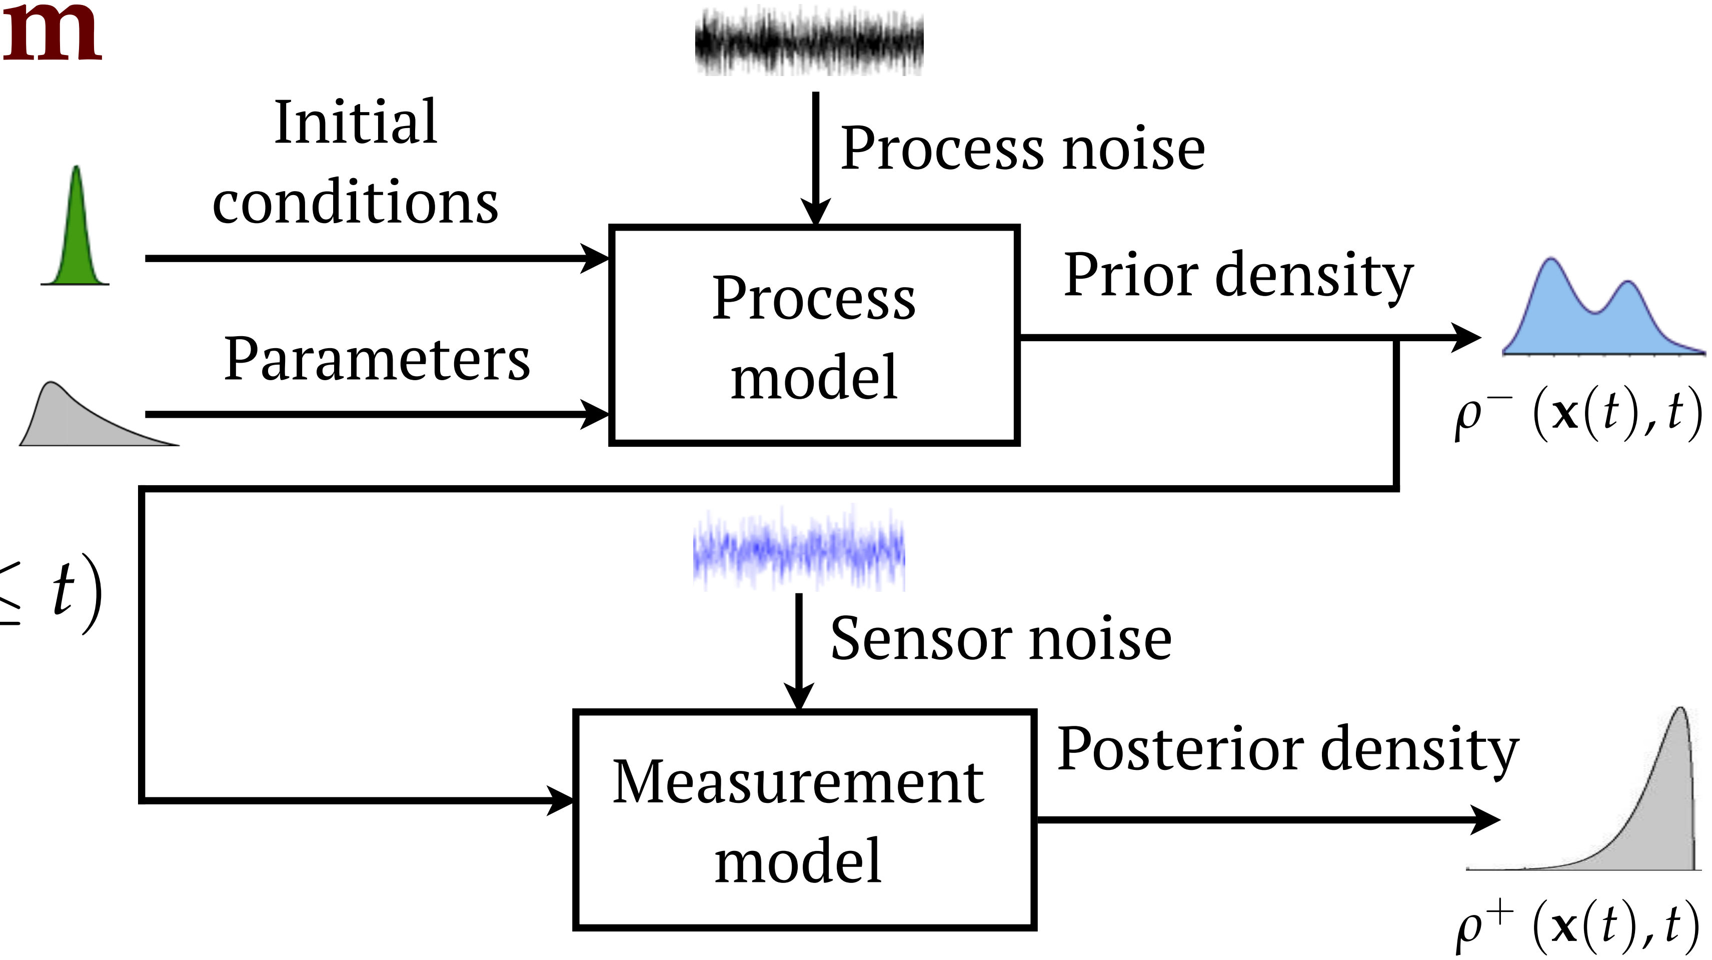

Trajectory flow:

$$\mathrm{d}\mathbf{X}(t) = \mathbf{f}(\mathbf{X}, t)\,\mathrm{d}t + \mathbf{g}(\mathbf{X}, t)\,\mathrm{d}\mathbf{w}(t), \qquad \mathrm{d}\mathbf{w}(t) \sim \mathcal{N}(0, \mathbf{Q}\mathrm{d}t)$$

$$\mathrm{d}\mathbf{Z}(t) = \mathbf{h}(\mathbf{X}, t)\,\mathrm{d}t + \mathrm{d}\mathbf{v}(t), \qquad \mathrm{d}\mathbf{v}(t) \sim \mathcal{N}(0, \mathbf{R}\mathrm{d}t)$$

Density flow:

$$\mathrm{d}\rho^{+} = \left[\mathcal{L}_{\mathrm{FP}}\mathrm{d}t + \left(\mathbf{h}(\mathbf{x}, t) - \mathbb{E}_{\rho^{+}}\{\mathbf{h}(\mathbf{x}, t)\}\right)^{\top}\mathbf{R}^{-1}\left(\mathrm{d}\mathbf{z}(t) - \mathbb{E}_{\rho^{+}}\{\mathbf{h}(\mathbf{x}, t)\}\mathrm{d}t\right)\right]\rho^{+}$$

Control Problem

Steer joint state PDF via feedback control over finite time horizon

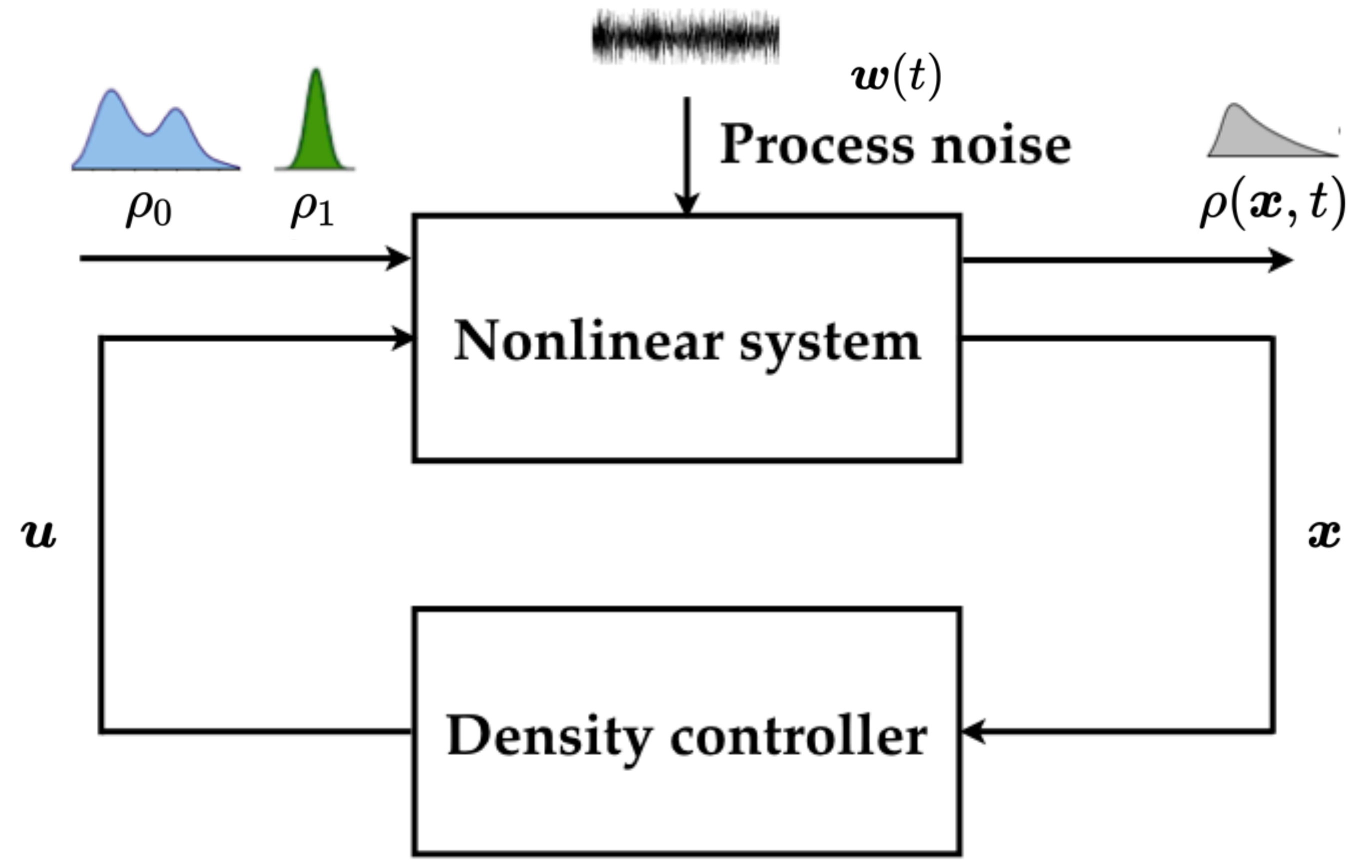

$$\underset{u \in \mathcal{U}}{\text{minimize}} \quad \mathbb{E}\left[\int_0^1 \|\boldsymbol{u}\|_2^2 \, \mathrm{d}t\right]$$

subject to

$$\mathrm{d}\boldsymbol{x} = \boldsymbol{f}\left(\boldsymbol{x}, \boldsymbol{u}, t\right) \, \mathrm{d}t + \boldsymbol{g}\left(\boldsymbol{x}, t\right) \, \mathrm{d}\boldsymbol{w},$$

$$\boldsymbol{x}\left(t = 0\right) \sim \rho_0, \quad \boldsymbol{x}\left(t = 1\right) \sim \rho_1$$

Mean Field Neural Network Learning Problem

Infinite width limit of fully connected NN:

Mei, Montanari and Nguyen, *Proceedings of the National Academy of Sciences*, 2018

Chizat and Bach, *NeurIPS*, 2018

Rotskoff and Vanden-Eijnden, *NeurIPS*, 2018

Sirignano and Spiliopoulos, *Stochastic Processes and their Applications*, 2020

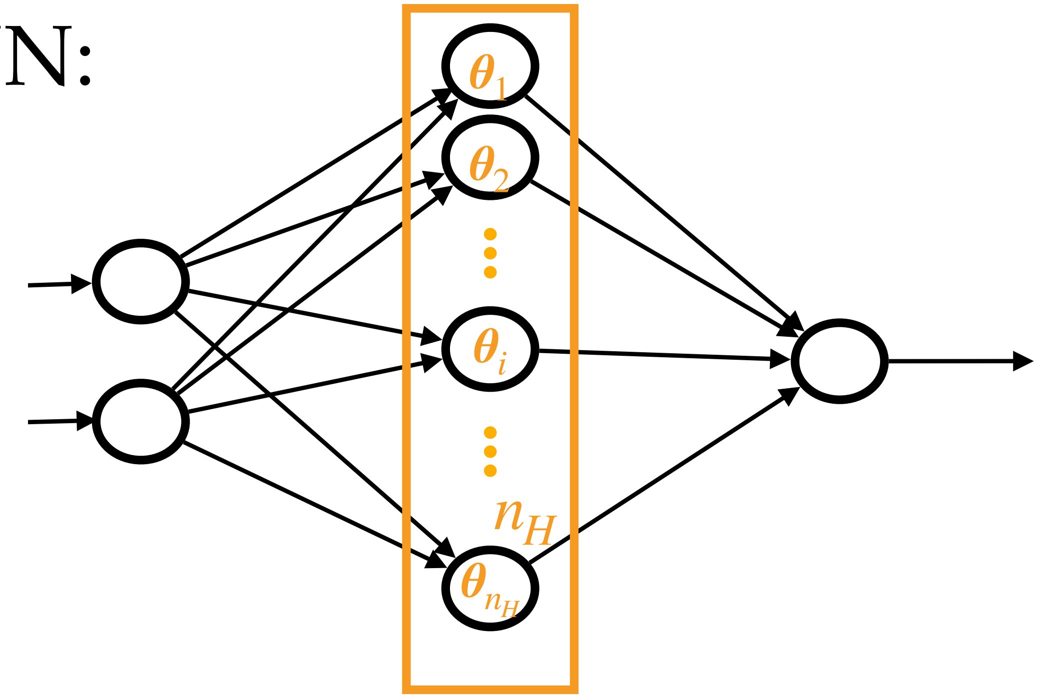

Mean field learning problem: $$\inf_{\rho \in \mathcal{P}_2(\mathbb{R}^p)} R\left(\int \Phi(\boldsymbol{x}, \boldsymbol{\theta})\rho(\boldsymbol{\theta})\mathrm{d}\boldsymbol{\theta}\right)$$

manifold of PDFs supported on $\mathbb{R}^p$ with finite second moments

PDF dynamics:

$$\frac{\partial \rho}{\partial t} = -\nabla^W R\left(\int \Phi\rho\right) = \nabla \cdot \left(\rho \nabla \frac{\delta}{\delta \rho} R\left(\int \Phi\rho\right)\right)$$

PDFs in Mars Entry-Descent-Landing

Prediction problem

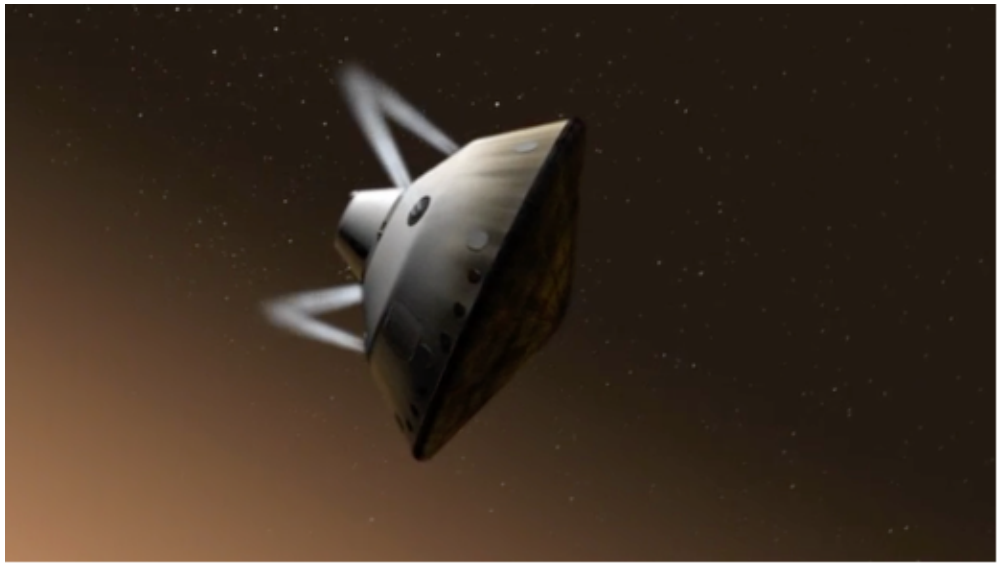

Predict heating rate uncertainty

Filtering problem

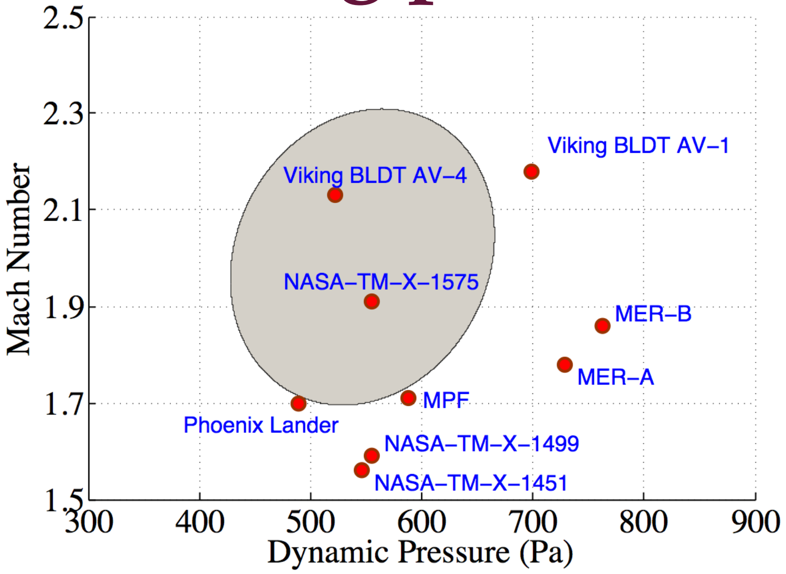

Estimate state to deploy parachute

Control problem

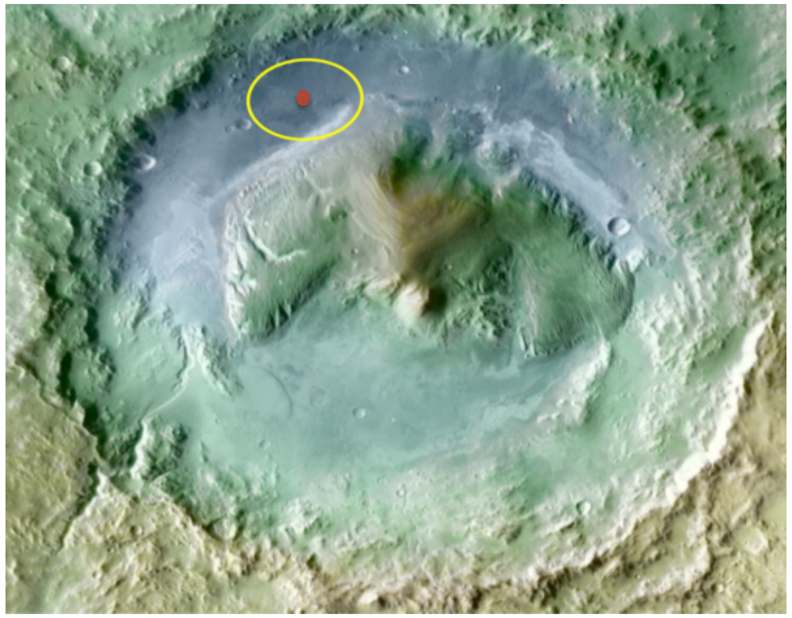

Gale Crater (4.49S, 137.42E)

Steer state PDF to achieve desired landing footprint accuracy

Learning problem

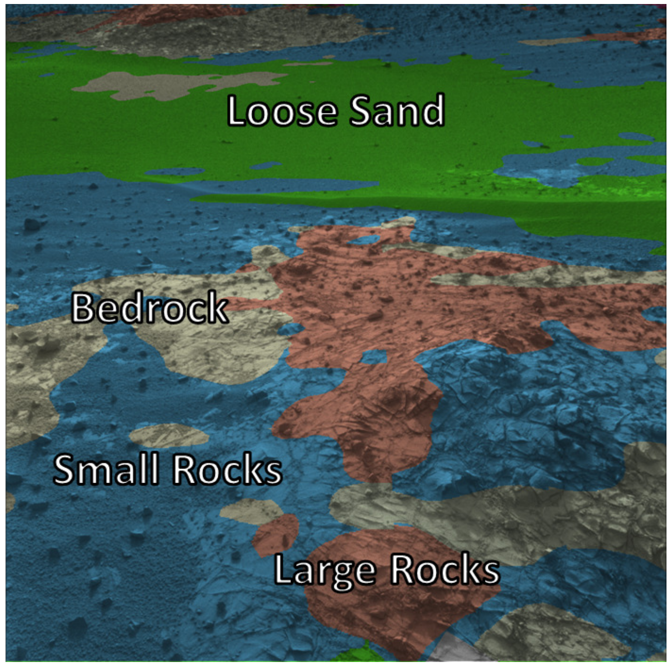

Learn surface feature from data

Solving prediction problem as Wasserstein gradient flow

What's New?

Main idea: Solve $\dfrac{\partial \rho}{\partial t} = \mathcal{L}_{\mathrm{FP}}\rho, \;\; \rho(x, t=0) = \rho_0$ **as gradient flow in** $\mathcal{P}_2(\mathcal{X})$

Infinite dimensional variational recursion:

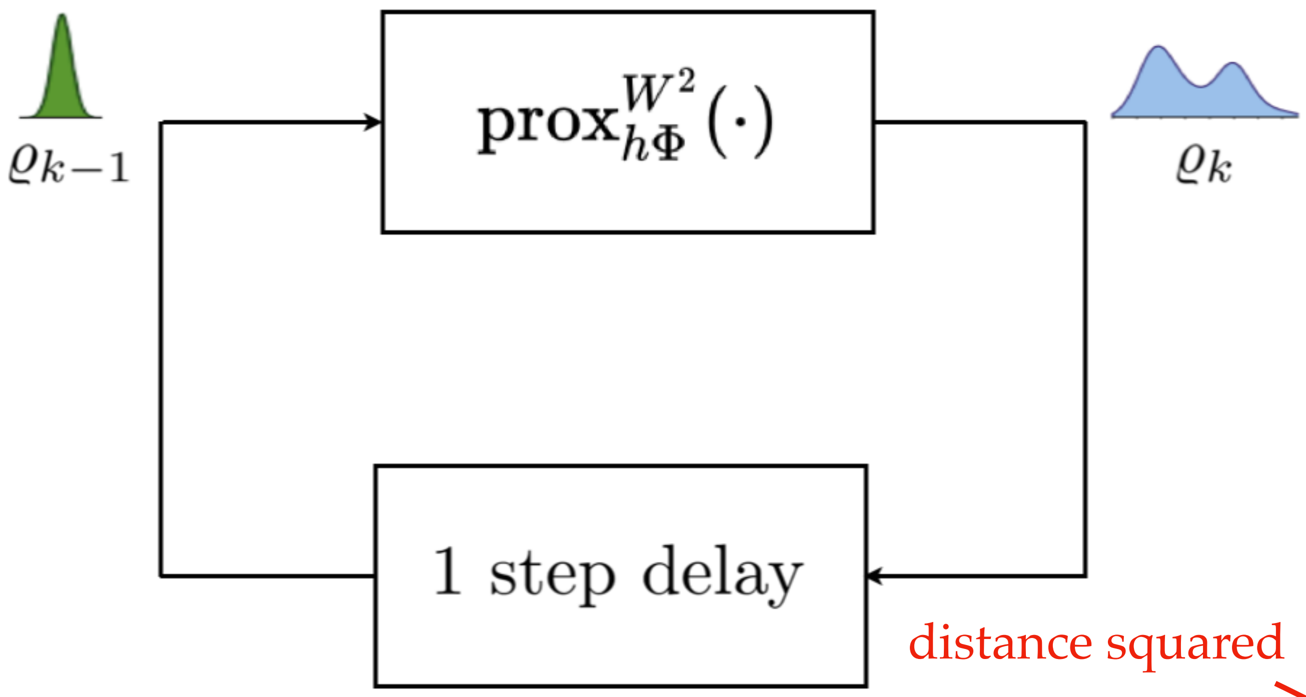

distance squared — time-step — energy-like functional

Proximal operator: $\varrho_k = \mathrm{prox}^{W^2}_{h\Phi}(\varrho_{k-1}) := \underset{\varrho \in \mathcal{P}_2(\mathcal{X})}{\arg\inf} \left\{ \frac{1}{2} W^2(\varrho, \varrho_{k-1}) + h\Phi(\varrho) \right\}$

Optimal transport cost: $W^2(\varrho, \varrho_{k-1}) := \inf_{\pi \in \Pi(\varrho, \varrho_{k-1})} \int_{\mathcal{X} \times \mathcal{X}} c(x, y)\, \mathrm{d}\pi(x, y)$

Free energy functional: $\Phi(\varrho) := \int_{\mathcal{X}} \psi \varrho \, \mathrm{d}x + \beta^{-1} \int_{\mathcal{X}} \varrho \log \varrho \, \mathrm{d}x$

Geometric Meaning of Gradient Flow

Gradient Flow in $\mathcal{X}$	Gradient Flow in $\mathcal{P}_2(\mathcal{X})$
$\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}t} = -\nabla\varphi(\boldsymbol{x}), \quad \boldsymbol{x}(0) = \boldsymbol{x}_0$	$\frac{\partial\rho}{\partial t} = -\nabla^W \Phi(\rho), \quad \rho(\boldsymbol{x}, 0) = \rho_0$
Recursion:	**Recursion:**
$\boldsymbol{x}_k = \boldsymbol{x}_{k-1} - h\nabla\varphi(\boldsymbol{x}_k) = \underset{\boldsymbol{x}\in\mathcal{X}}{\arg\min} \left\{\frac{1}{2}\Vert\boldsymbol{x} - \boldsymbol{x}_{k-1}\Vert_2^2 + h\varphi(\boldsymbol{x})\right\} =: \mathrm{prox}_{h\varphi}^{\Vert\cdot\Vert_2}(\boldsymbol{x}_{k-1})$	$\rho_k = \rho(\cdot, t = kh) = \underset{\rho\in\mathcal{P}_2(\mathcal{X})}{\arg\min} \left\{\frac{1}{2}W^2(\rho, \rho_{k-1}) + h\Phi(\rho)\right\} =: \mathrm{prox}_{h\Phi}^{W^2}(\rho_{k-1})$
Convergence:	**Convergence:**
$\boldsymbol{x}_k \to \boldsymbol{x}(t = kh) \quad \text{as} \quad h \downarrow 0$	$\rho_k \to \rho(\cdot, t = kh) \quad \text{as} \quad h \downarrow 0$
φ **as Lyapunov function:**	Φ **as Lyapunov functional:**
$\frac{\mathrm{d}}{\mathrm{d}t}\varphi = -\Vert\nabla\varphi\Vert_2^2 \leq 0$	$\frac{\mathrm{d}}{\mathrm{d}t}\Phi = -\mathbb{E}_\rho\left[\left\Vert\nabla\frac{\delta\Phi}{\delta\rho}\right\Vert_2^2\right] \leq 0$

Geometric Meaning of Gradient Flow

Gradient Flow in $\mathcal{X}$

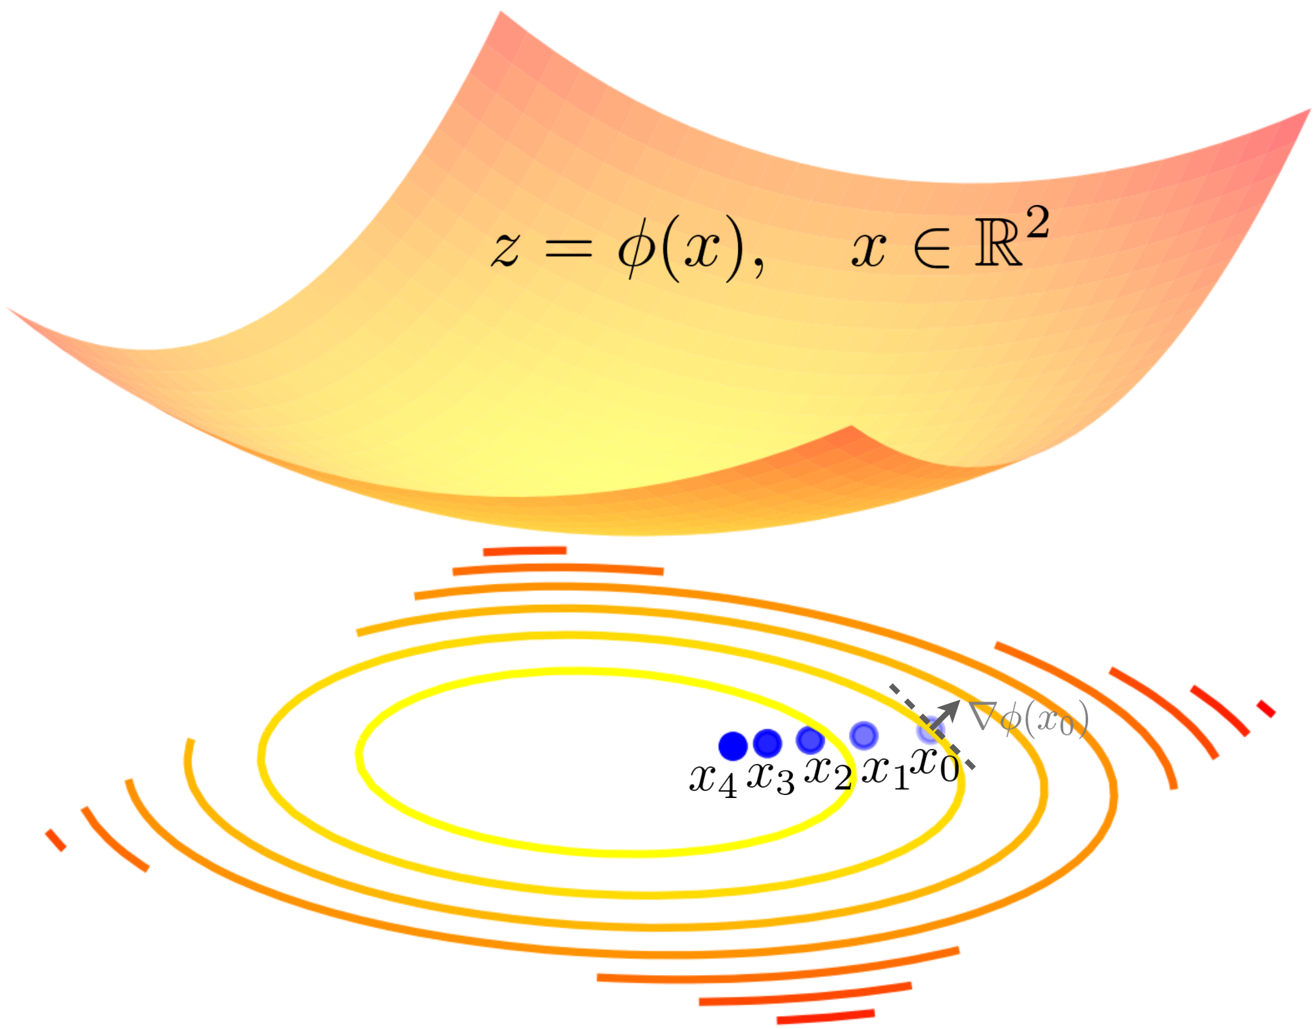

Gradient Flow in $\mathcal{P}_2(\mathcal{X})$

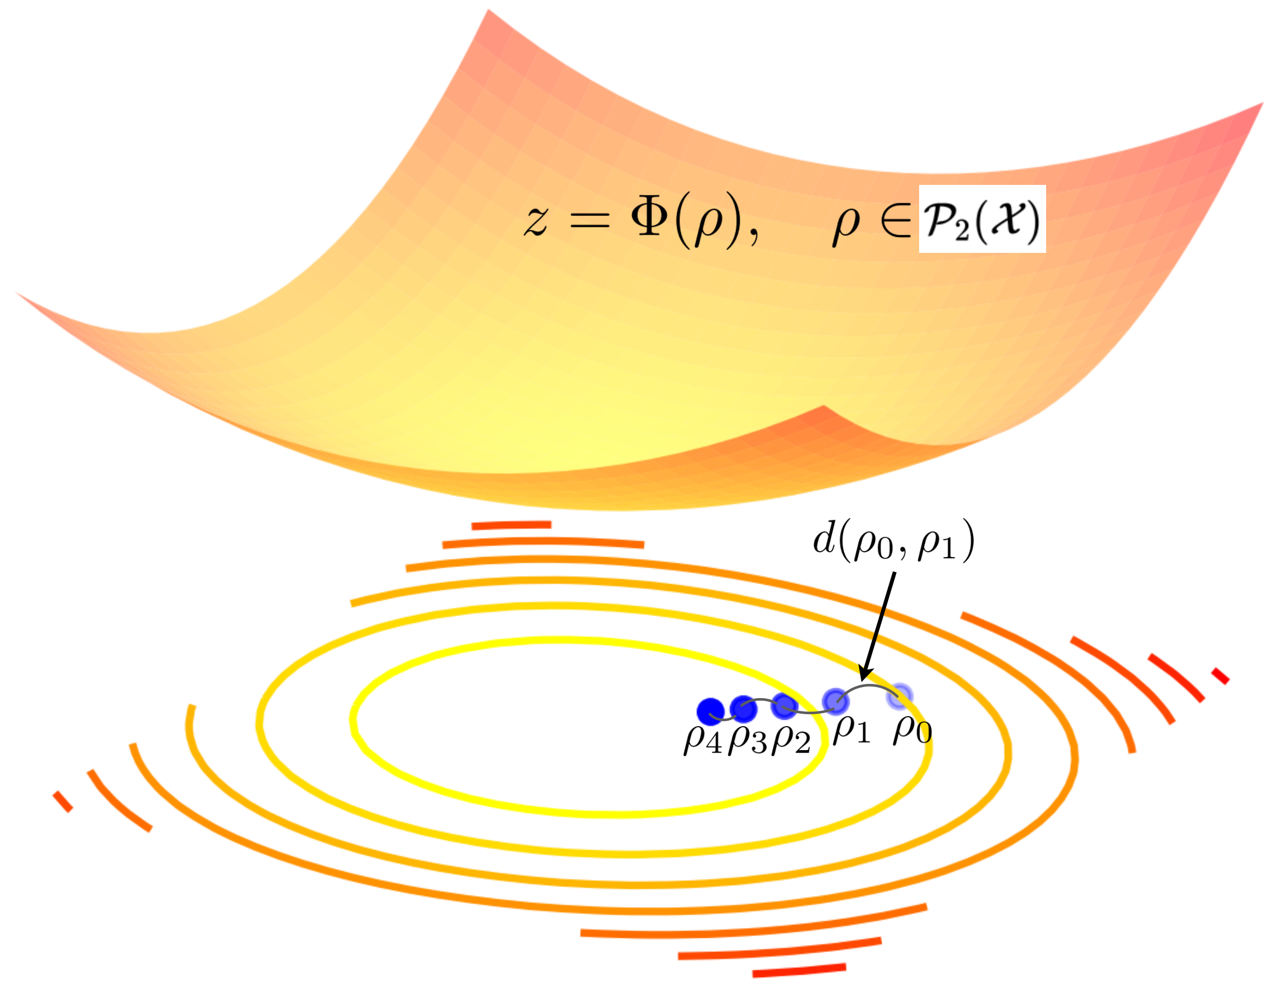

Algorithm: Gradient Ascent on the Dual Space

Uncertainty propagation via point clouds

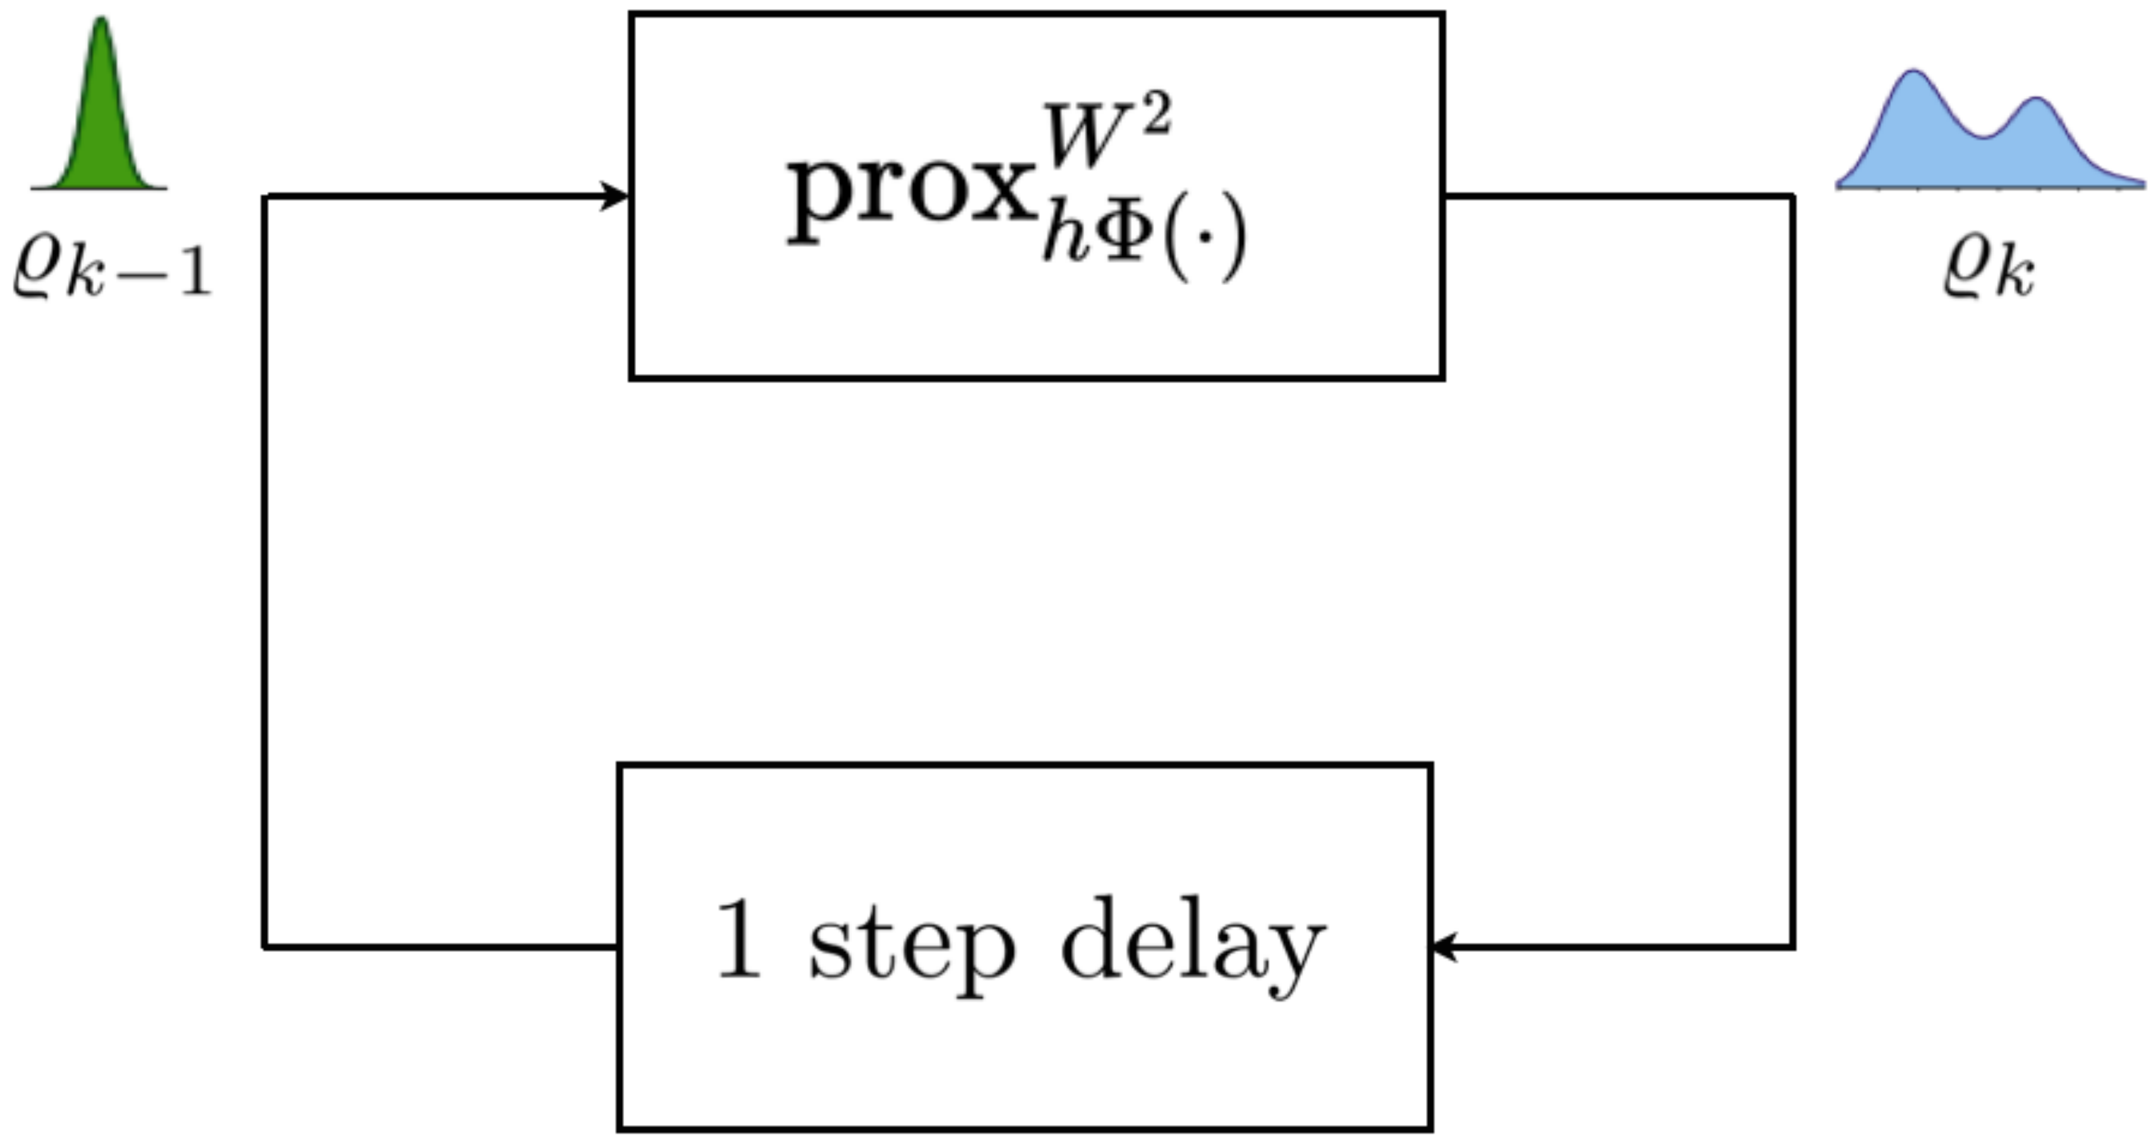

No spatial discretization or function approximation

Algorithm: Gradient Ascent on the Dual Space

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (\nabla \psi \rho) + \beta^{-1} \Delta \rho$$

$\Updownarrow$ **Proximal Recursion**

$$\rho_k = \rho(\boldsymbol{x}, t = kh) = \underset{\rho \in \mathcal{P}_2(\mathbb{R}^n)}{\arg\inf} \left\{ \frac{1}{2} W^2(\rho, \rho_{k-1}) + h\, \Phi(\rho) \right\}$$

$\Downarrow$ **Discrete Primal Formulation**

$$\boldsymbol{\varrho}_k = \underset{\boldsymbol{\varrho}}{\arg\min} \left\{ \min_{\boldsymbol{M} \in \Pi(\boldsymbol{\varrho}_{k-1}, \boldsymbol{\varrho})} \frac{1}{2} \langle \boldsymbol{C}_k, \boldsymbol{M} \rangle + h \left\langle \boldsymbol{\psi}_{k-1} + \beta^{-1} \log \boldsymbol{\varrho}, \boldsymbol{\varrho} \right\rangle \right\}$$

$\Downarrow$ **Entropic Regularization**

$$\boldsymbol{\varrho}_k = \underset{\boldsymbol{\varrho}}{\arg\min} \left\{ \min_{\boldsymbol{M} \in \Pi(\boldsymbol{\varrho}_{k-1}, \boldsymbol{\varrho})} \frac{1}{2} \langle \boldsymbol{C}_k, \boldsymbol{M} \rangle + \epsilon H(\boldsymbol{M}) + h \left\langle \boldsymbol{\psi}_{k-1} + \beta^{-1} \log \boldsymbol{\varrho}, \boldsymbol{\varrho} \right\rangle \right\}$$

$\Updownarrow$ **Dualization**

$$\boldsymbol{\lambda}_0^{\text{opt}}, \boldsymbol{\lambda}_1^{\text{opt}} = \underset{\boldsymbol{\lambda}_0, \boldsymbol{\lambda}_1 \geq 0}{\arg\max} \Bigg\{ \langle \boldsymbol{\lambda}_0, \boldsymbol{\varrho}_{k-1} \rangle - F^{\star}(-\boldsymbol{\lambda}_1) - \frac{\epsilon}{h} \left(\exp(\boldsymbol{\lambda}_0^{\top} h/\epsilon) \exp(-\boldsymbol{C}_k/2\epsilon) \exp(\boldsymbol{\lambda}_1 h/\epsilon) \right) \Bigg\}$$

Recursion on the Cone

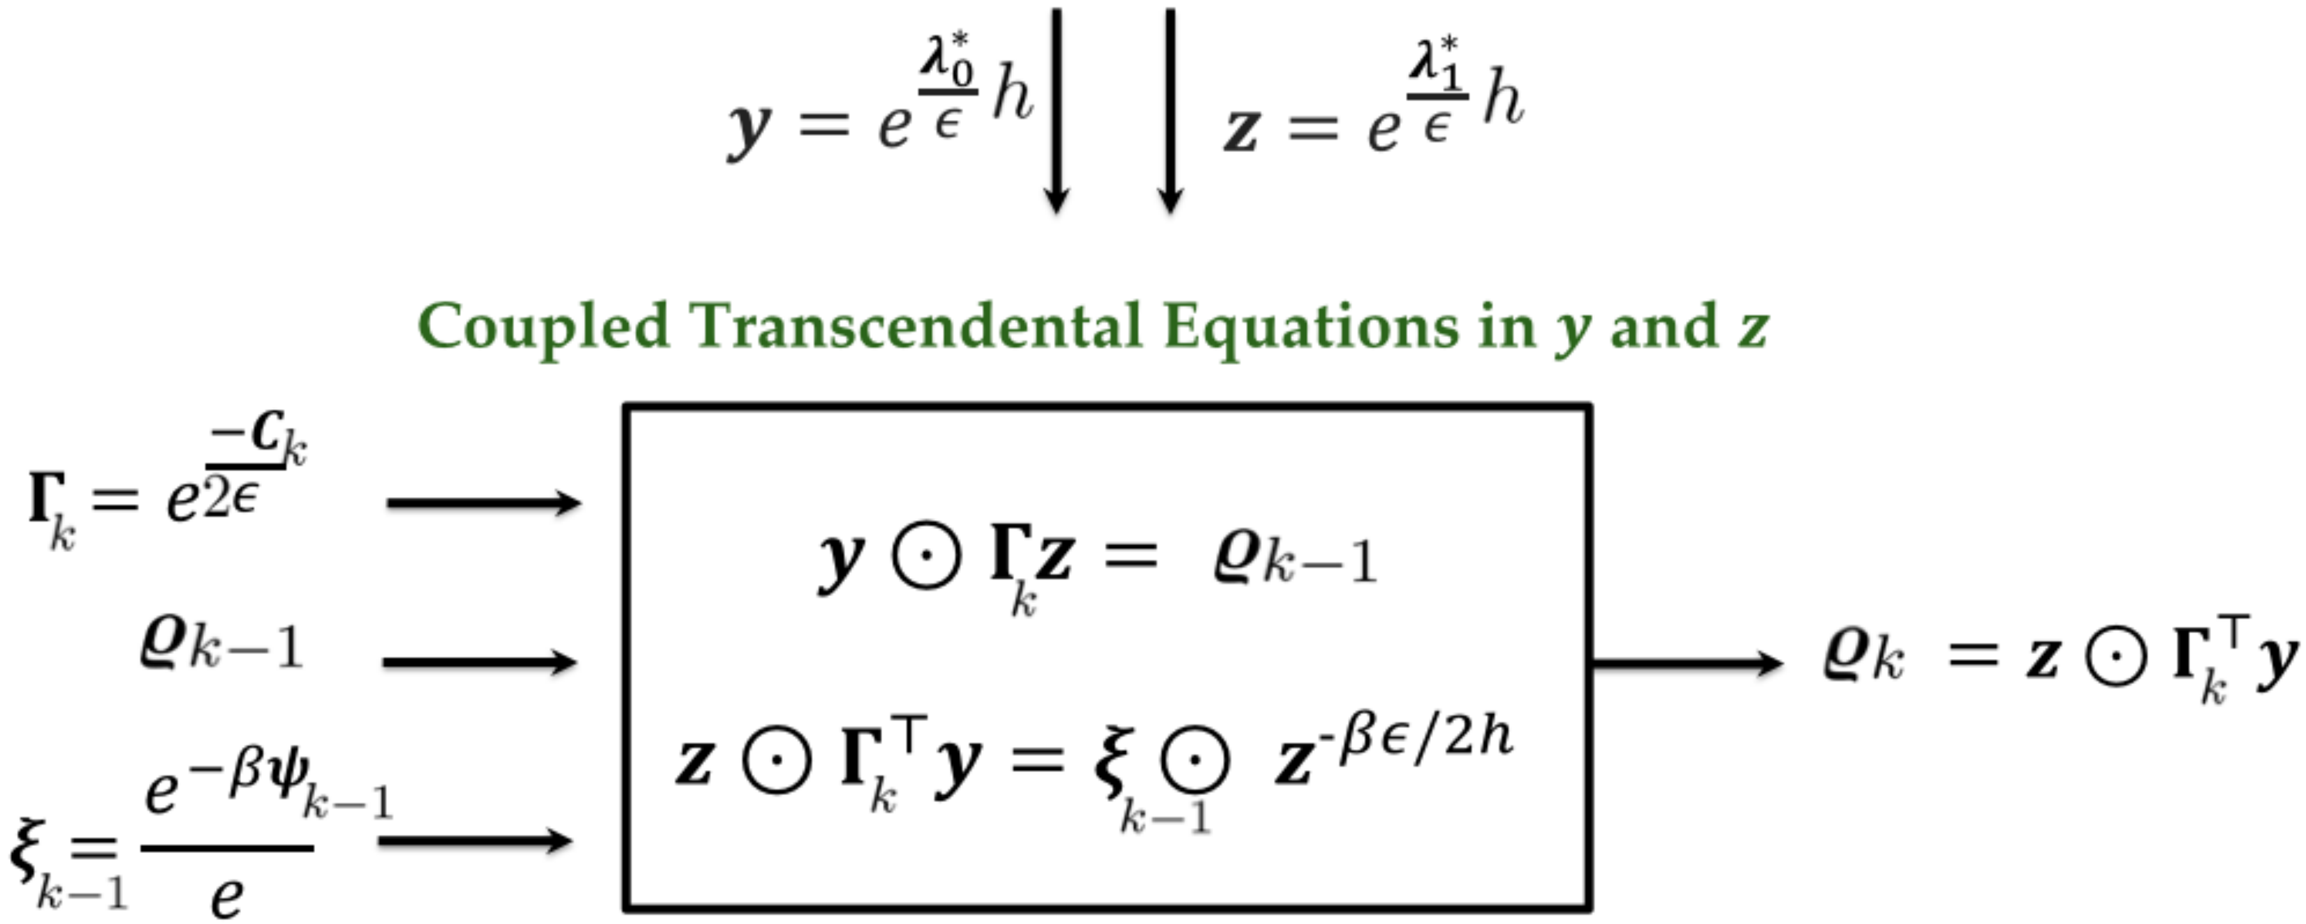

Theorem: Consider the recursion on the cone $\mathbb{R}^n_{\geq 0} \times \mathbb{R}^n_{\geq 0}$

$$\boldsymbol{y} \odot (\boldsymbol{\Gamma}_k \boldsymbol{z}) = \varrho_{k-1}, \quad \boldsymbol{z} \odot \left(\boldsymbol{\Gamma}_k{}^\top \boldsymbol{y}\right) = \boldsymbol{\xi}_{k-1} \odot \boldsymbol{z}^{-\frac{\beta\epsilon}{h}},$$

Then the solution $(\boldsymbol{y}^*, \boldsymbol{z}^*)$ gives the proximal update $\varrho_k = \boldsymbol{z}^* \odot (\boldsymbol{\Gamma}_k{}^\top \boldsymbol{y}^*)$

K.F. Caluya and A.H., Gradient flow algorithms for density propagation in stochastic systems, *IEEE TAC* 2019.

Algorithmic Setup

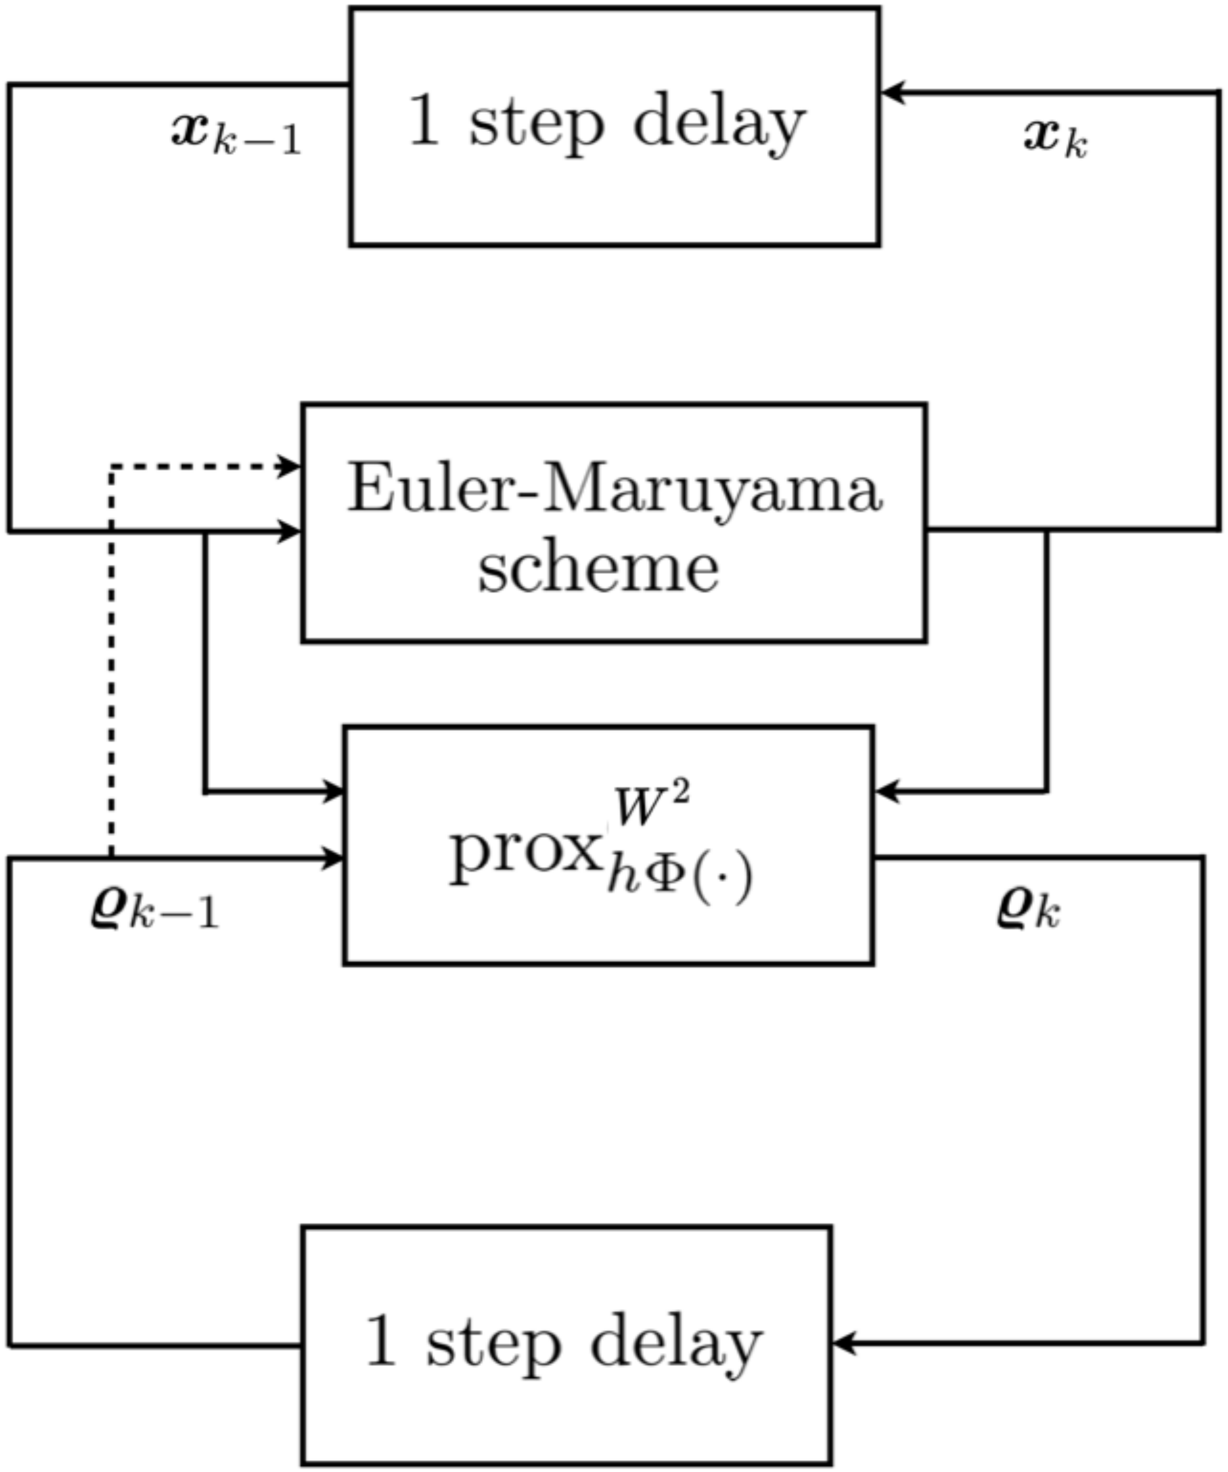

Theorem: Block co-ordinate iteration of $(\mathbf{y}, \mathbf{z})$ recursion is contractive on $\mathbb{R}^n_{>0} \times \mathbb{R}^n_{>0}$.

Proximal Prediction: 2D Linear Gaussian

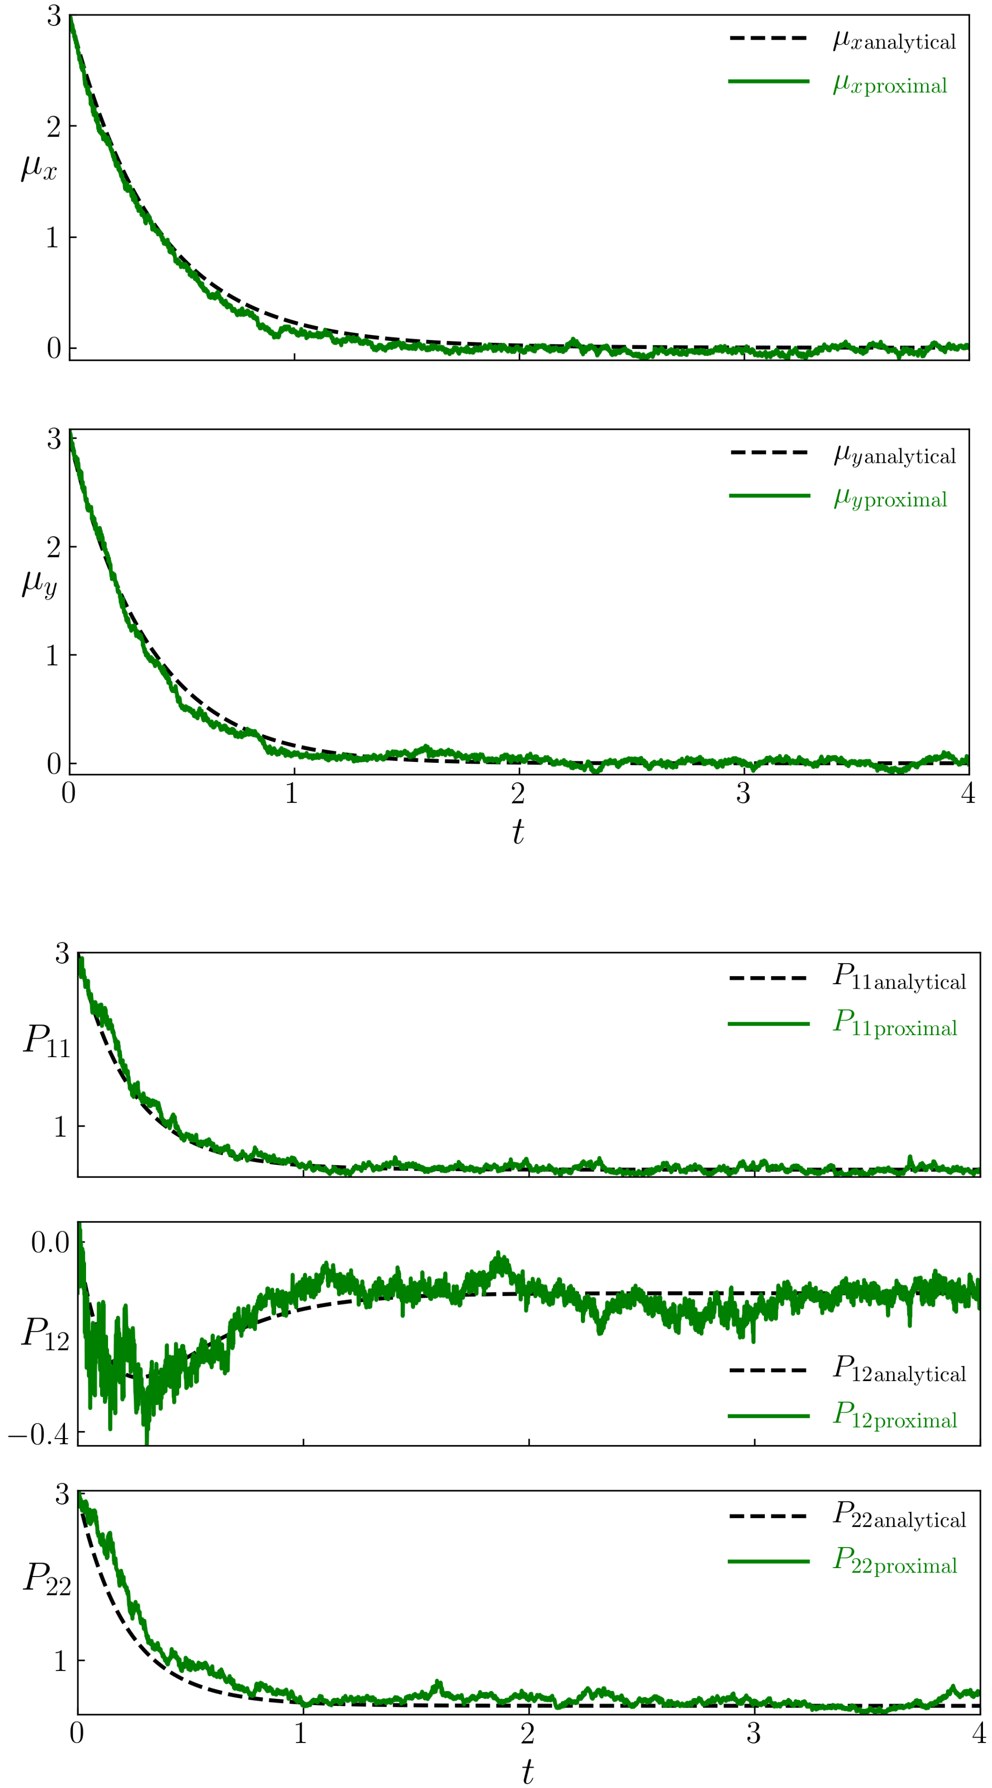

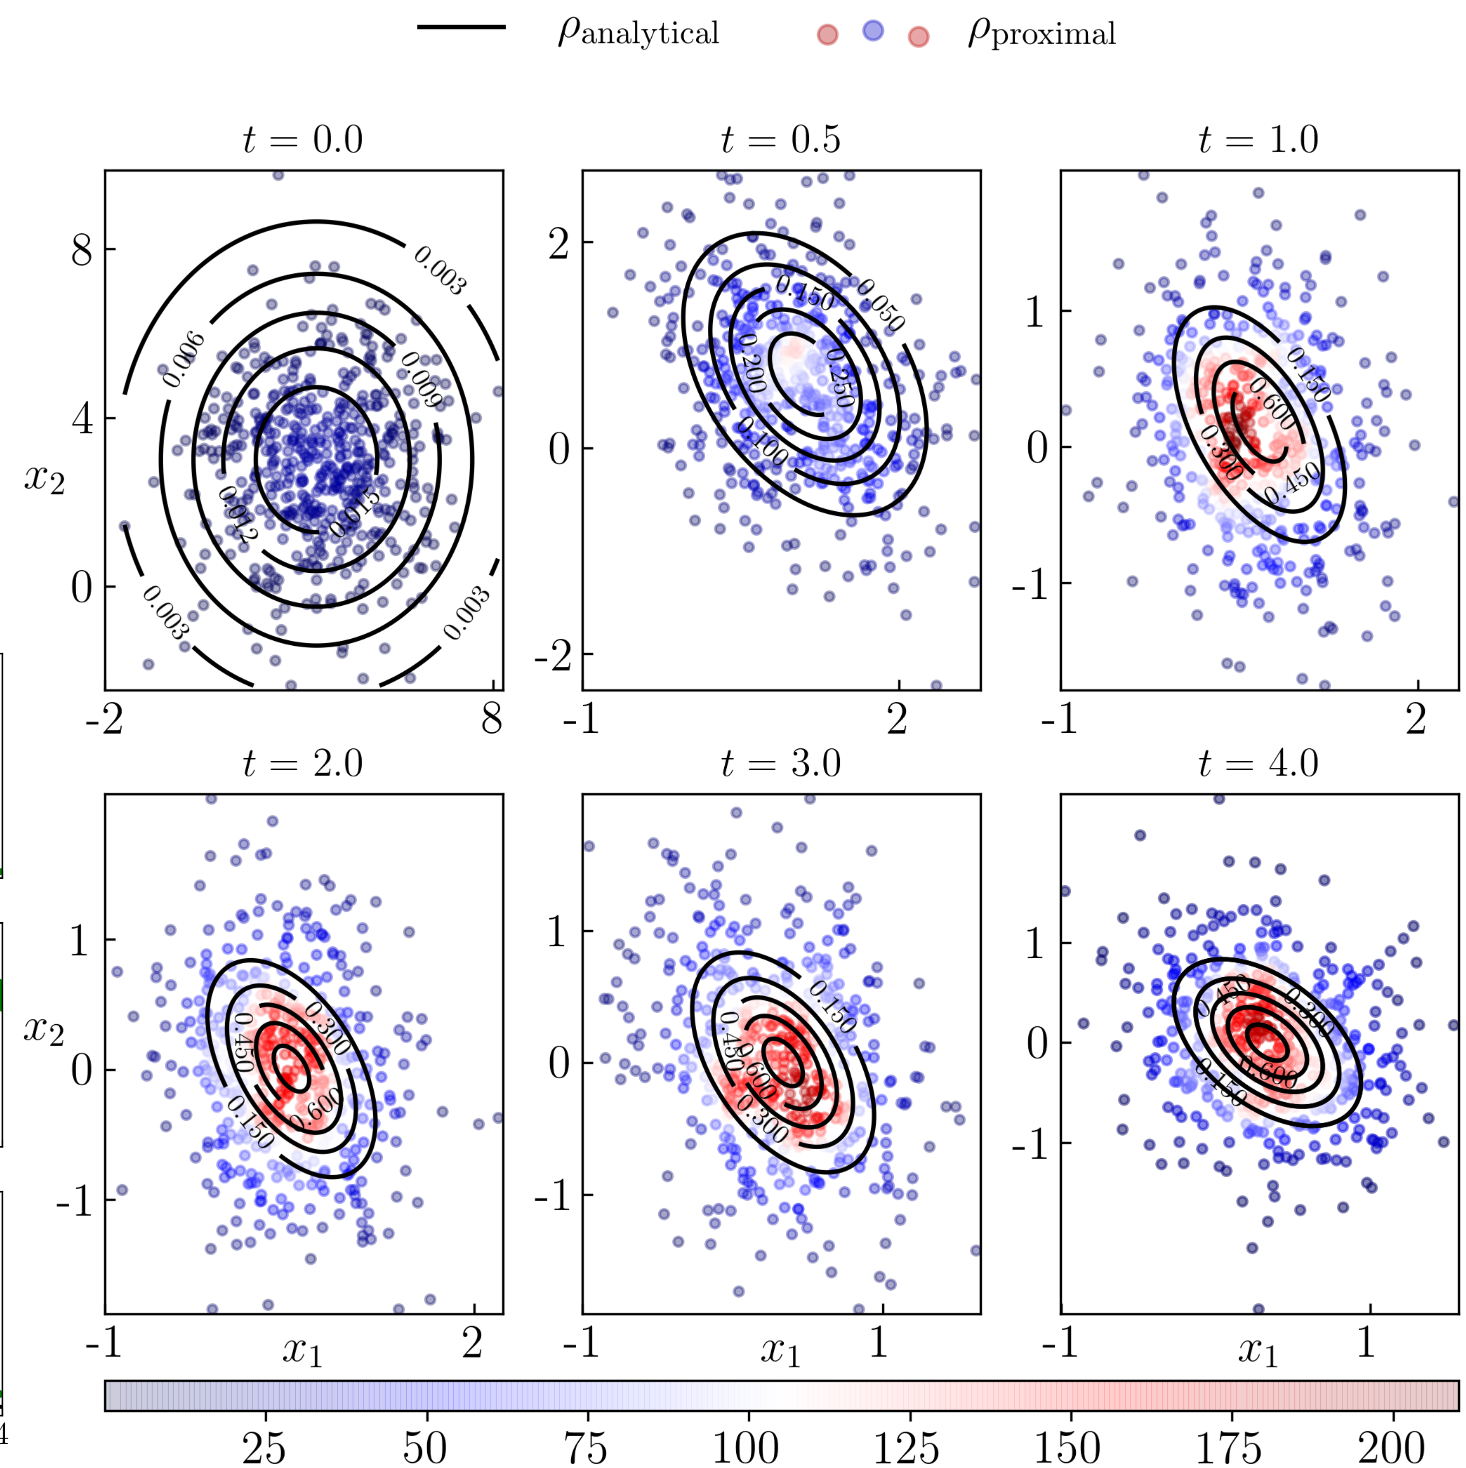

Proximal Prediction: Nonlinear Non-Gaussian

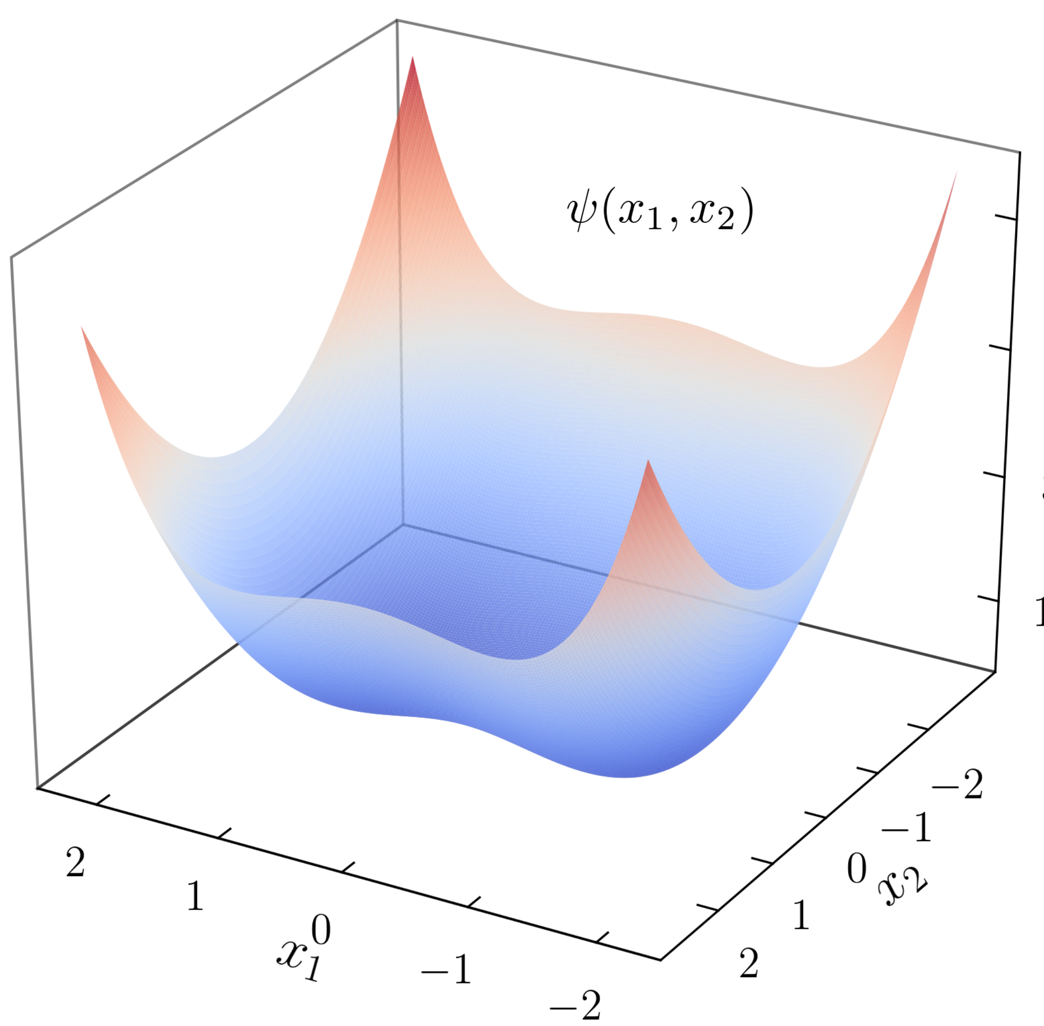

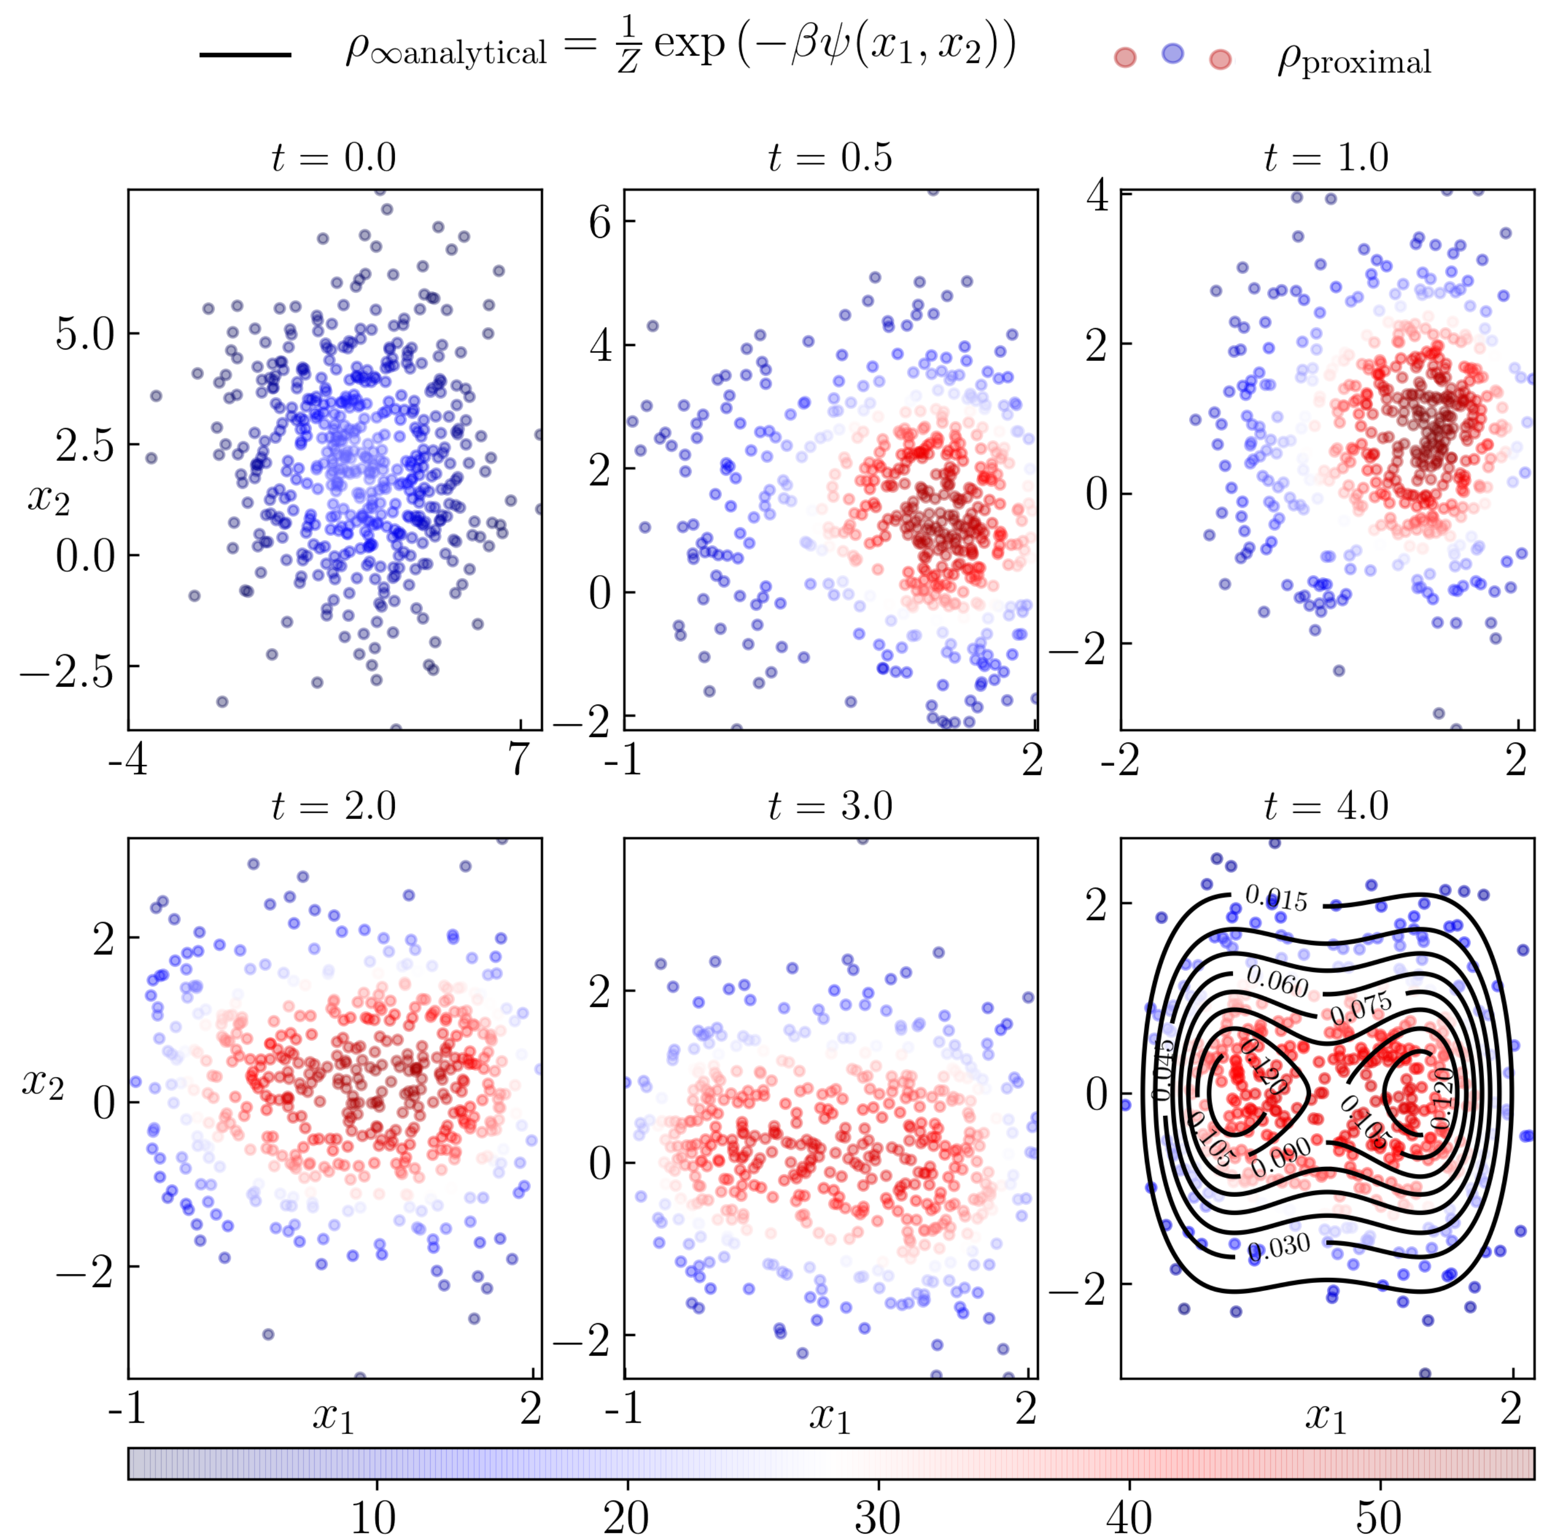

Computational Time: Nonlinear Non-Gaussian

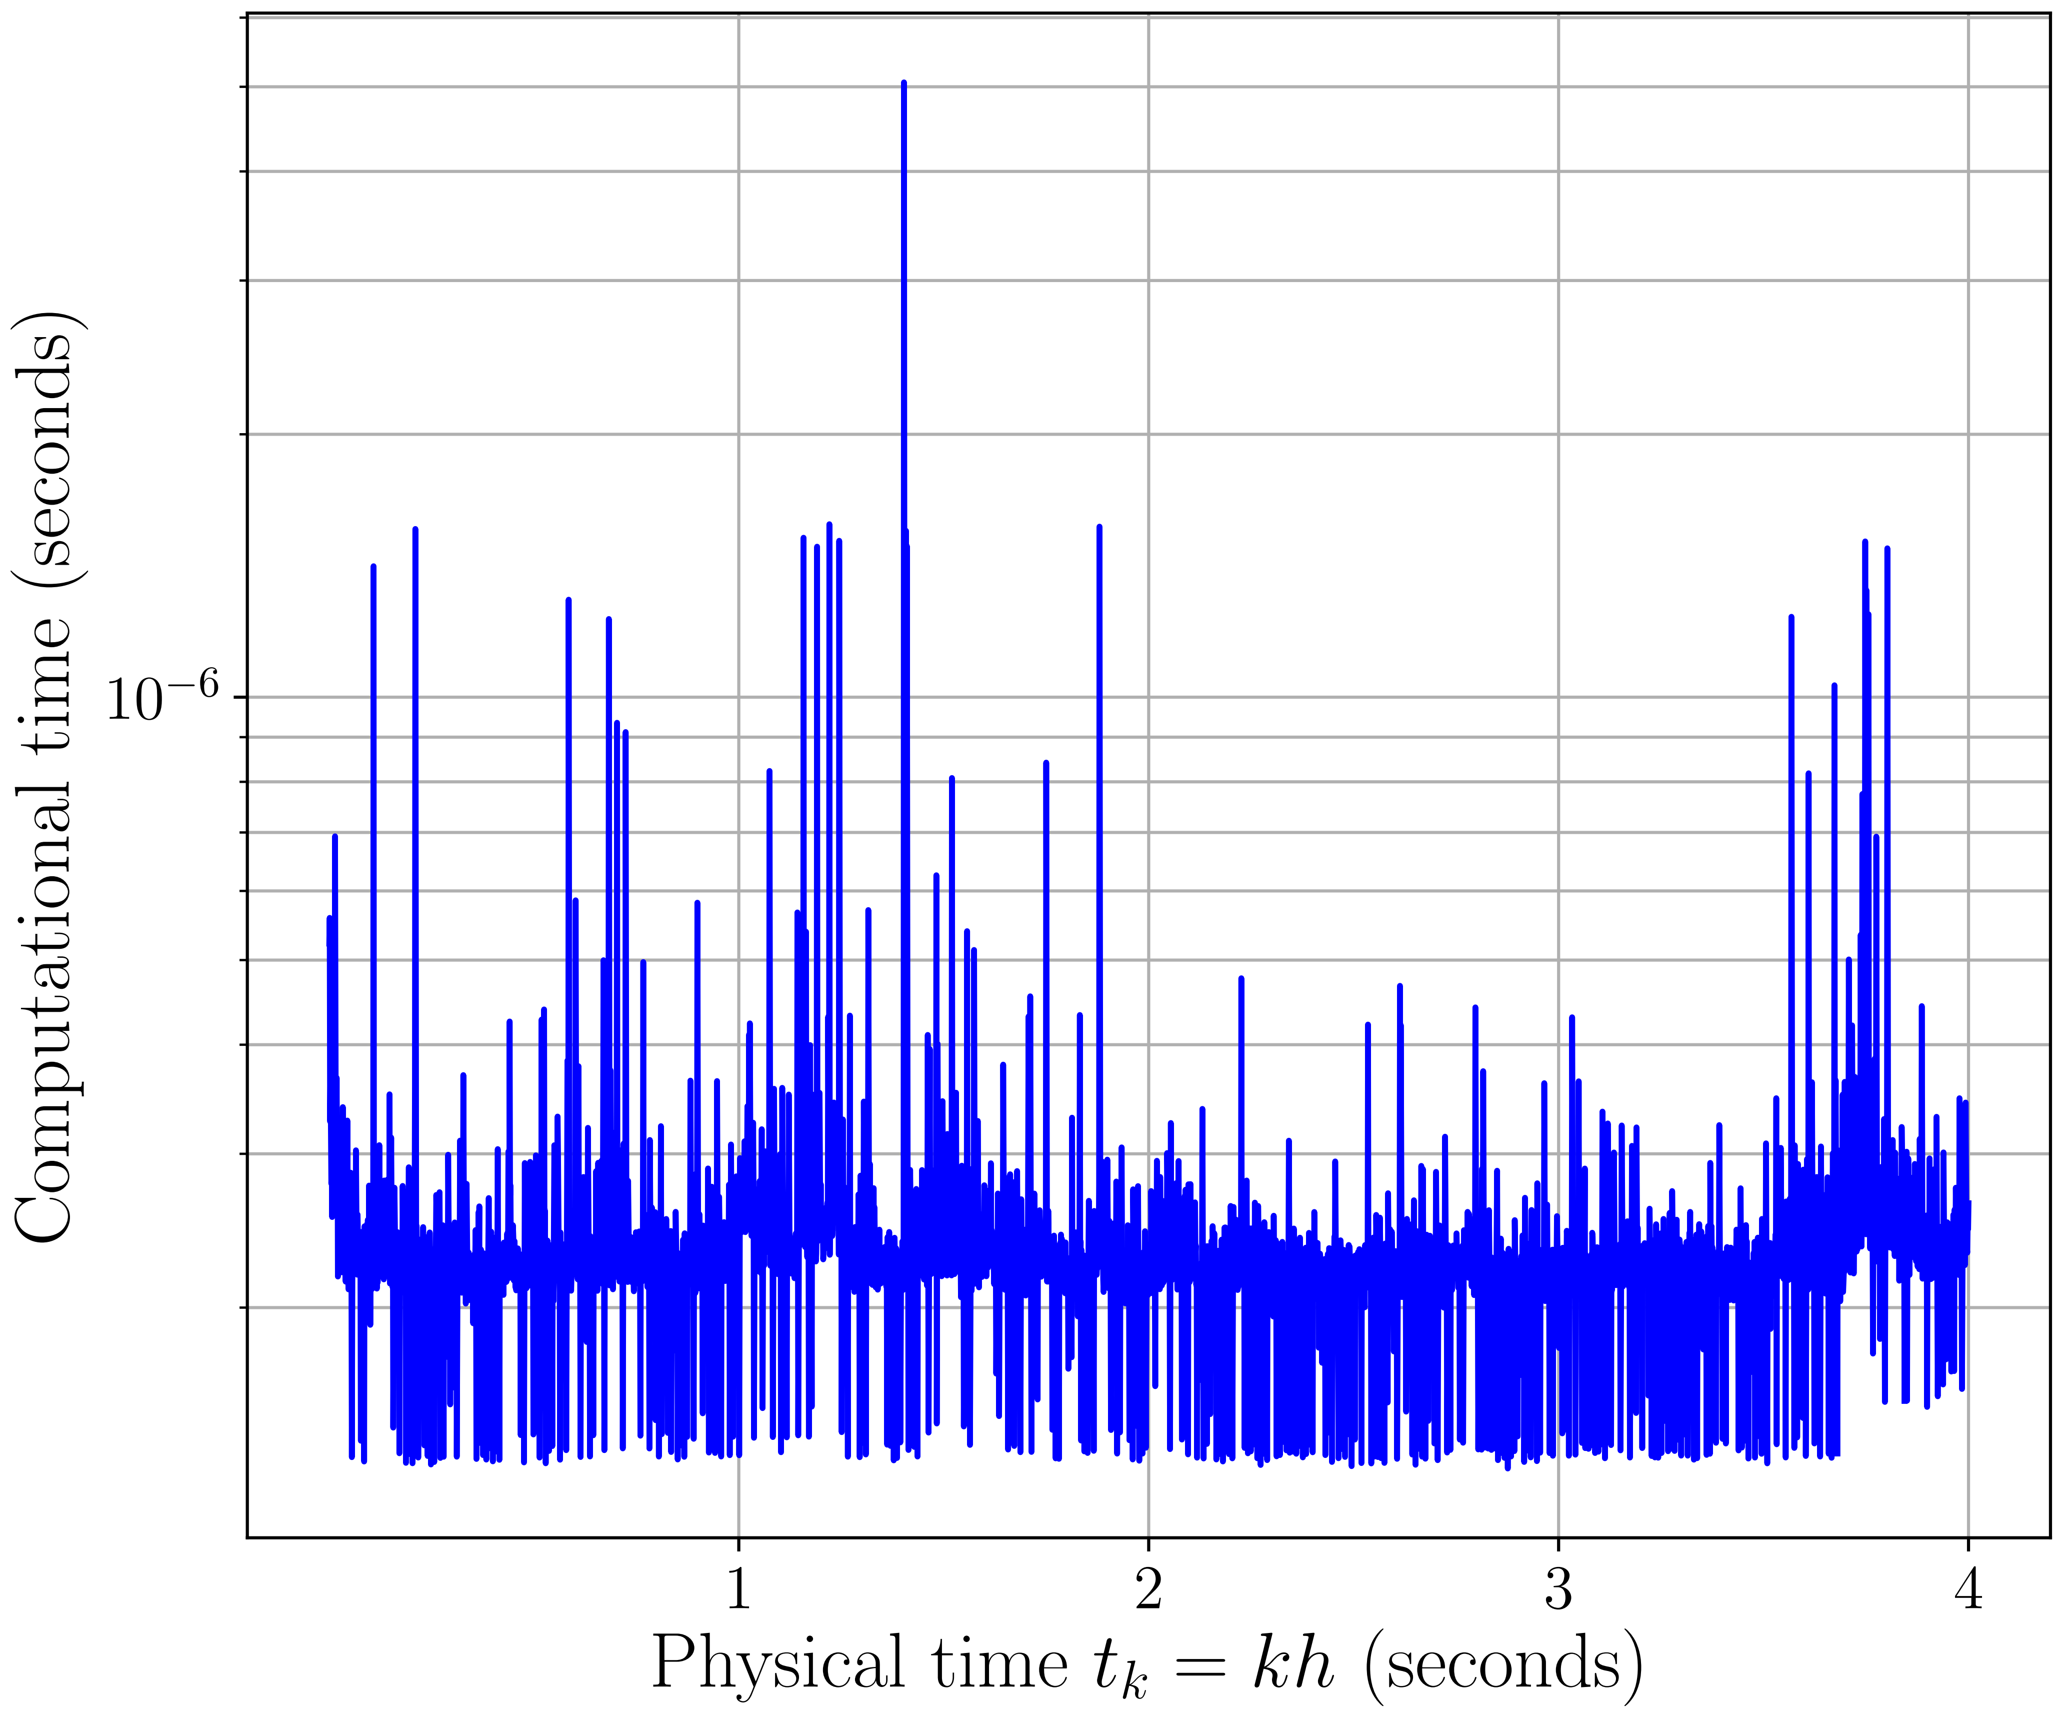

Proximal Prediction: Satellite in Geocentric Orbit

Here, $\mathcal{X} \equiv \mathbb{R}^6$

$$\begin{pmatrix} \mathrm{d}x \\ \mathrm{d}y \\ \mathrm{d}z \\ \mathrm{d}v_x \\ \mathrm{d}v_y \\ \mathrm{d}v_z \end{pmatrix} = \begin{pmatrix} v_x \\ v_y \\ v_z \\ -\dfrac{\mu x}{r^3} + (f_x)_{\text{pert}} - \gamma v_x \\ -\dfrac{\mu y}{r^3} + (f_y)_{\text{pert}} - \gamma v_y \\ -\dfrac{\mu z}{r^3} + (f_z)_{\text{pert}} - \gamma v_z \end{pmatrix} \mathrm{d}t + \sqrt{2\beta^{-1}\gamma} \begin{pmatrix} 0 \\ 0 \\ 0 \\ \mathrm{d}w_1 \\ \mathrm{d}w_2 \\ \mathrm{d}w_3 \end{pmatrix},$$

$$\begin{pmatrix} f_x \\ f_y \\ f_z \end{pmatrix}_{\text{pert}} = \begin{pmatrix} s\theta\, c\phi & c\theta\, c\phi & -s\phi \\ s\theta\, s\phi & c\theta\, s\phi & c\phi \\ c\theta & -s\theta & 0 \end{pmatrix} \begin{pmatrix} \dfrac{k}{2r^4}\left(3(s\theta)^2 - 1\right) \\ -\dfrac{k}{r^5} s\theta\, c\theta \\ 0 \end{pmatrix}, k := 3J_2 R_{\mathrm{E}}^2, \mu = \text{constant}$$

Computational Time: Satellite in Geocentric Orbit

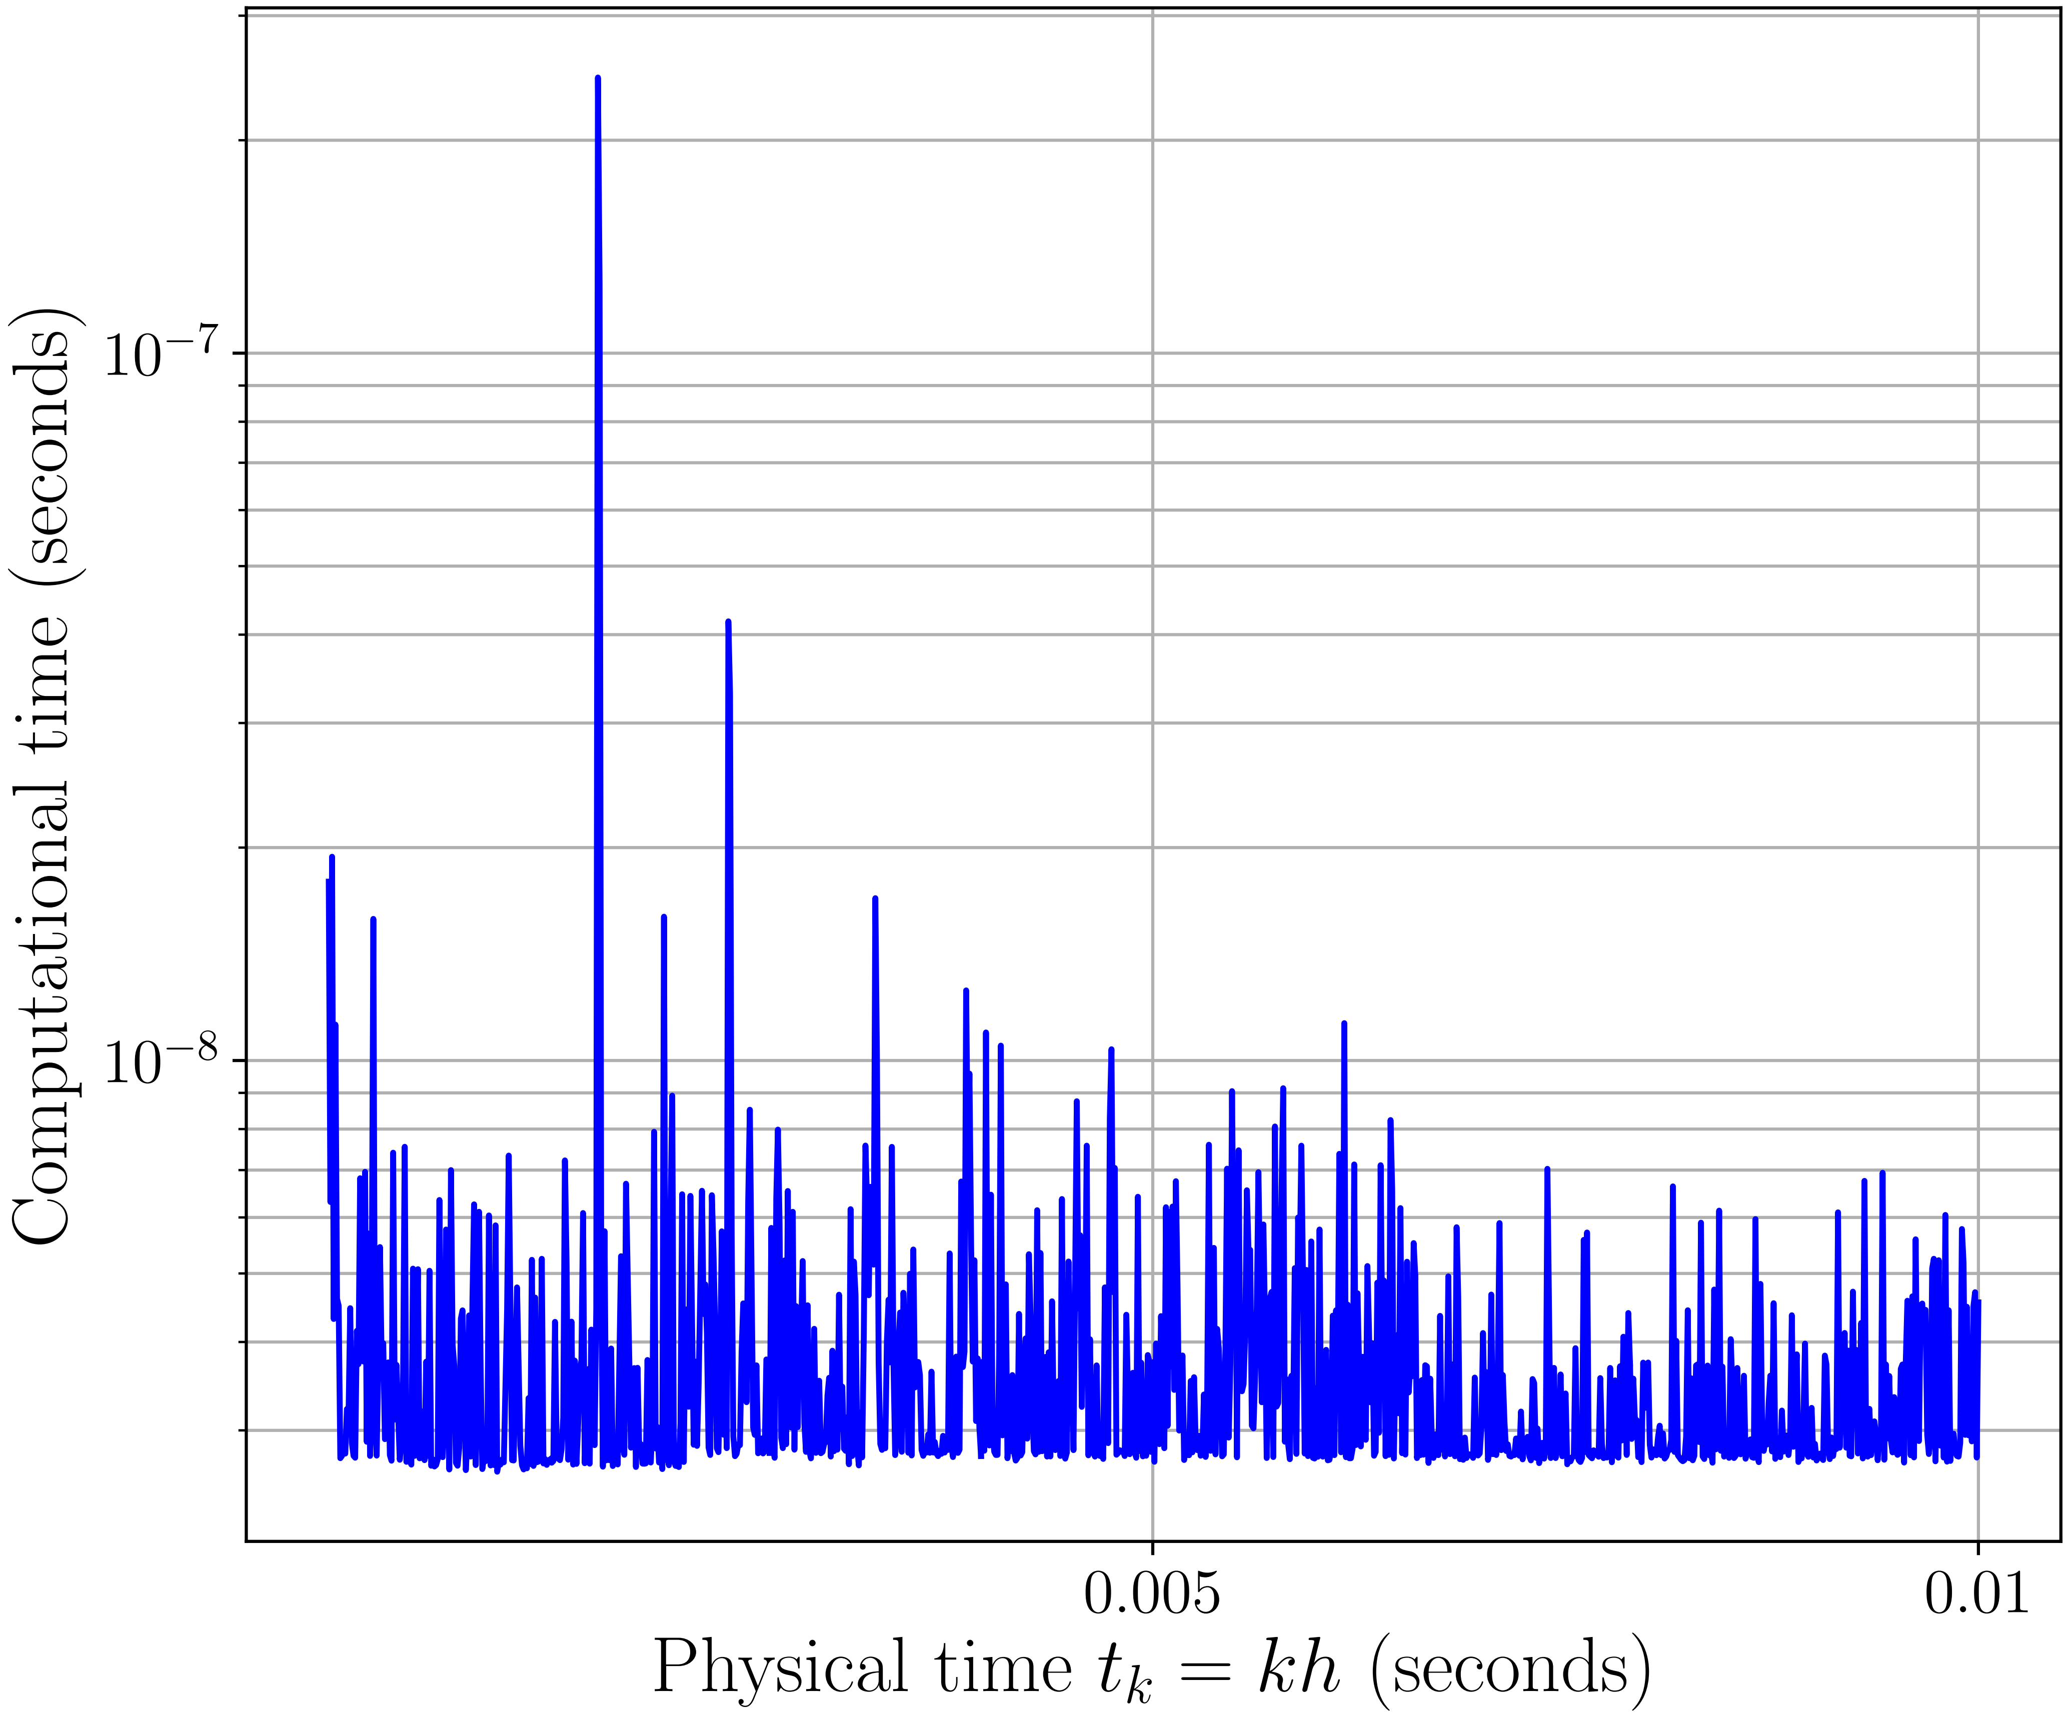

Extensions and Applications

Networked nonlinear power system dynamics with $O(100)$ states

A.H., K.F. Caluya, P. Ojaghi, and X. Geng, Stochastic uncertainty propagation in power system dynamics with measure-valued proximal recursions, *IEEE Transactions on Power Systems*, 2022.

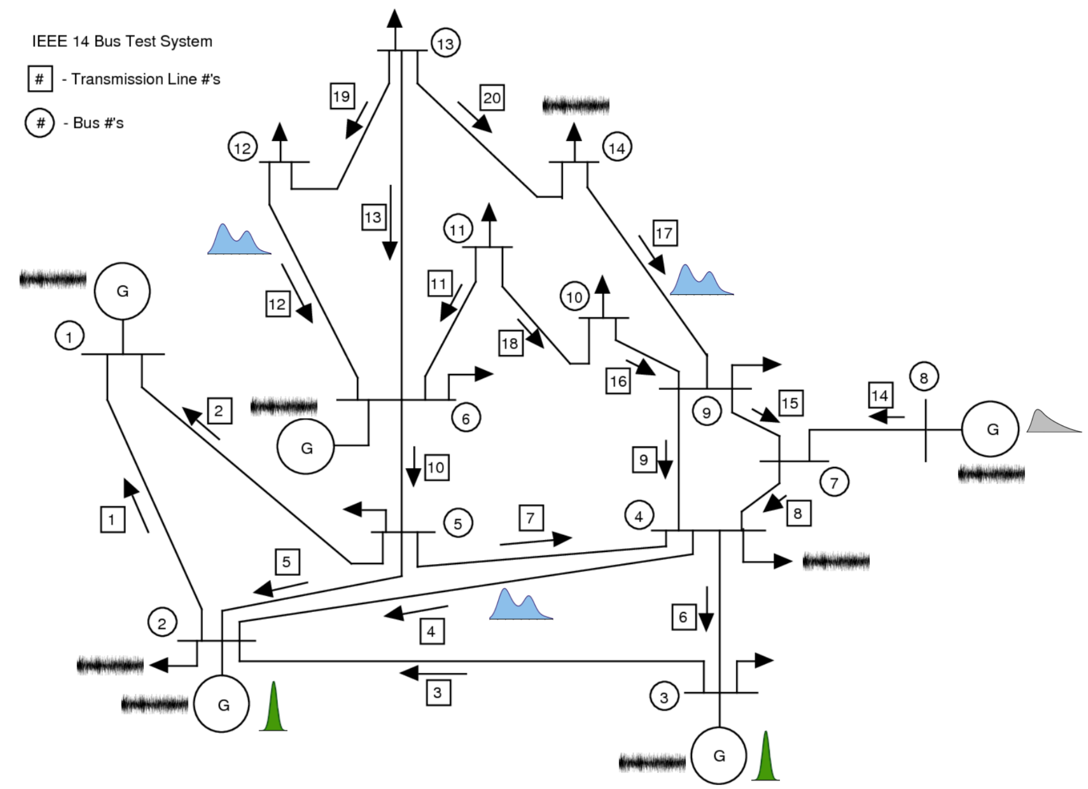

Mean field learning in NN

A.M.H. Teter, I. Nodozi, and A.H., Proximal mean field learning in shallow neural networks, *arXiv:2210:13879*, 2022.

Case study: Wisconsin Breast Cancer Diagnostic (WBCD) Data Set

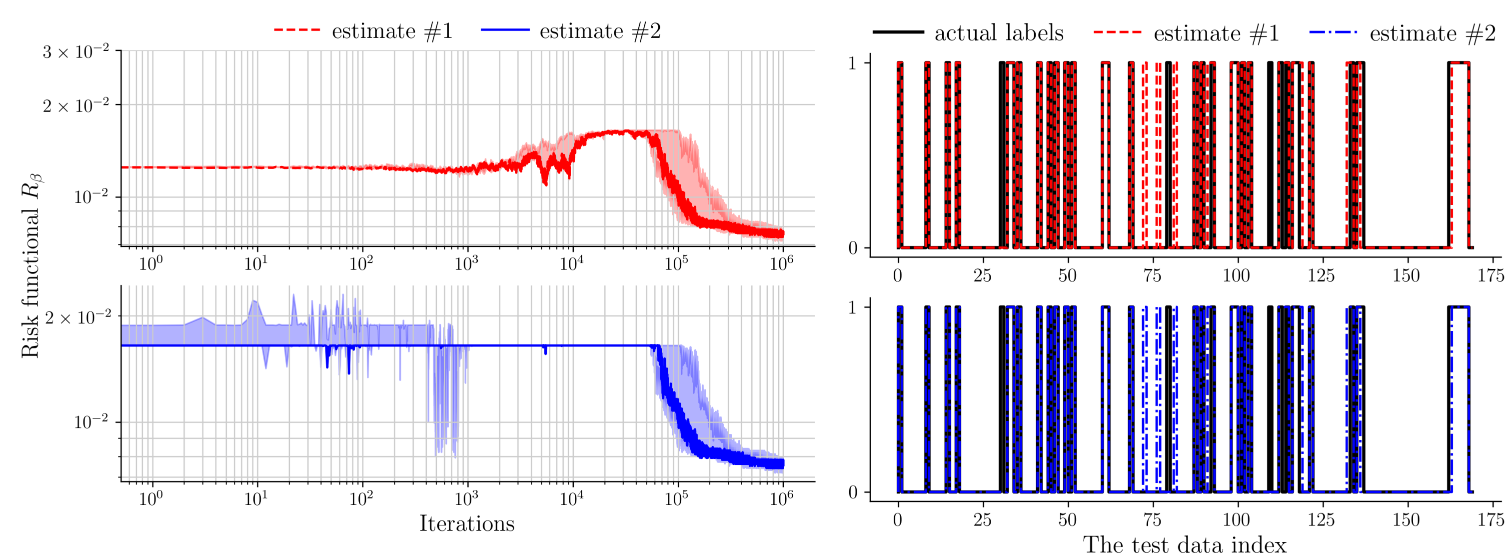

GPU: Jetson TX2 NVIDIA Pascal GPU 256 CUDA cores, 64 bit NVIDIA Denver + ARM Cortex A57 CPUs (≈ 2 hrs runtime)

Classification accuracy for the WBDC dataset

β	Estimate #1	Estimate #2
0.03	91.17%	92.35%
0.05	92.94%	92.94%
0.07	78.23%	92.94%

Solving filtering as generalized gradient flow

What's New?

Main idea: Solve the Kushner-Stratonovich SPDE

$$\mathrm{d}\rho^{+} = \left[\mathcal{L}_{\mathrm{FP}}\mathrm{d}t + \mathcal{L}\left(\mathrm{d}z, \mathrm{d}t, \rho^{+}\right)\right]\rho^{+}, \;\; \rho(x, t=0) = \rho_0 \text{ as gradient flow in } \mathcal{P}_2(\mathcal{X})$$

Recursion of {deterministic $\circ$ stochastic} proximal operators:

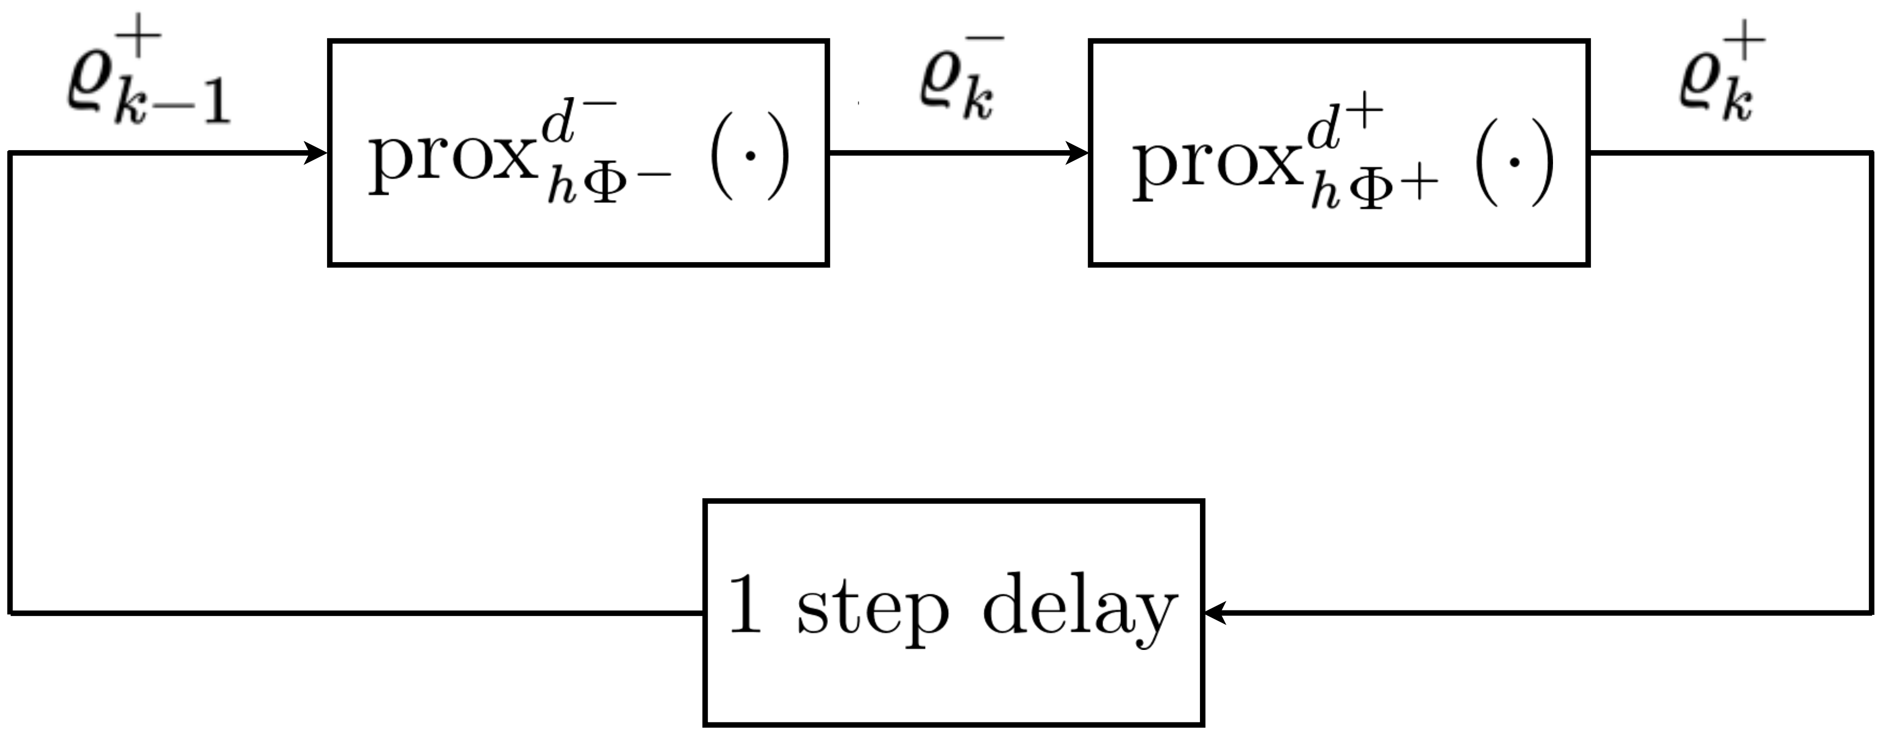

Convergence: $\varrho_k^{+}(h) \rightarrow \rho^{+}(x, t=kh) \quad \text{as} \quad h \downarrow 0$

For prior, as before: $d^{-} \equiv W^2, \quad \Phi^{-} \equiv \mathbb{E}_{\varrho}\left[\psi + \beta^{-1}\log\varrho\right]$

For posterior: $d^{+} \equiv d_{\mathrm{FR}}^2$ or $D_{\mathrm{KL}}, \quad \Phi^{+} \equiv \dfrac{1}{2}\mathbb{E}_{\varrho^{+}}\left[(y_k - h(x))^{\top} R^{-1} (y_k - h(x))\right]$

Explicit Recovery of the Kalman-Bucy Filter

Model:

$$\mathrm{d}\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t)\mathrm{d}t + \mathbf{B}\mathrm{d}\mathbf{w}(t), \qquad \mathrm{d}\mathbf{w}(t) \sim \mathcal{N}(0, \mathbf{Q}\mathrm{d}t)$$

$$\mathrm{d}\mathbf{z}(t) = \mathbf{C}\mathbf{x}(t)\mathrm{d}t + \mathrm{d}\mathbf{v}(t), \qquad \mathrm{d}\mathbf{v}(t) \sim \mathcal{N}(0, \mathbf{R}\mathrm{d}t)$$

Given $\mathbf{x}(0) \sim \mathcal{N}(\mu_0, \mathbf{P}_0)$, want to recover:

$$\mathbf{P}^{+}\mathbf{C}\mathbf{R}^{-1}$$

$$\mathrm{d}\mu^{+}(t) = \mathbf{A}\mu^{+}(t)\mathrm{d}t + \mathbf{K}(t) \left(\mathrm{d}\mathbf{z}(t) - \mathbf{C}\mu^{+}(t)\mathrm{d}t\right),$$

$$\dot{\mathbf{P}}^{+}(t) = \mathbf{A}\mathbf{P}^{+}(t) + \mathbf{P}^{+}(t)\mathbf{A}^{\top} + \mathbf{B}\mathbf{Q}\mathbf{B}^{\top} - \mathbf{K}(t)\mathbf{R}\mathbf{K}(t)^{\top}.$$

A.H. and T.T. Georgiou, Gradient Flows in Uncertainty Propagation and Filtering of Linear Gaussian Systems, *CDC 2017*.

A.H. and T.T. Georgiou, Gradient Flows in Filtering and Fisher-Rao Geometry, *ACC 2018*.

Explicit Recovery of the Wonham Filter

Model:

$$x(t) \sim \text{Markov}(Q),$$
$$\mathrm{d}z(t) = h(x(t))\,\mathrm{d}t \;+\; \sigma_v(t)\mathrm{d}v(t)$$

J.SIAM Control
Ser. A, Vol. 2, No. 3
Printed in U.S.A., 1965

SOME APPLICATIONS OF STOCHASTIC DIFFERENTIAL EQUATIONS TO OPTIMAL NONLINEAR FILTERING*

W. M. WONHAM†

State space: $\Omega := \{a_1, \ldots, a_m\}$

Posterior $\pi^+(t) := \{\pi_1^+(t), \ldots, \pi_m^+(t)\}$ **solves the nonlinear SDE:**

$$\mathrm{d}\pi^+(t) = \pi^+(t)Q\,\mathrm{d}t \;+\; \frac{1}{(\sigma_v(t))^2}\pi^+(t)\Big(H - \widehat{h}(t)I\Big)\Big(\mathrm{d}z(t) - \widehat{h}(t)\mathrm{d}t\Big),$$

where $H := \mathrm{diag}(h(a_1), \ldots, h(a_m)), \quad \widehat{h}(t) := \sum_{i=1}^{m} h(a_i)\pi_i^+(t),$

Initial condition: $\pi^+(t=0) = \pi_0,$

By defn. $\pi^+(t) = \mathbb{P}(x(t) = a_i \mid z(s), 0 \leq s \leq t)$

A.H. and T.T. Georgiou, Proximal Recursion for the Wonham Filter, *CDC 2019*.

Numerical Results for the Wonham Filter

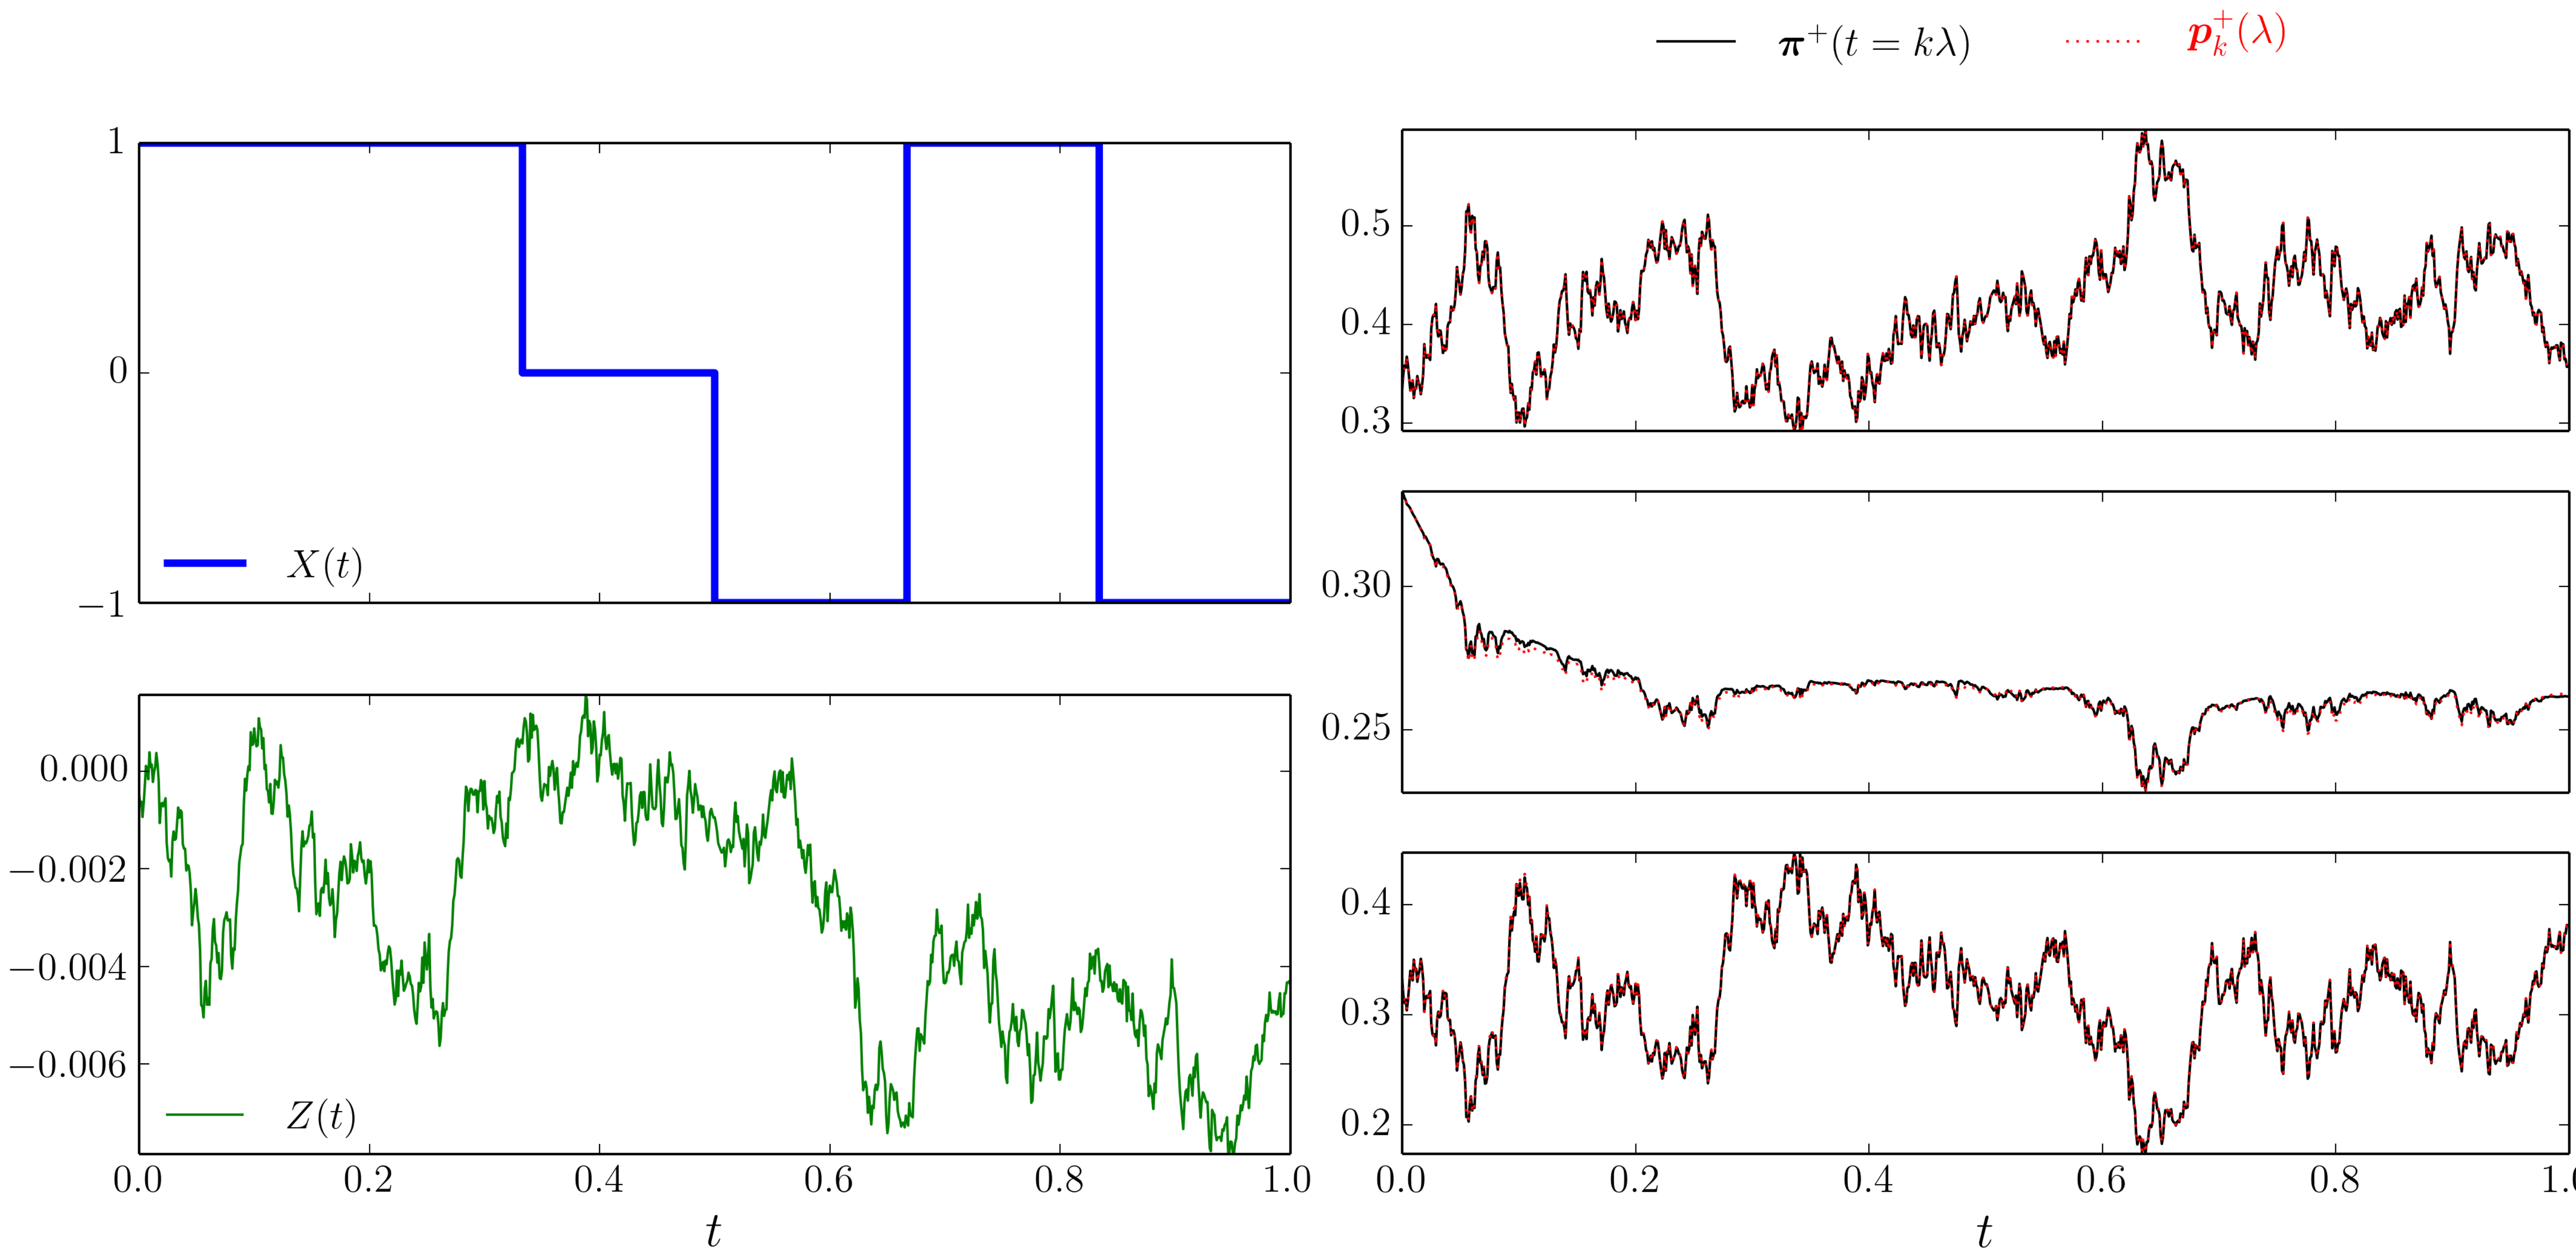

A.H. and T.T. Georgiou, Proximal Recursion for the Wonham Filter, *CDC 2019*.

Solving density control as generalized gradient flow

State Feedback Density Steering

Steer joint state PDF via feedback control over finite time horizon

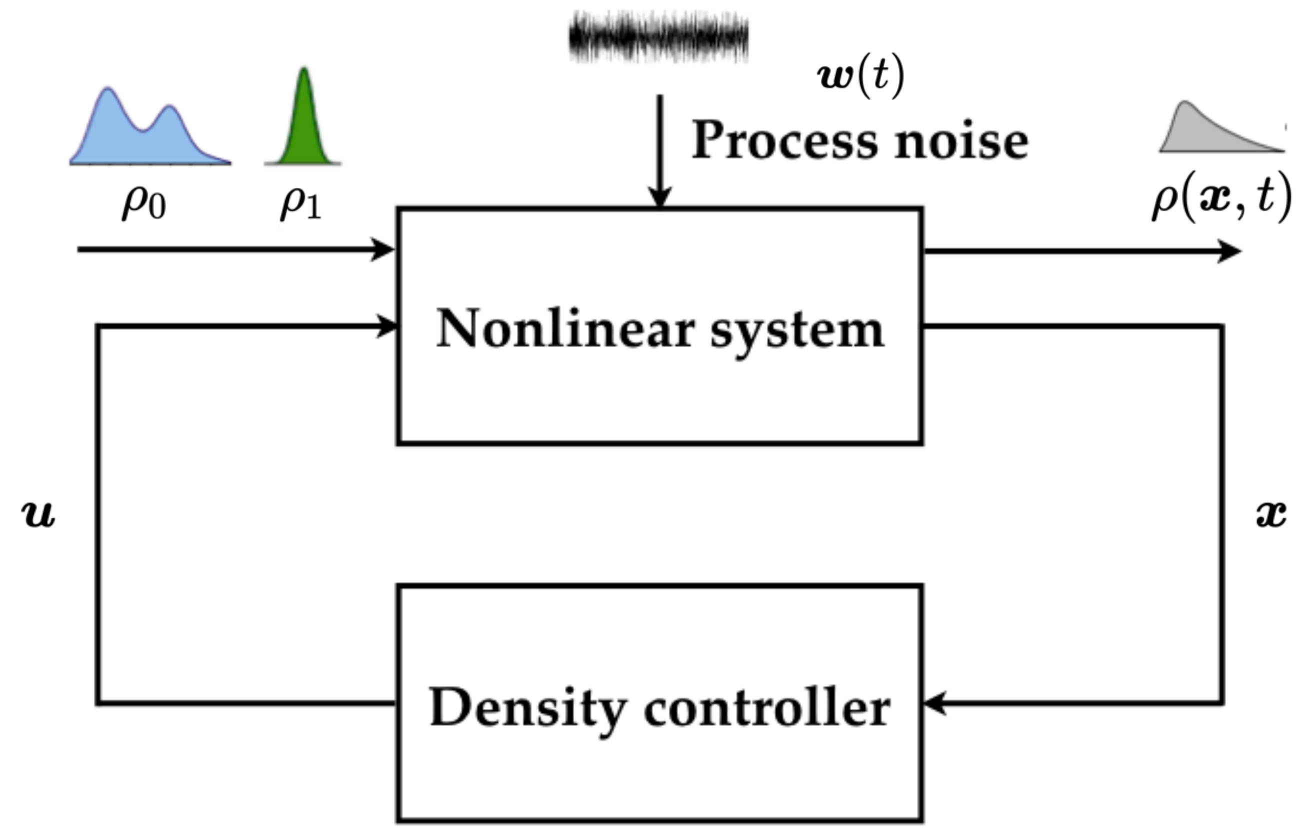

Common scenario: $\boldsymbol{G} \equiv \boldsymbol{B}$

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

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

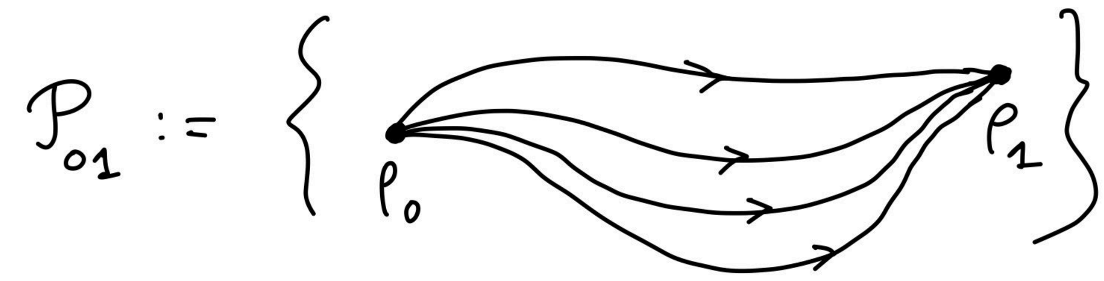

Diffusion tensor: $\boldsymbol{D} := \boldsymbol{G}\boldsymbol{G}^\top$

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

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

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

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

Coupled nonlinear PDEs + linear boundary conditions

Controlled Fokker-Planck or Kolmogorov's forward PDE

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

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

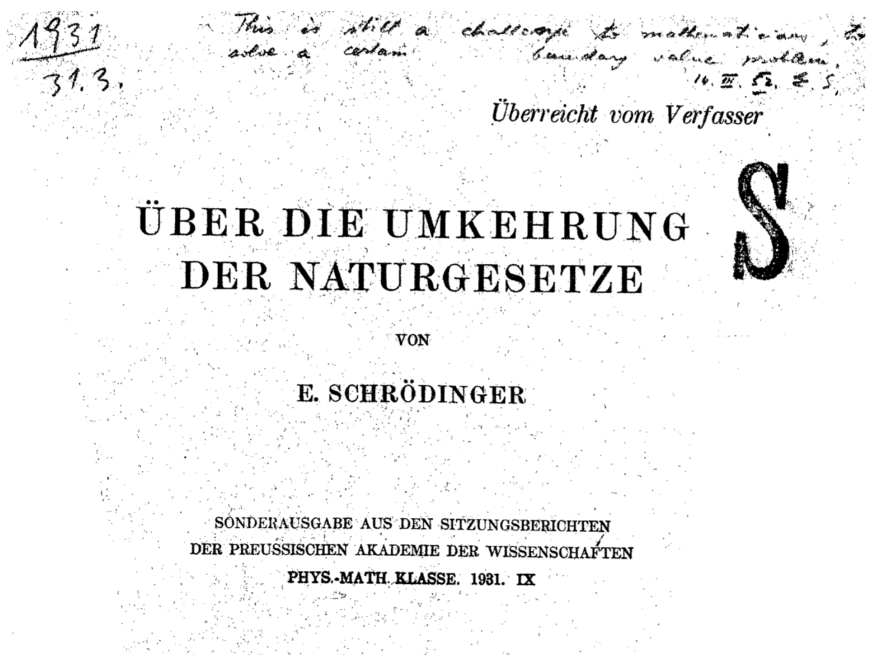

1931
31.3.
This is still a challenge to mathematicians, to solve a certain boundary value problem.
14.III.52. E.S.

Überreicht vom Verfasser

ÜBER DIE UMKEHRUNG
DER NATURGESETZE

VON

E. SCHRÖDINGER

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

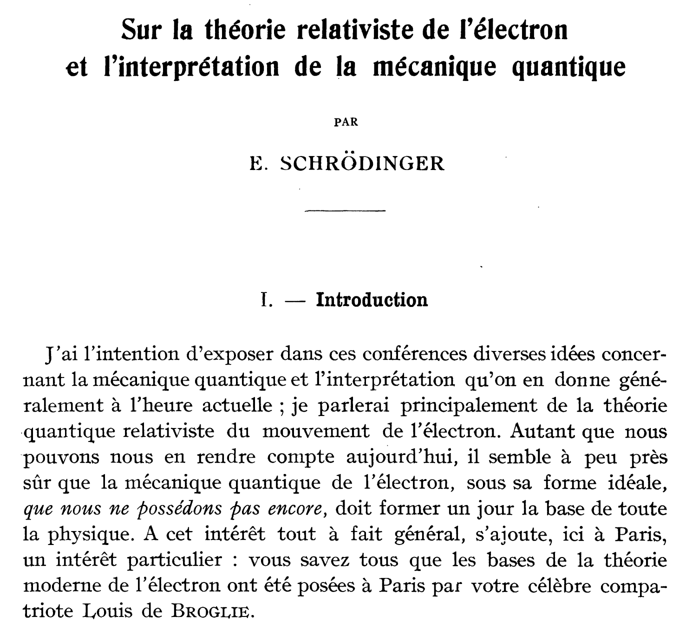

Sur la théorie relativiste de l'électron
et l'interprétation de la mécanique quantique

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:

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

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

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

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

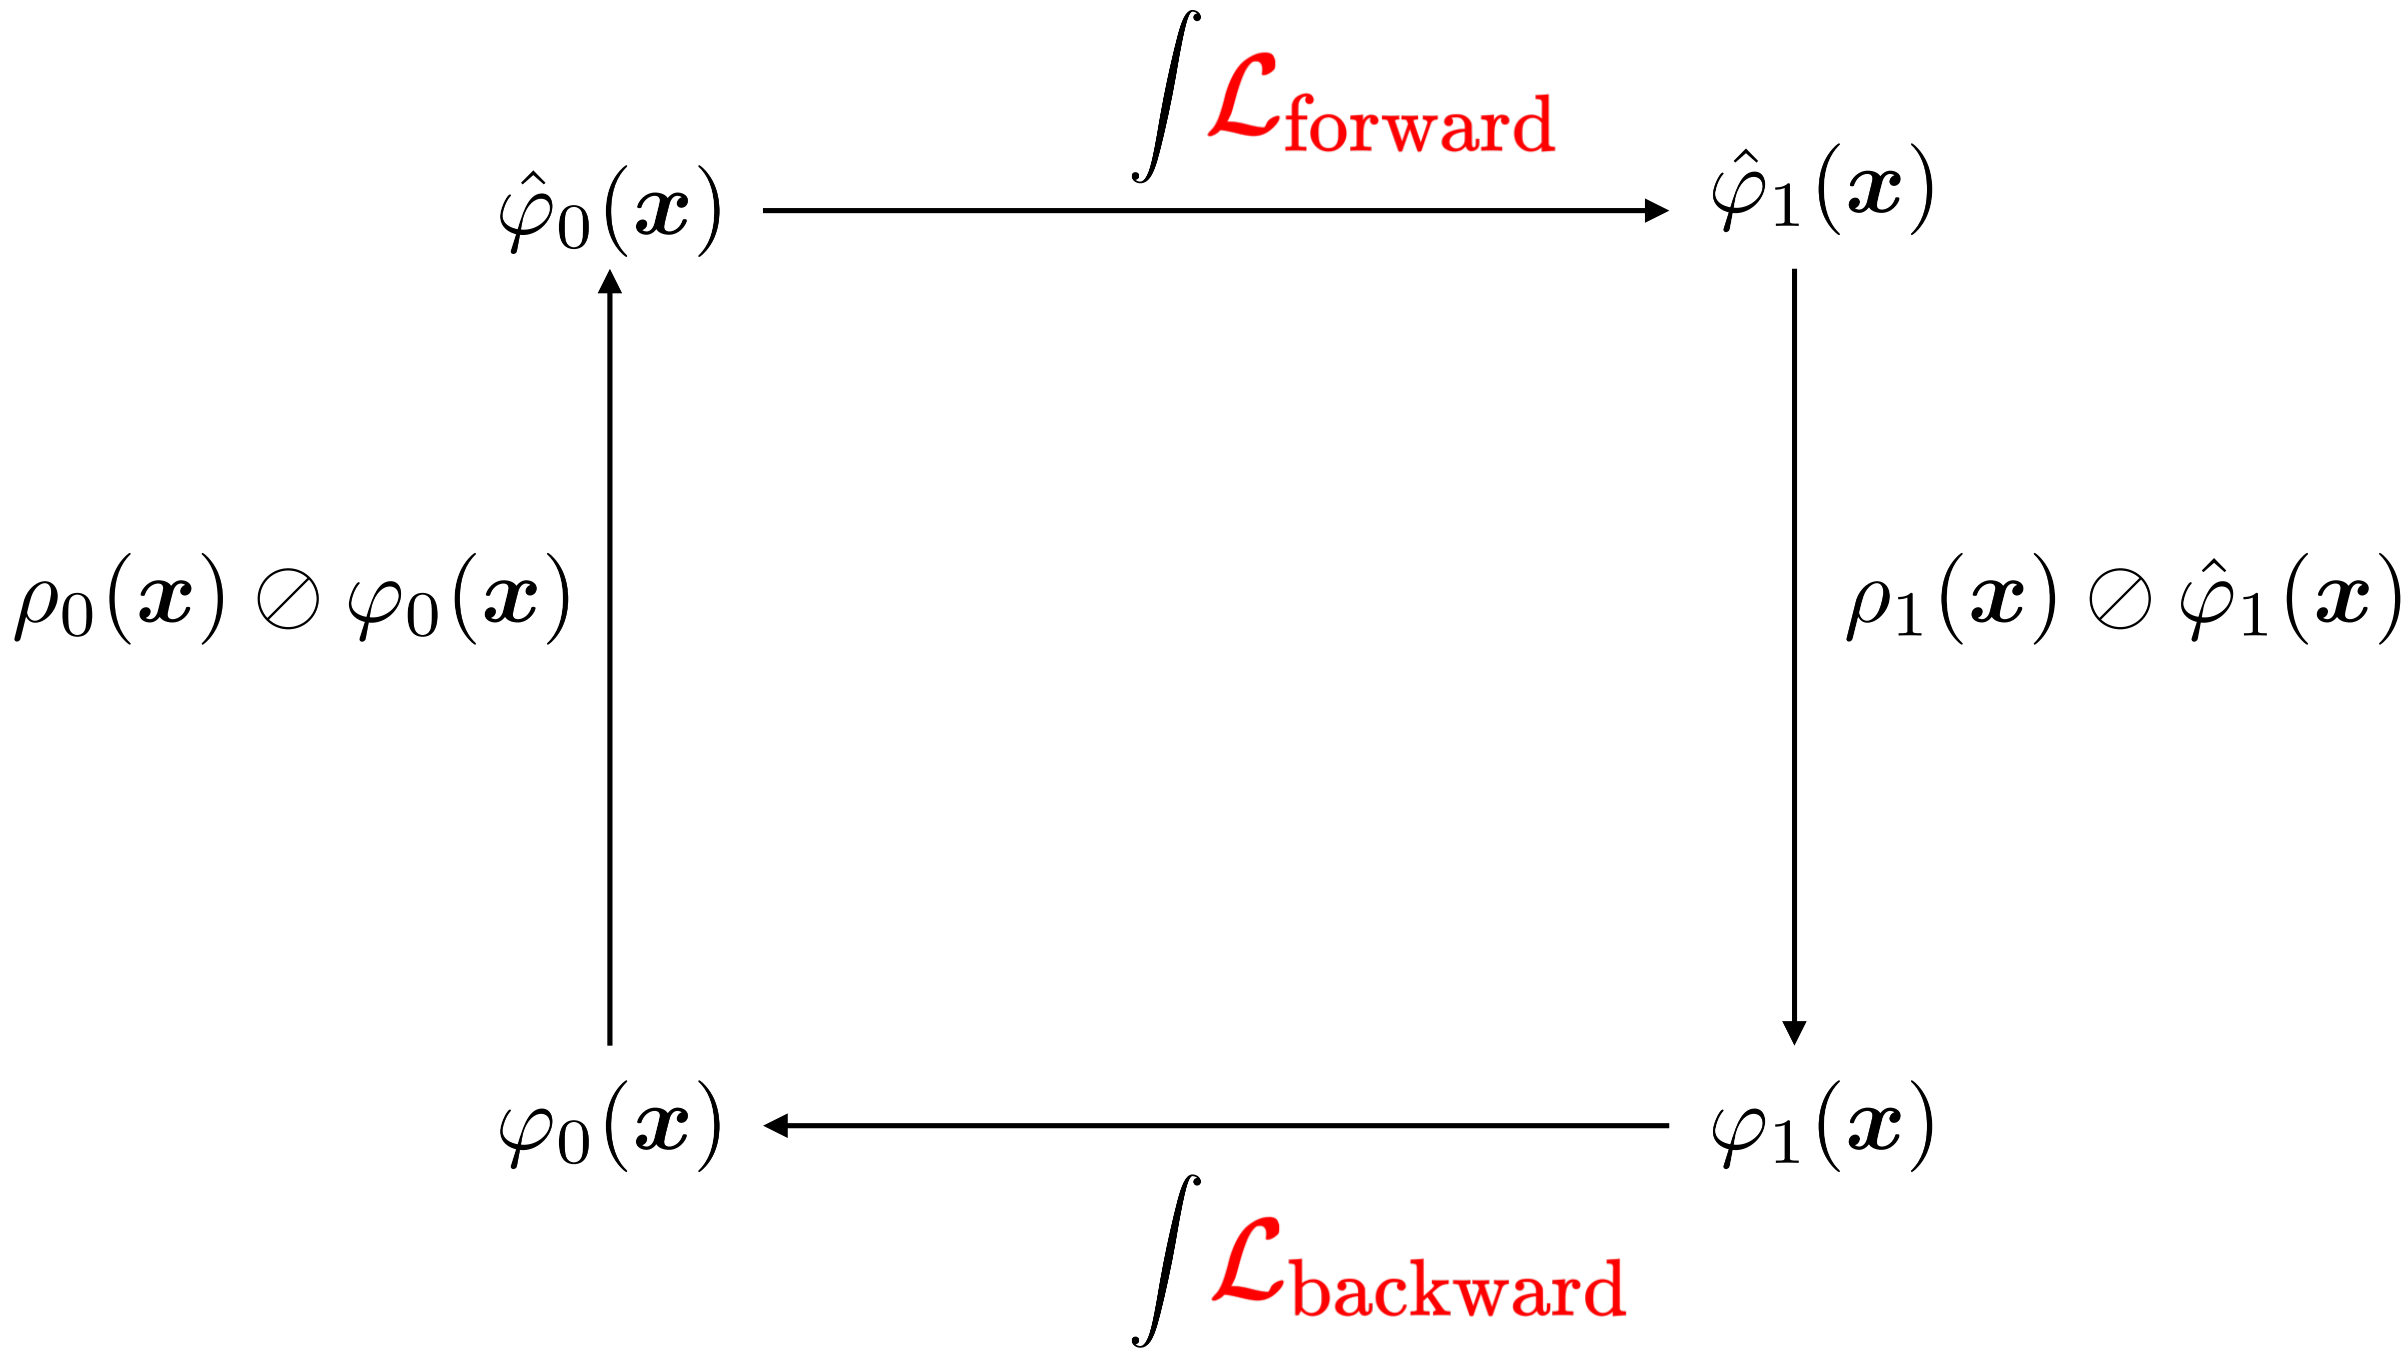

This recursion is contractive in the Hilbert metric!!

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

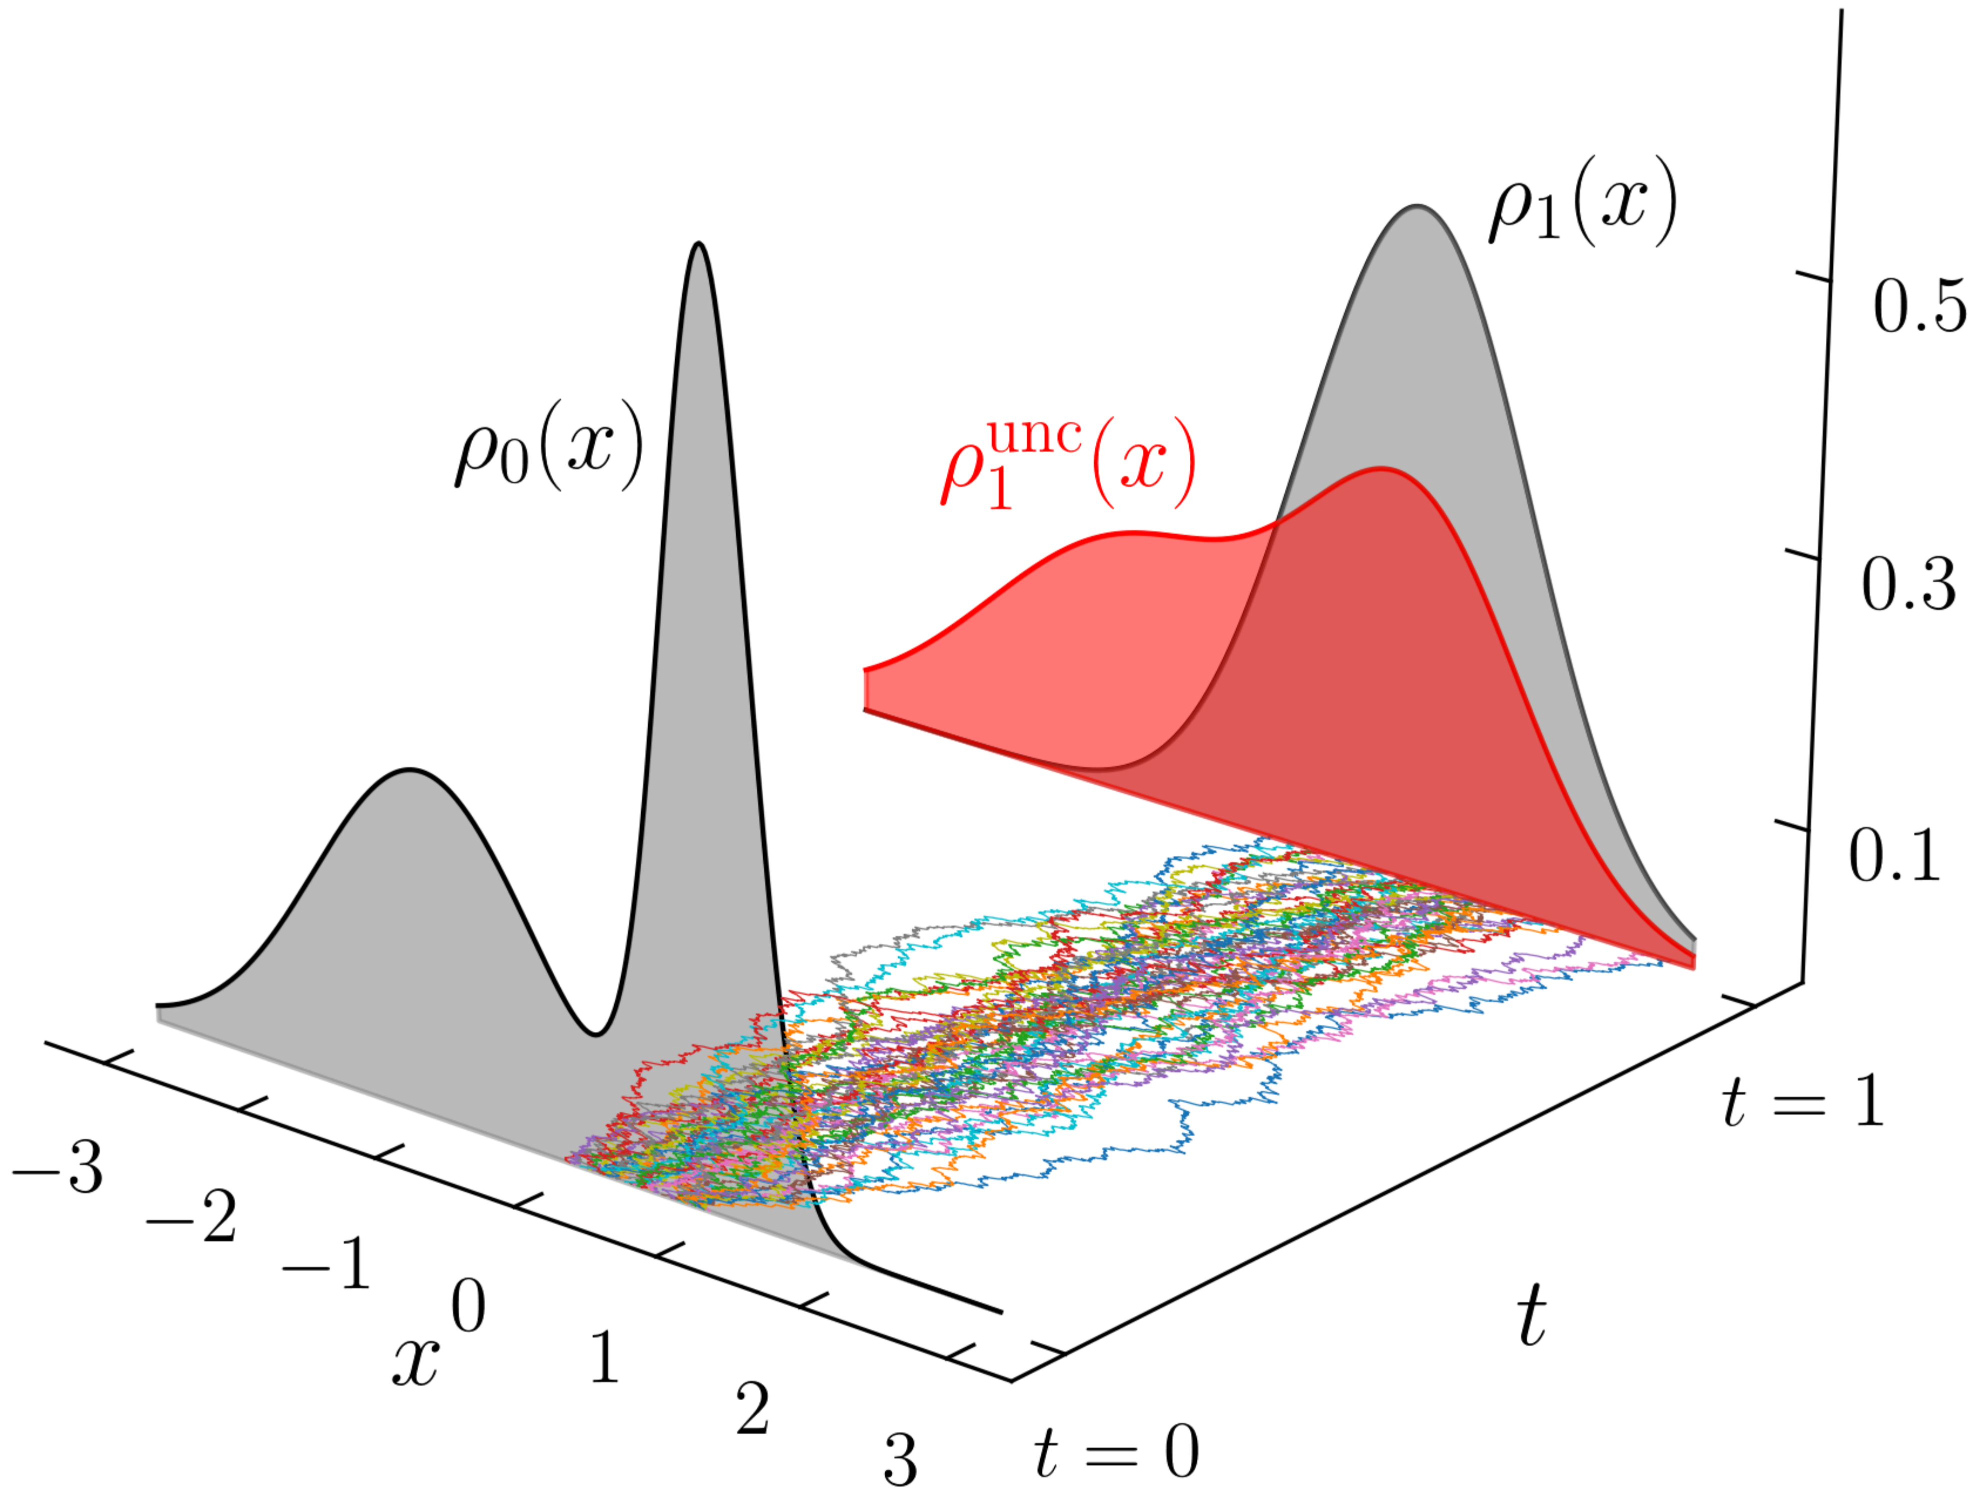

Zero prior dynamics

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

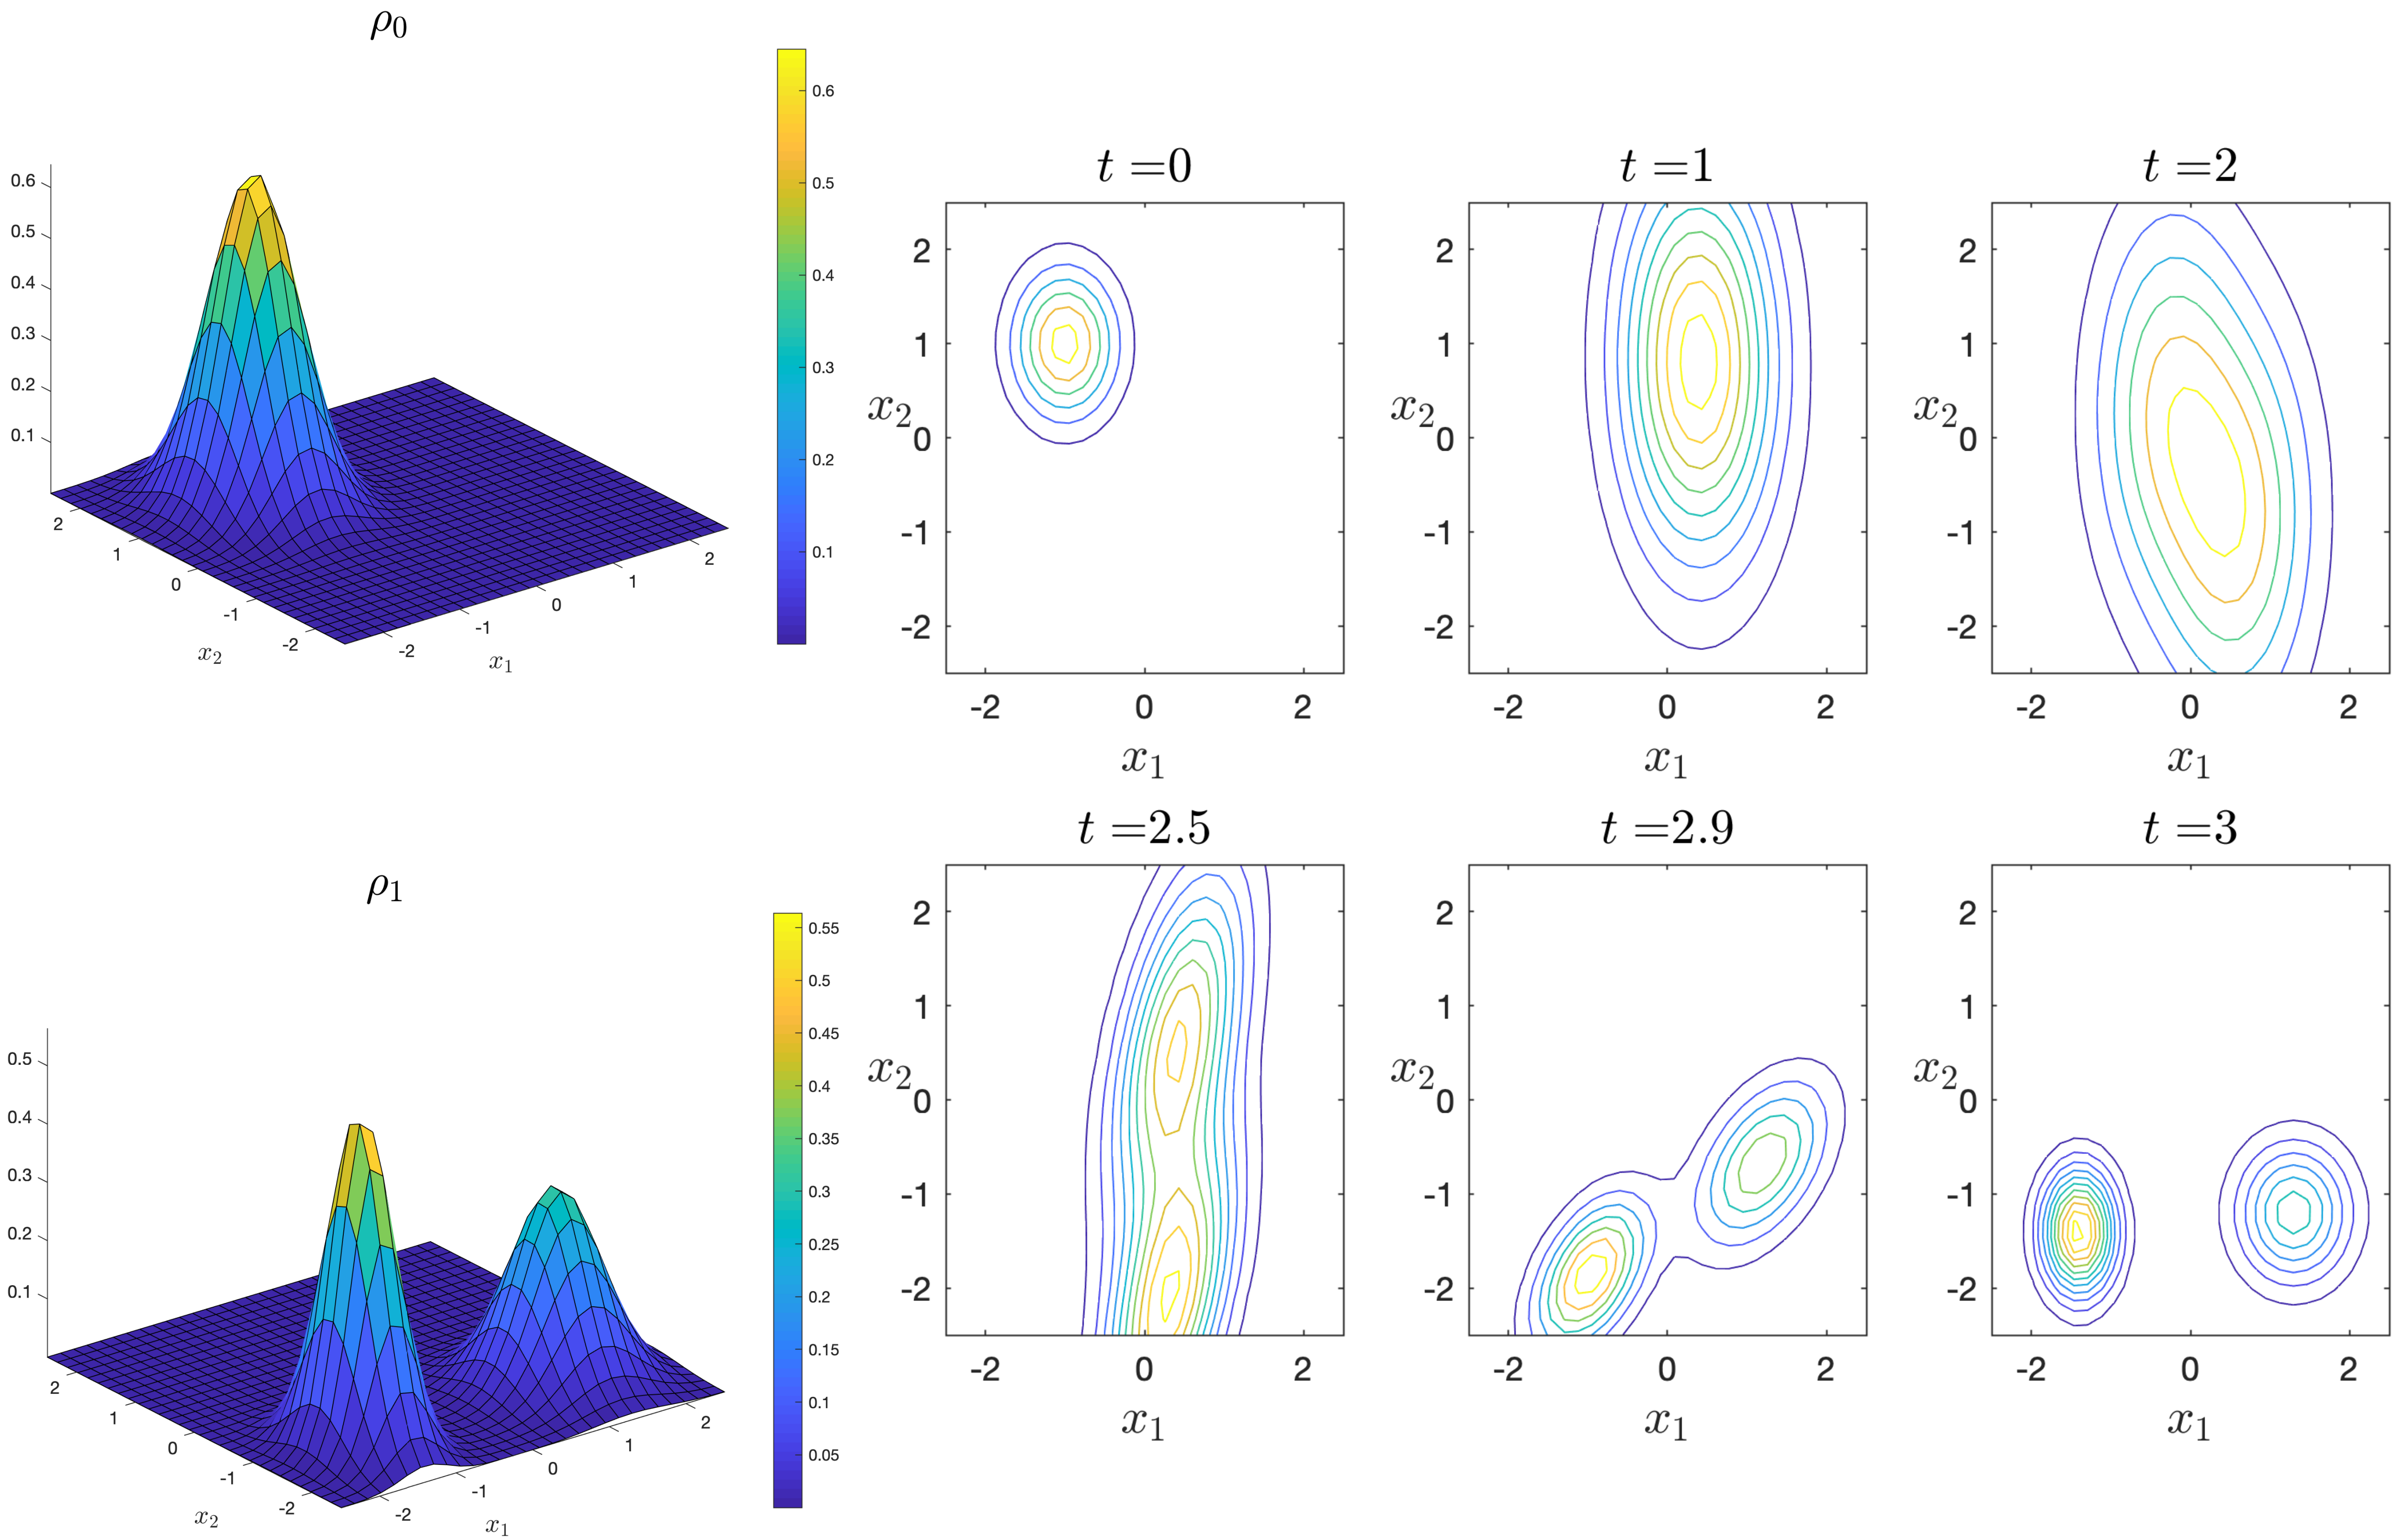

Linear prior dynamics

In general …

Need (uncontrolled) forward AND backward
Kolmogorov solvers

Bad news: Need two different solvers

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

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

***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: Nonlinear Grad. Drift

Uncontrolled joint PDF evolution:

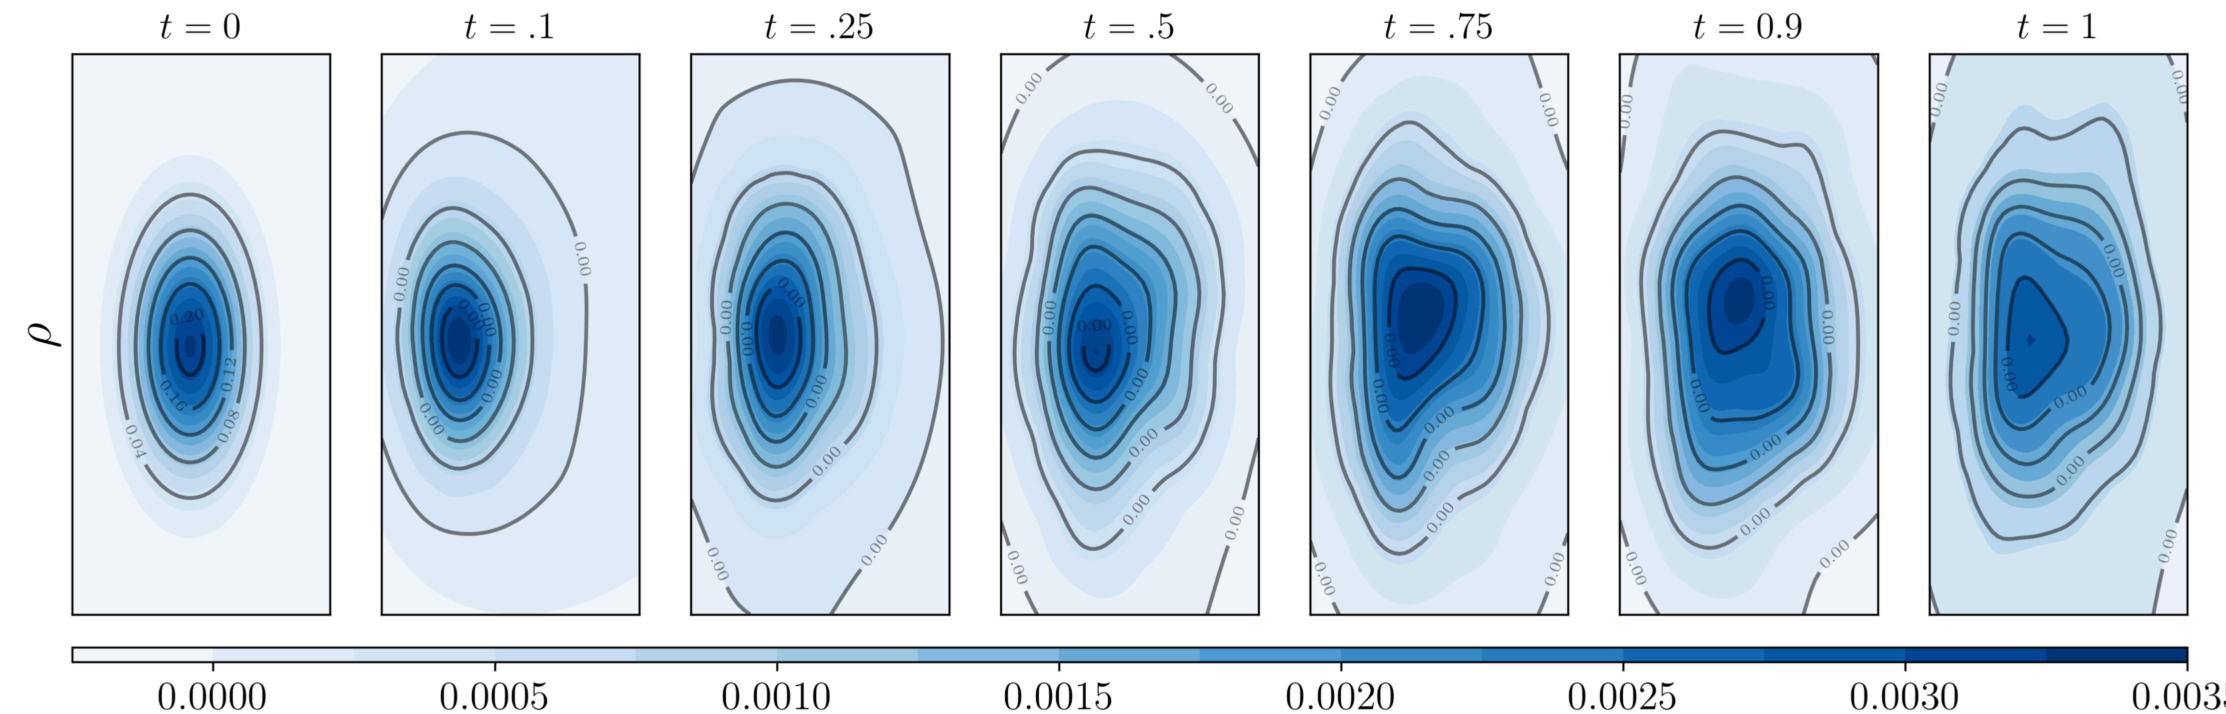

Optimal controlled joint PDF evolution:

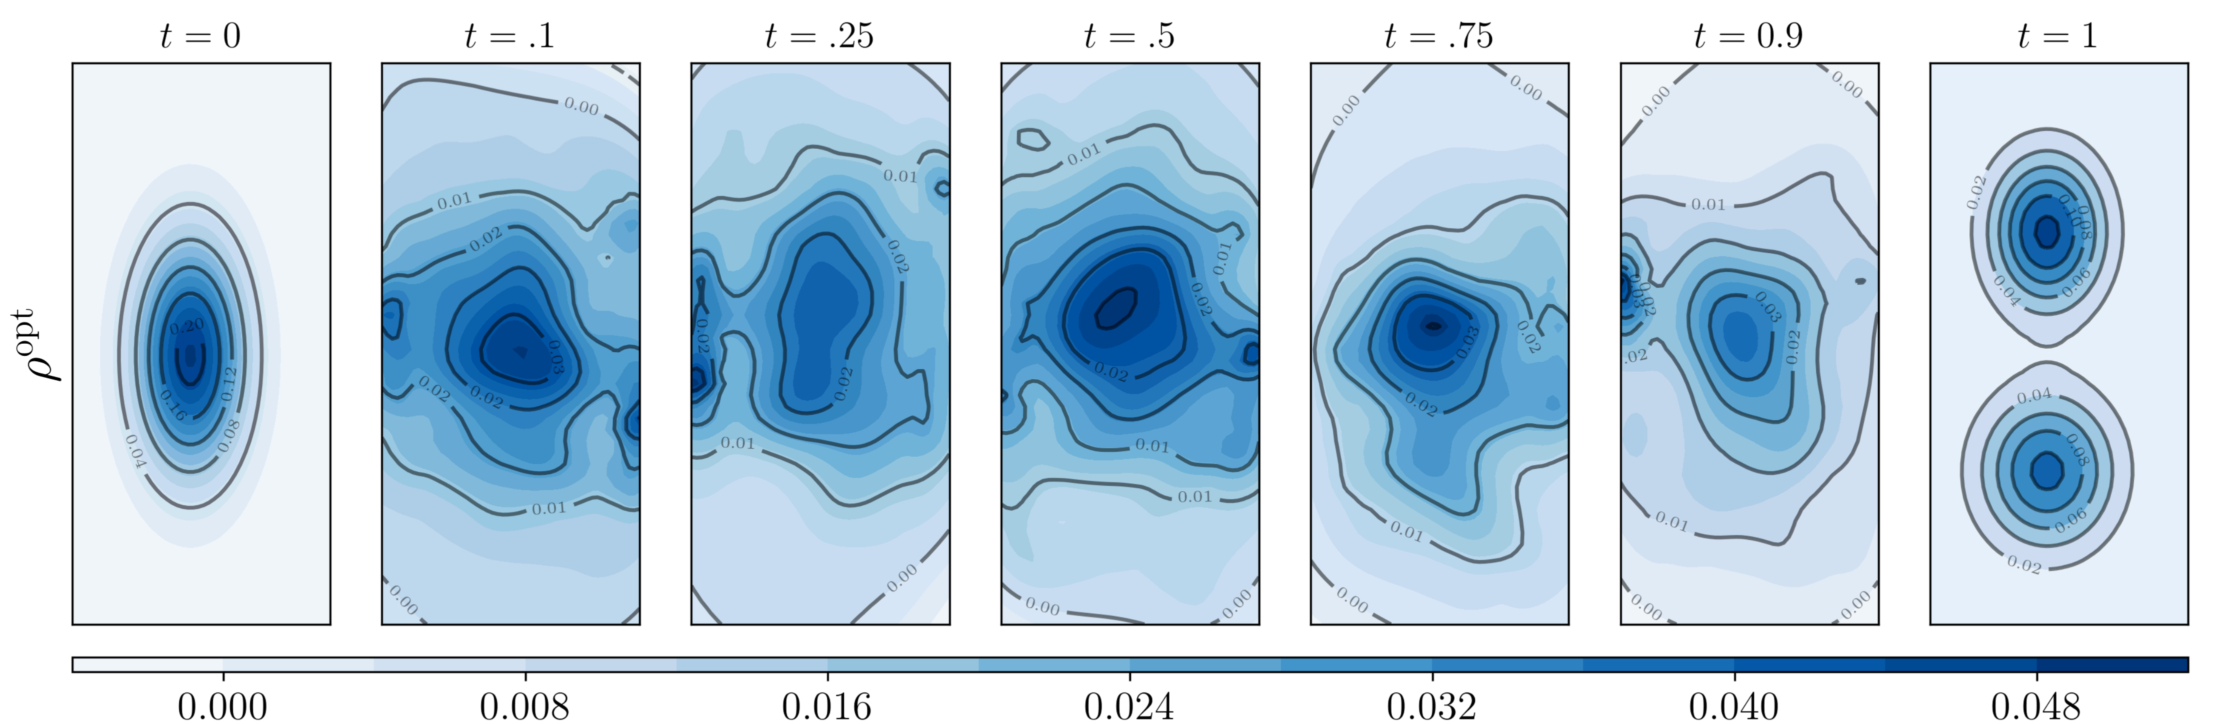

Feedback Density Control: Mixed Conservative-Dissipative Drift

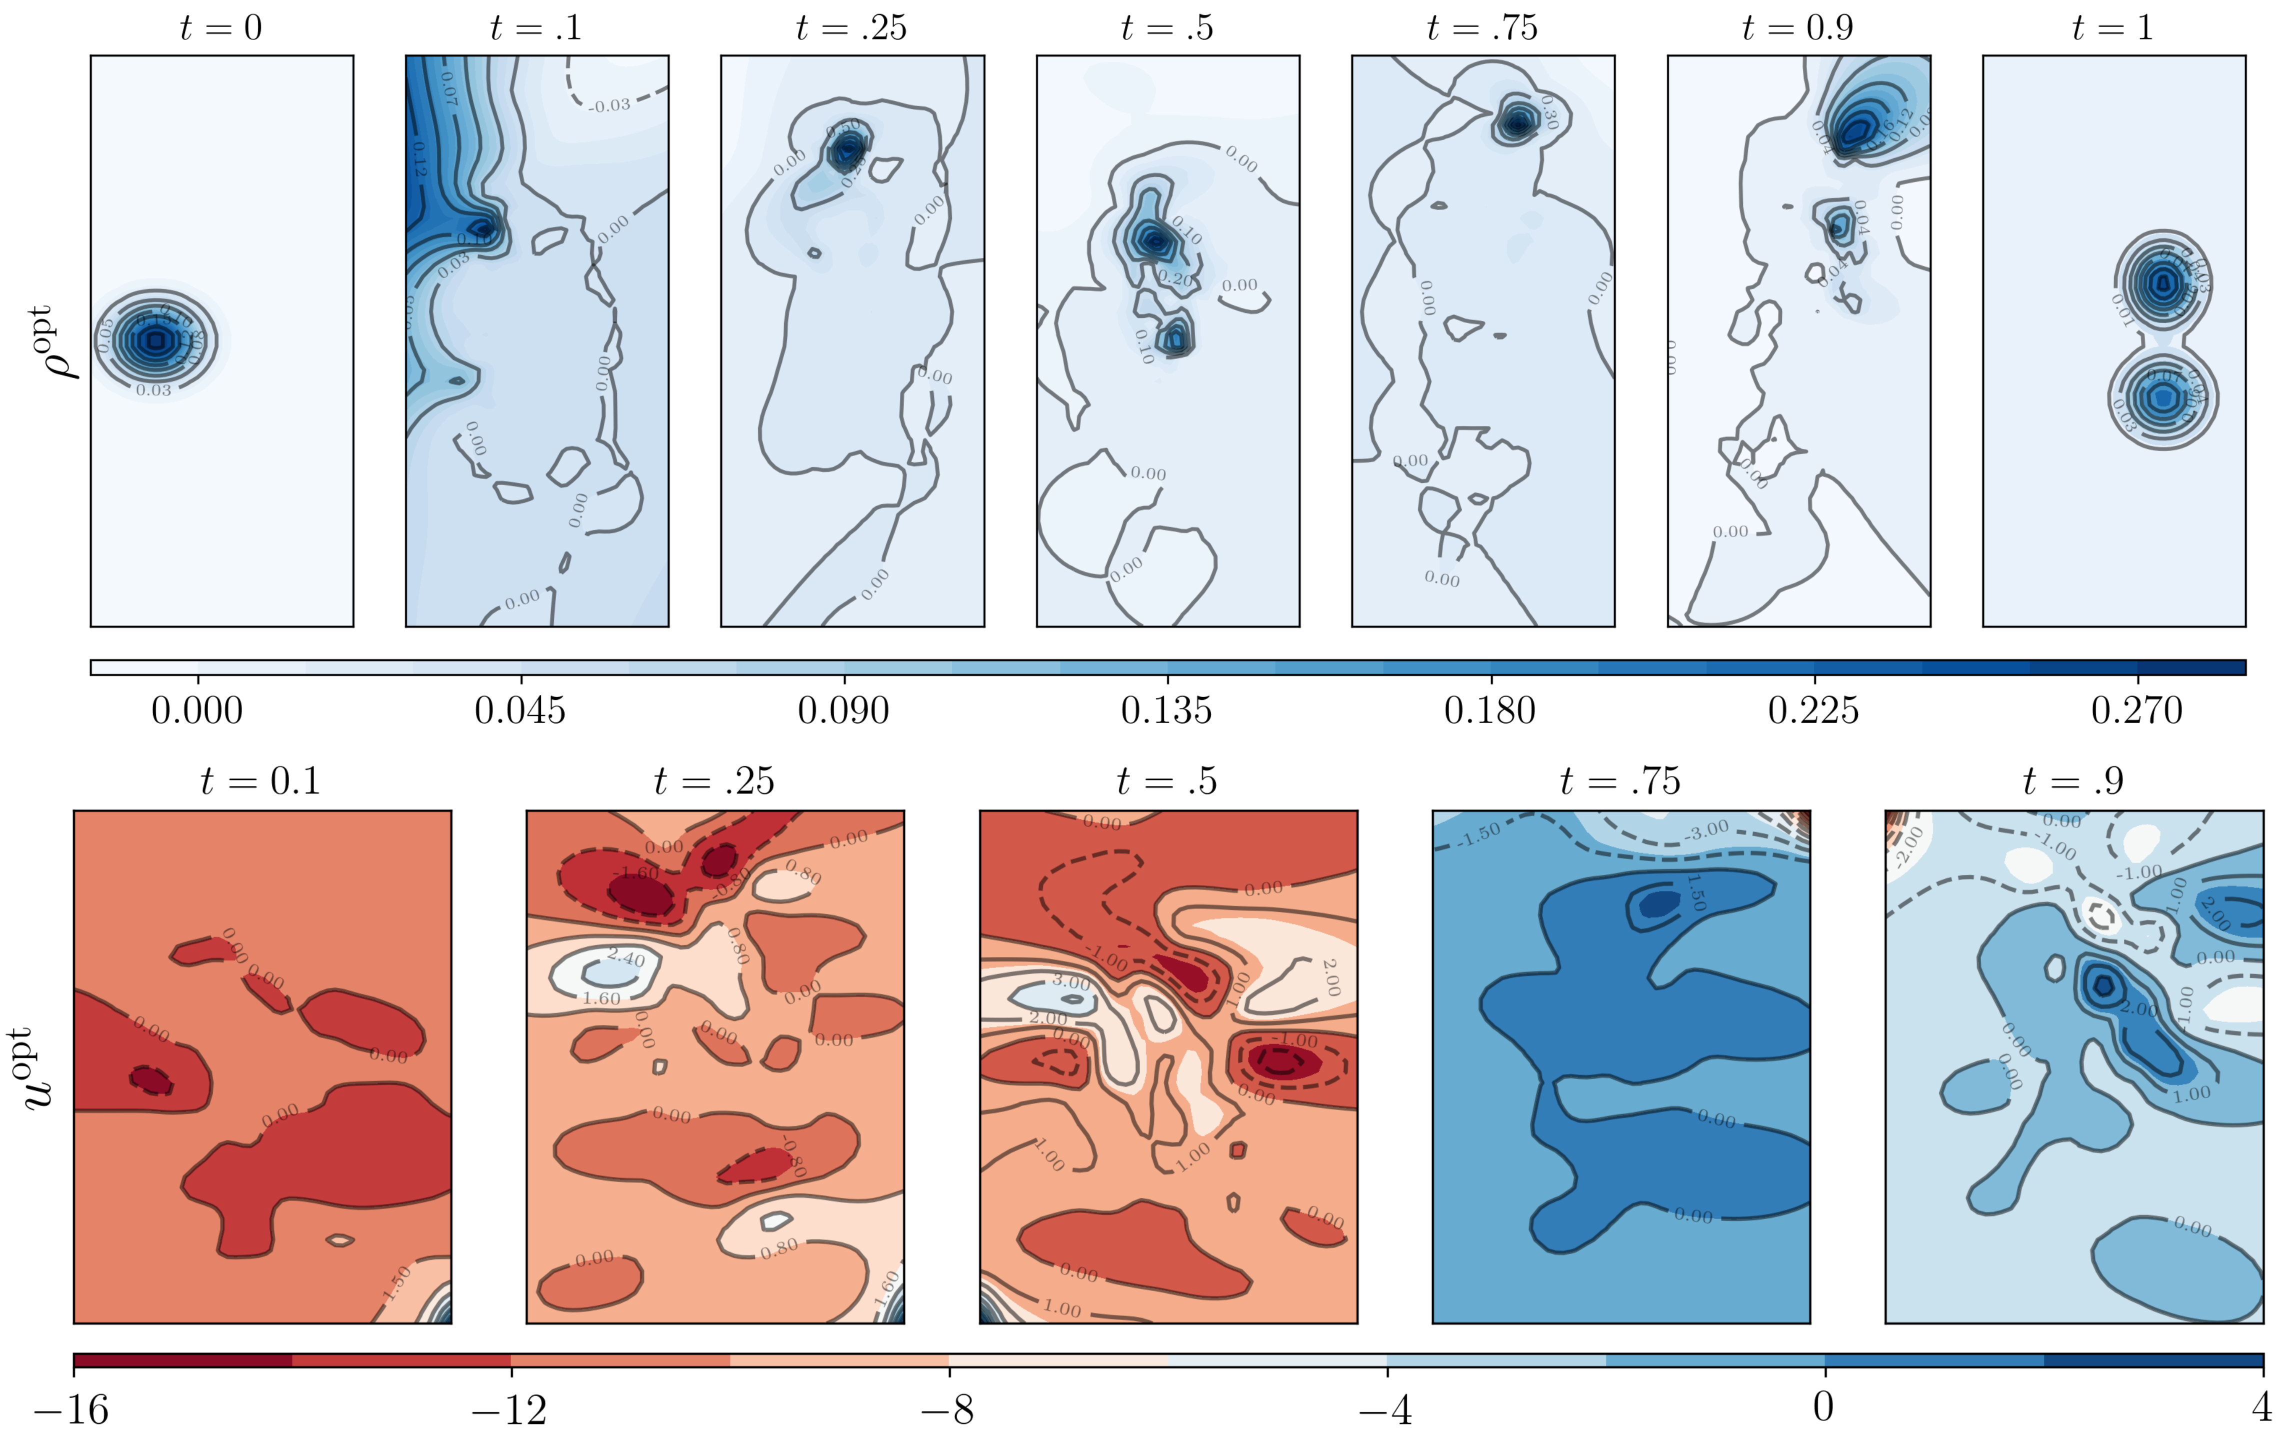

K.F. Caluya and A.H., Wasserstein proximal algorithms for the Schrödinger bridge problem: density control with nonlinear drift, *IEEE TAC* 2021.

Nonlinear Density Steering with Deterministic Path Constraints

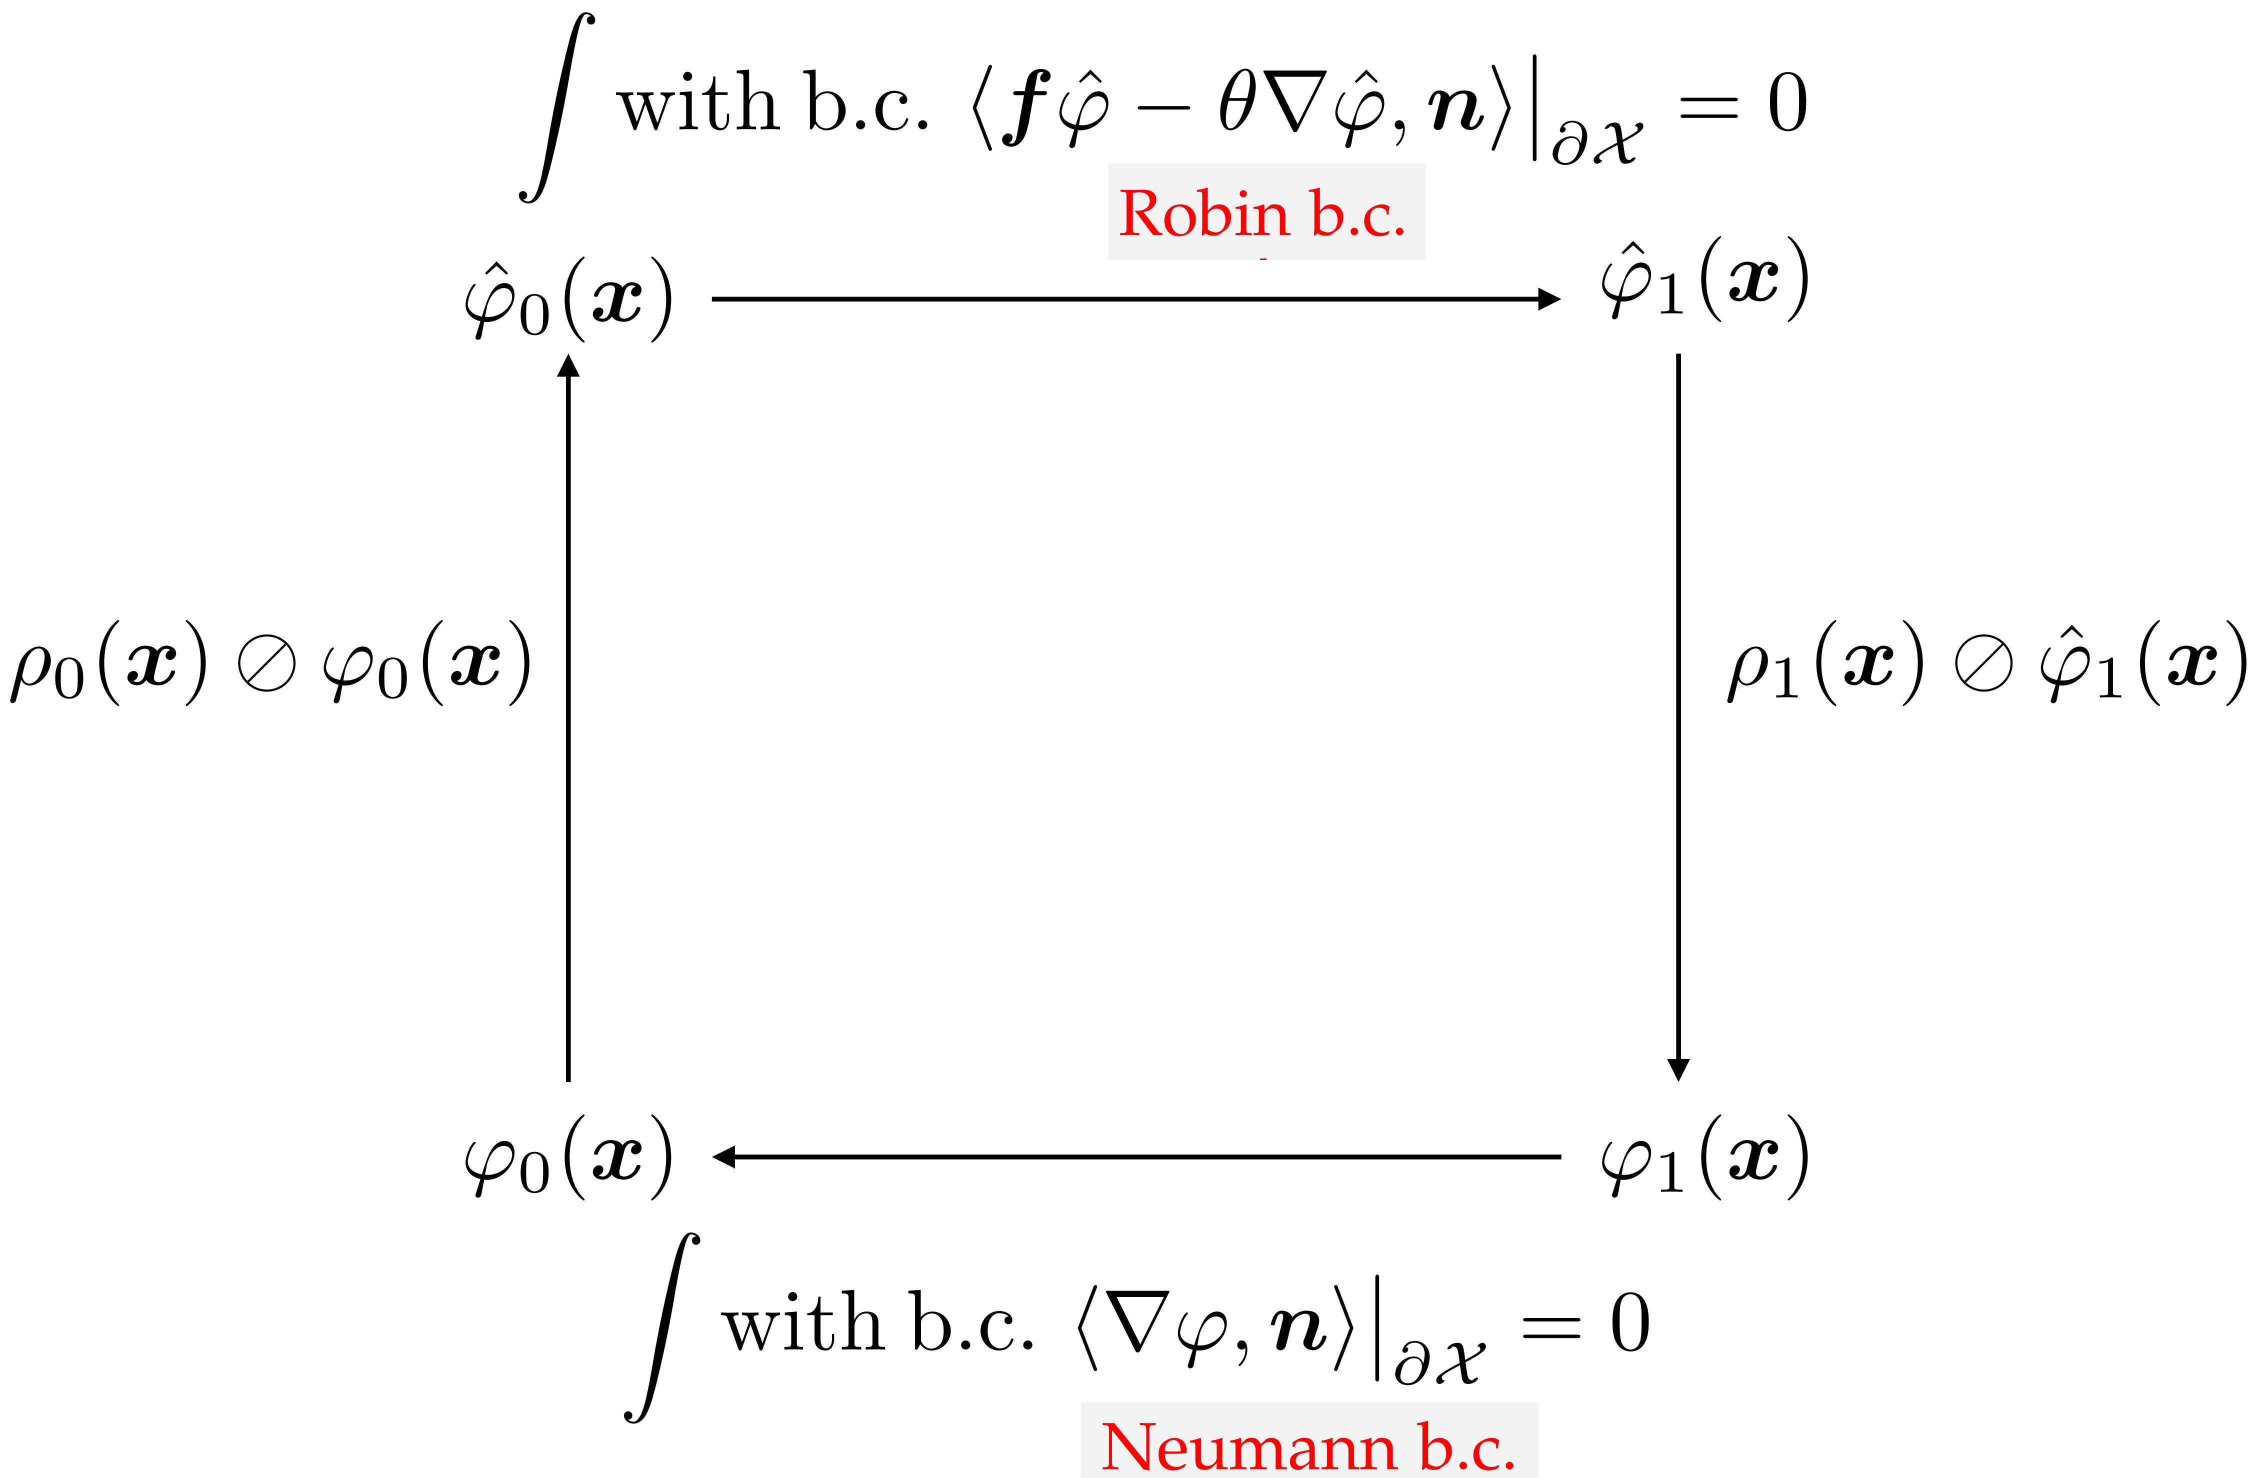

— K.F. Caluya, and A.H., Reflected Schrödinger Bridge: Density Control with Path Constraints, *ACC 2021*.

Nonlinear Density Steering with Deterministic Path Constraints

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

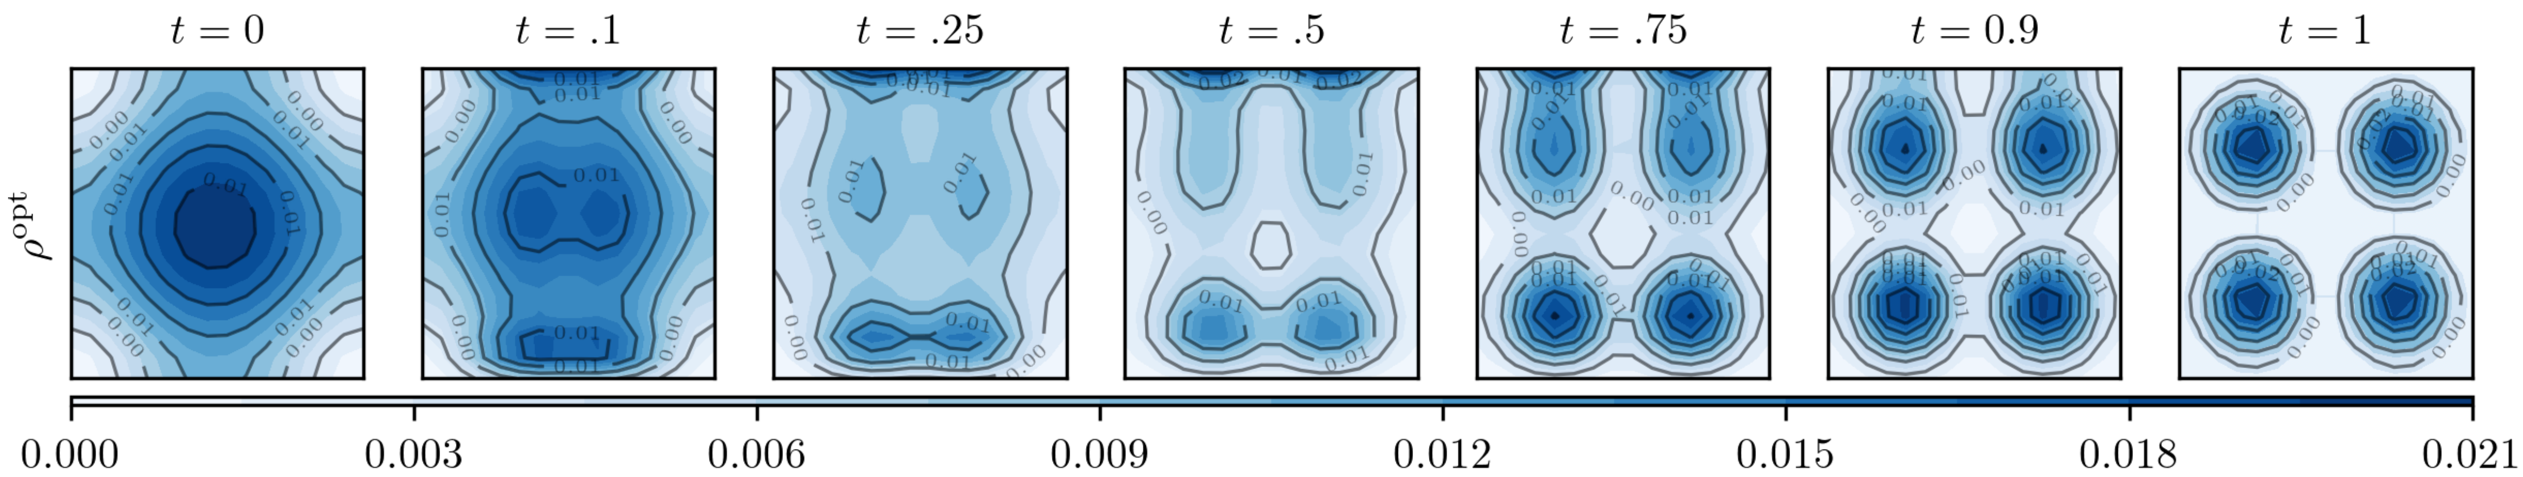

Uncontrolled state PDFs:

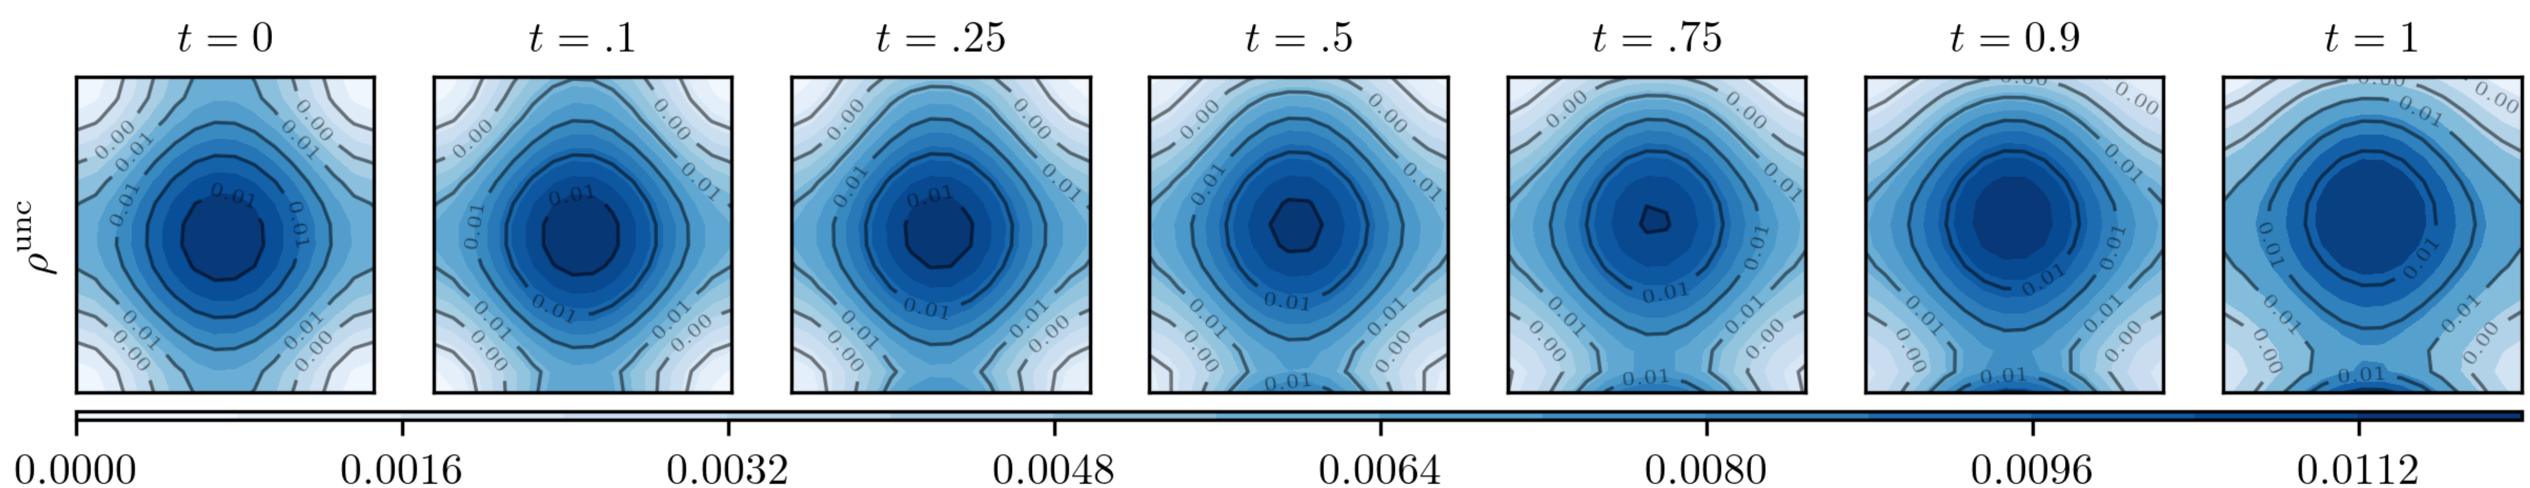

— K.F. Caluya, and A.H., Reflected Schrödinger Bridge: Density Control with Path Constraints, *ACC 2021*.

Application: Multi-lane Automated Driving

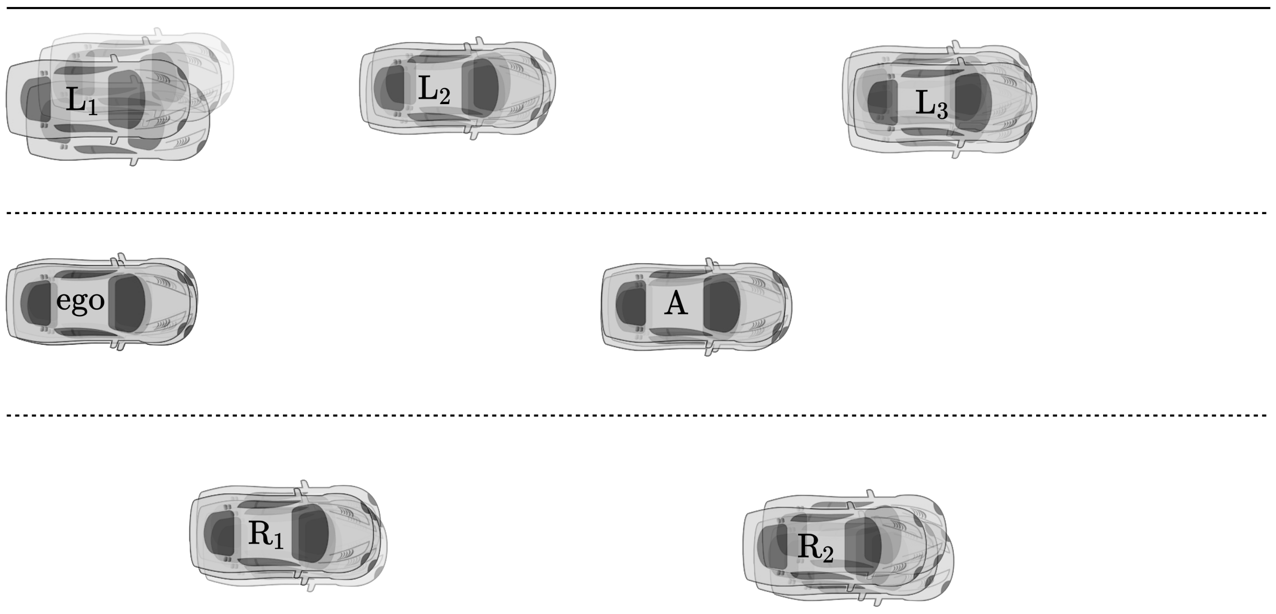

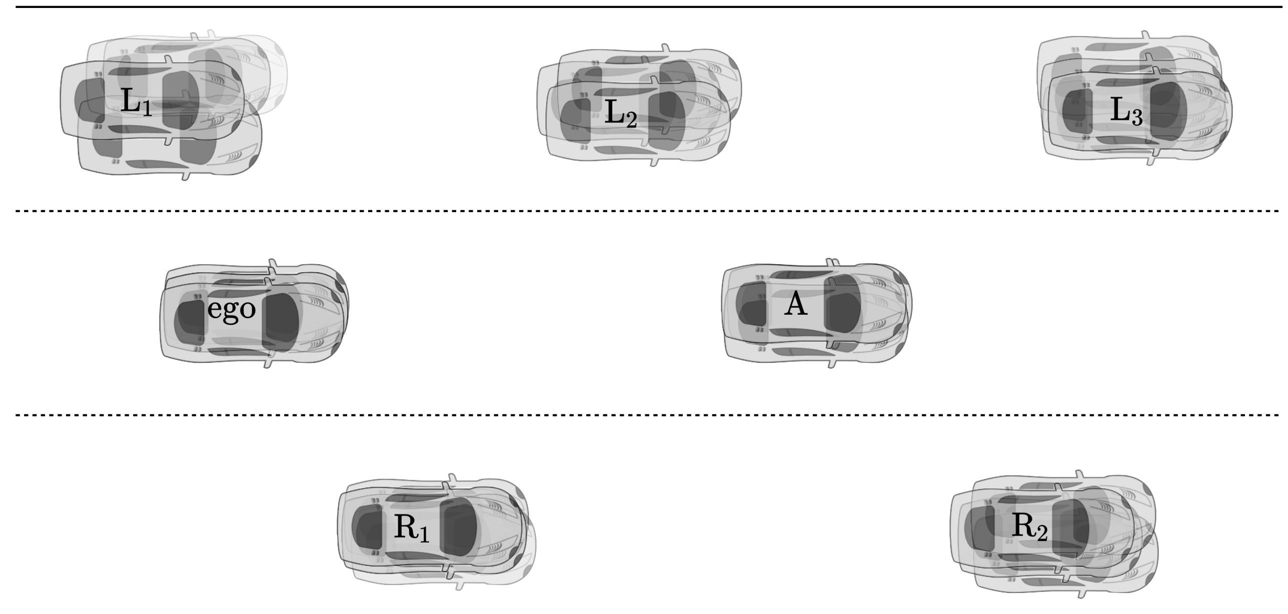

t_0 → t_1

x marginals

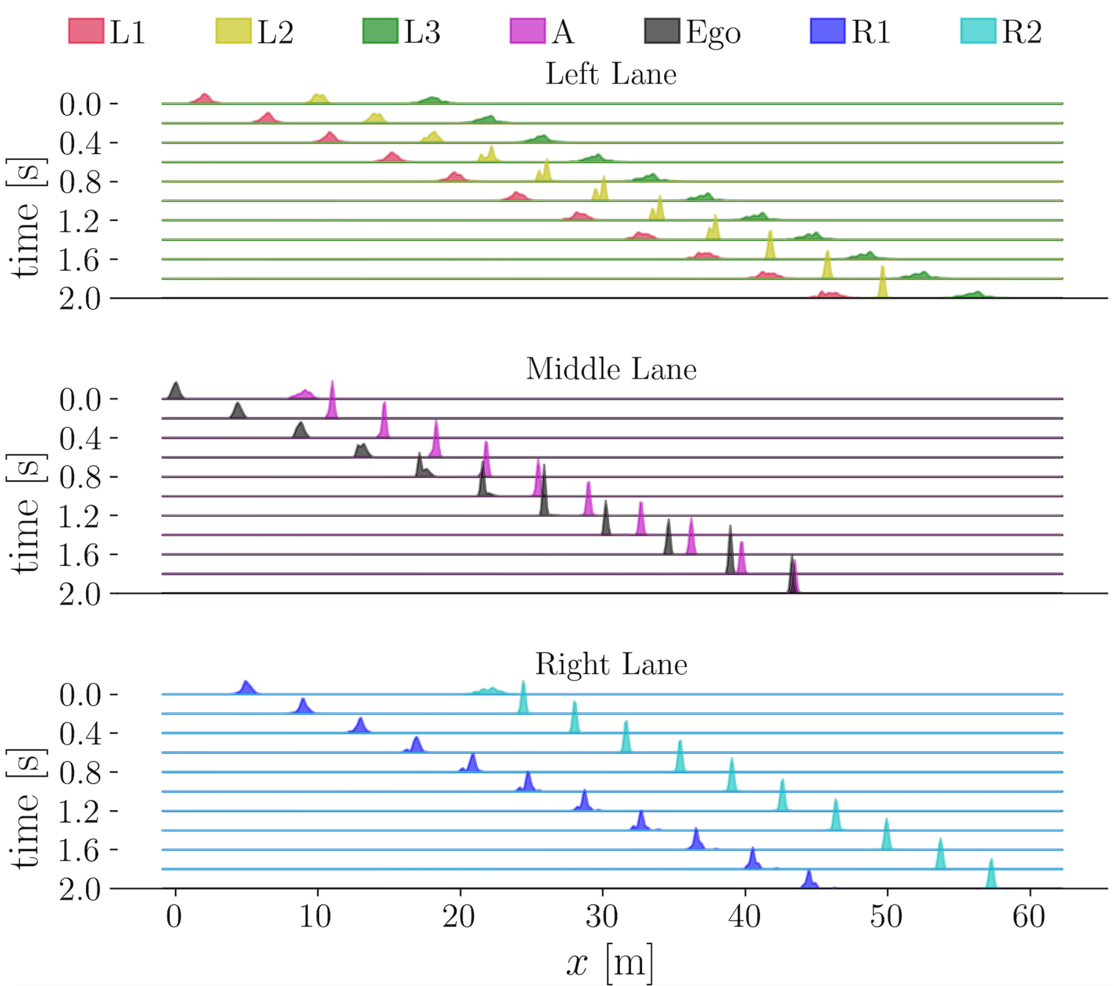

y marginals

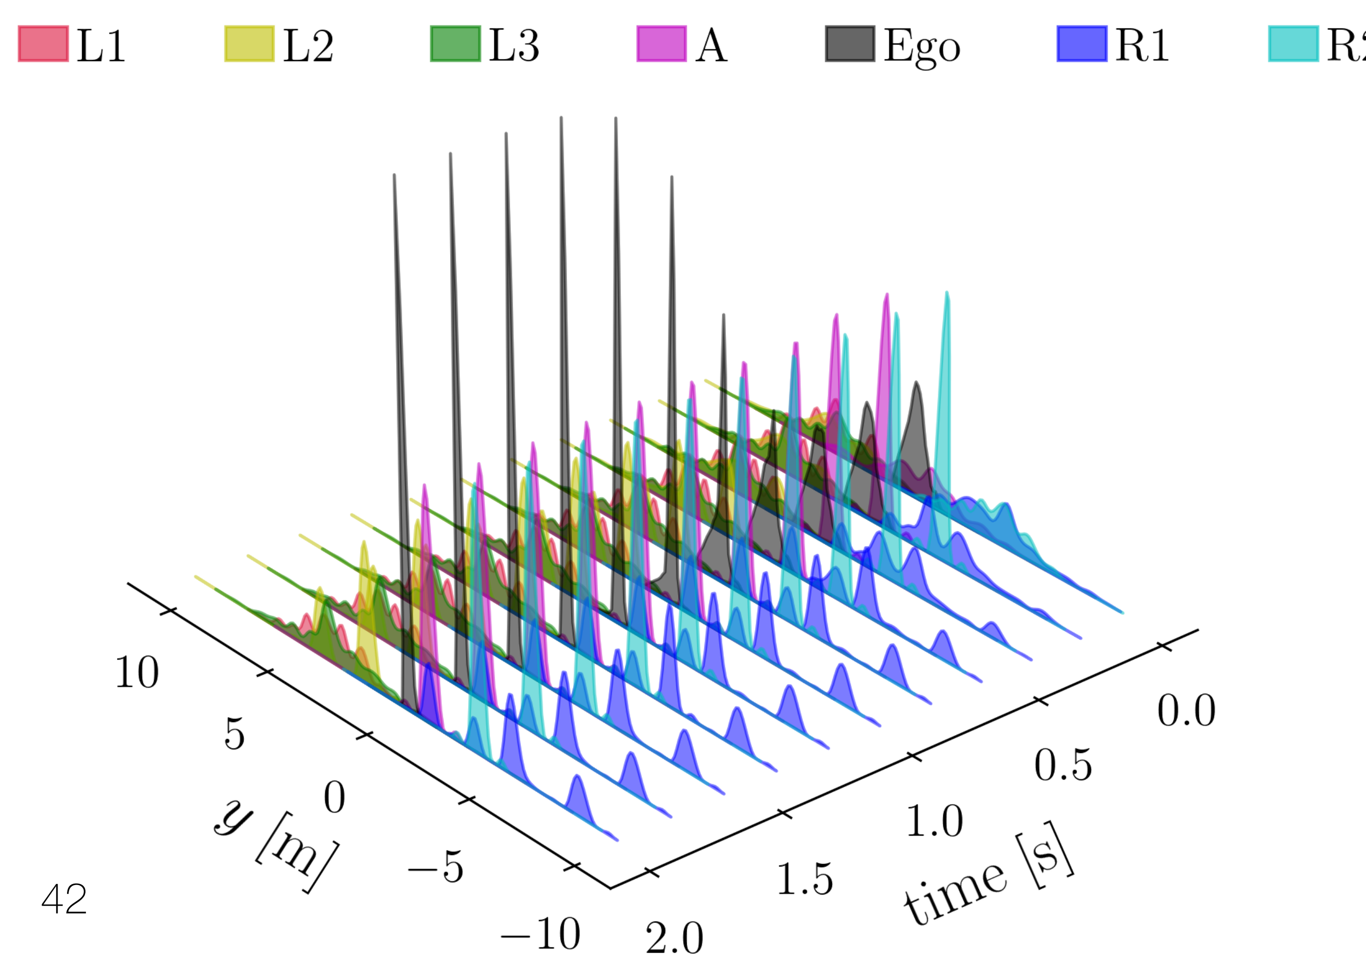

Application: Multi-lane Automated Driving

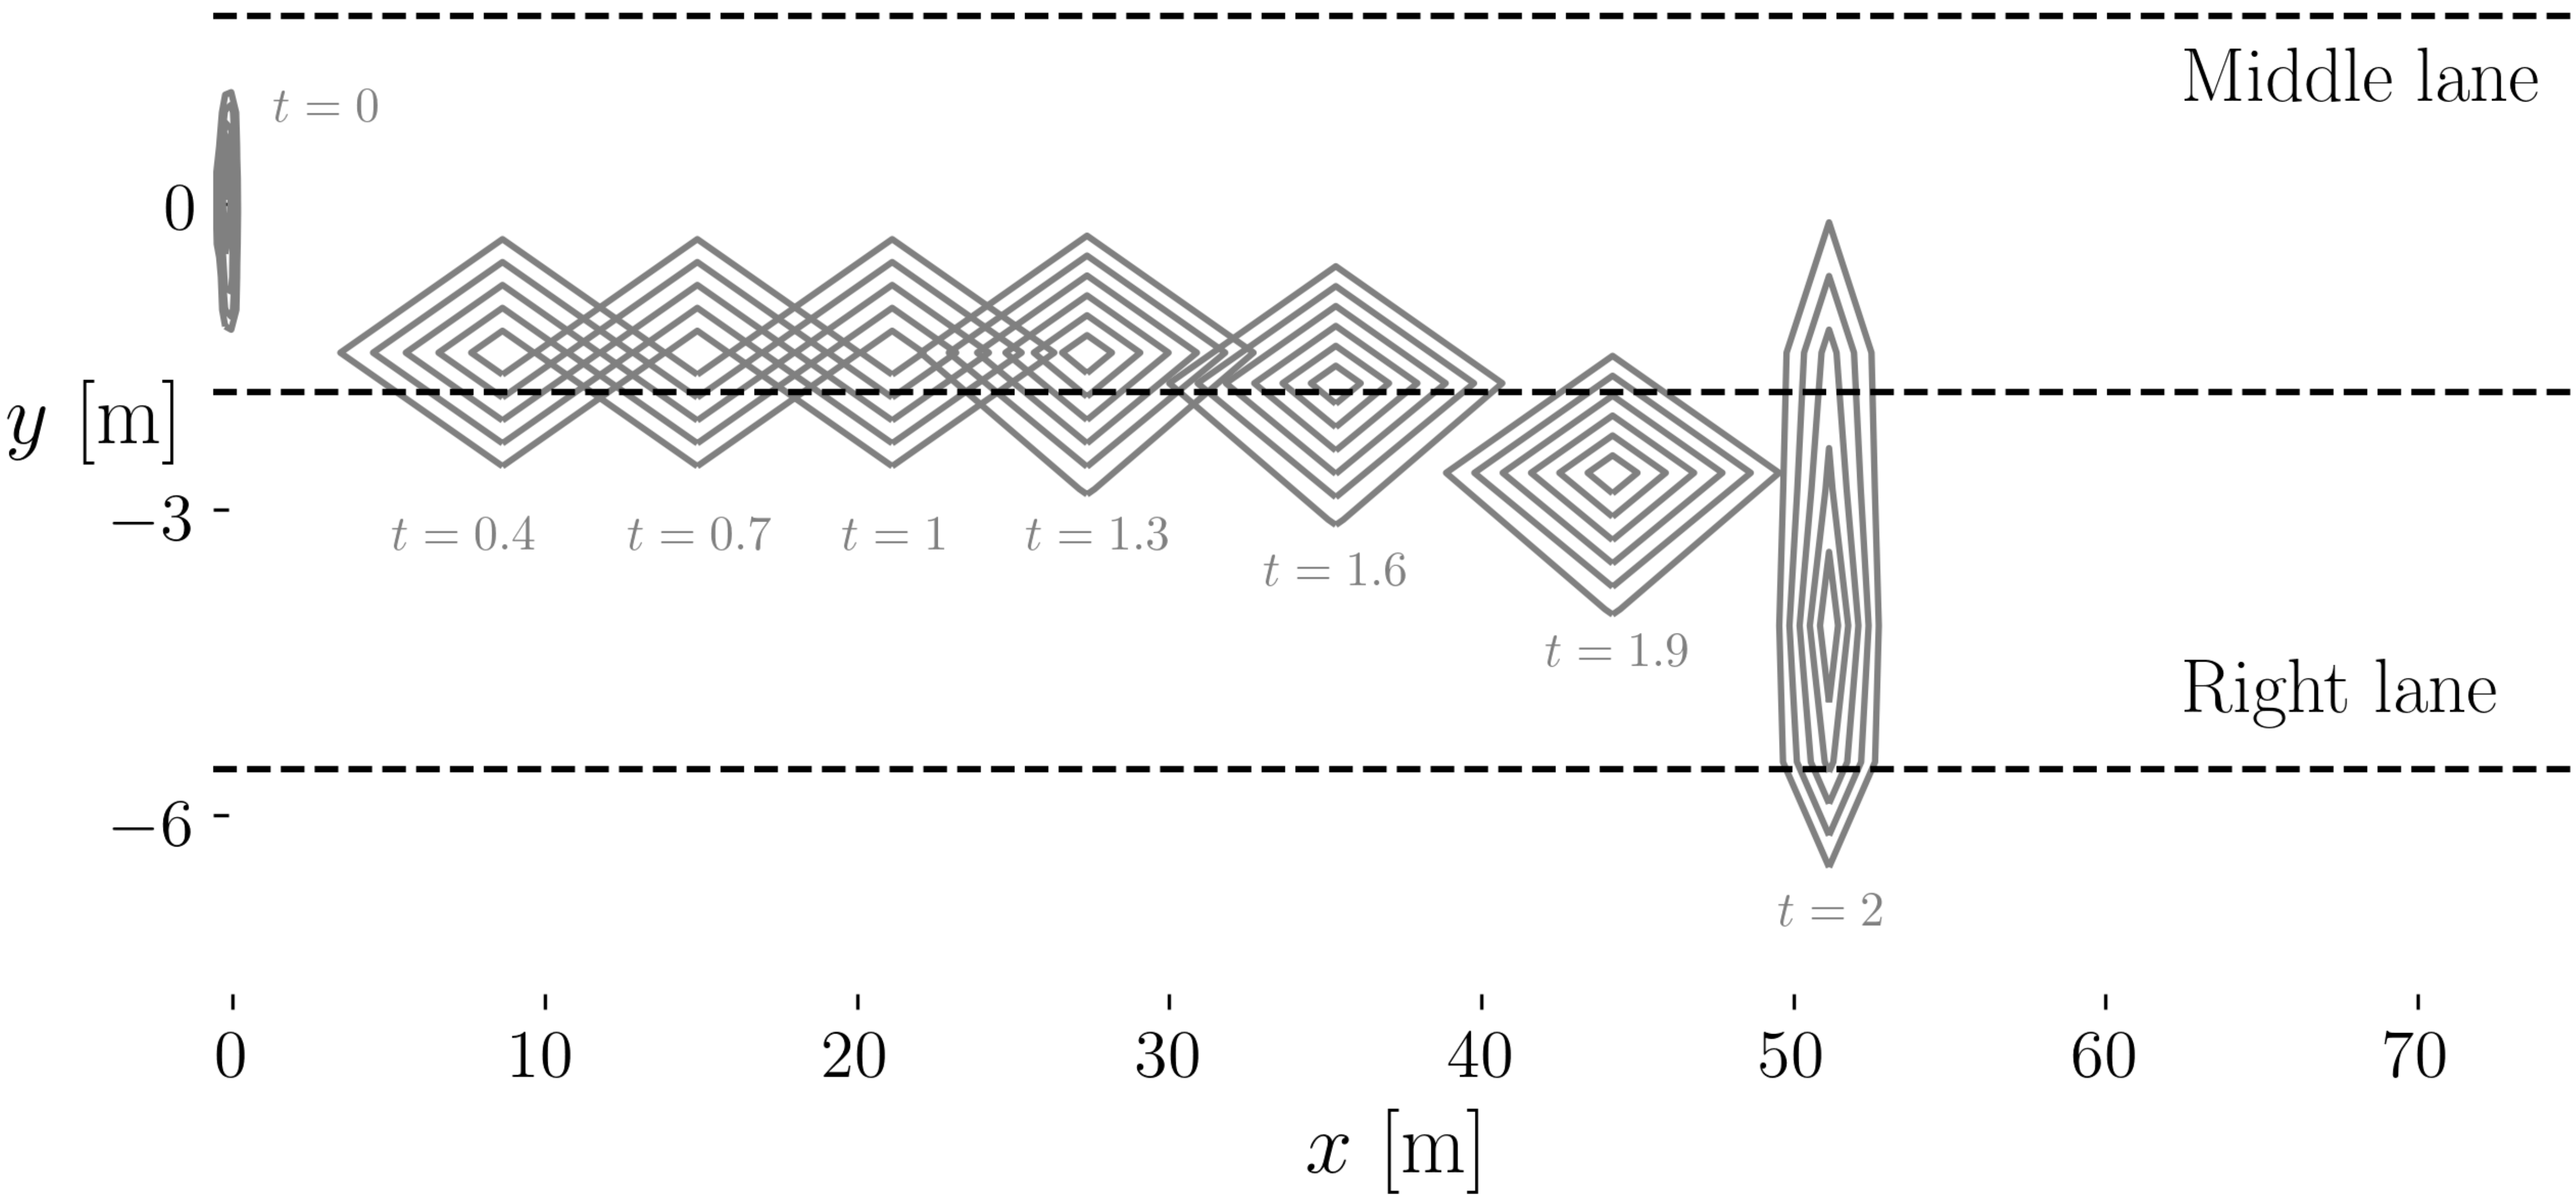

— S. Haddad, K.F. Caluya, A.H., and B. Singh, Prediction and Optimal Feedback Steering of Probability Density Functions for Safe Automated Driving, *IEEE Control Systems Letters, 2020*.

— 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: Controlled SA

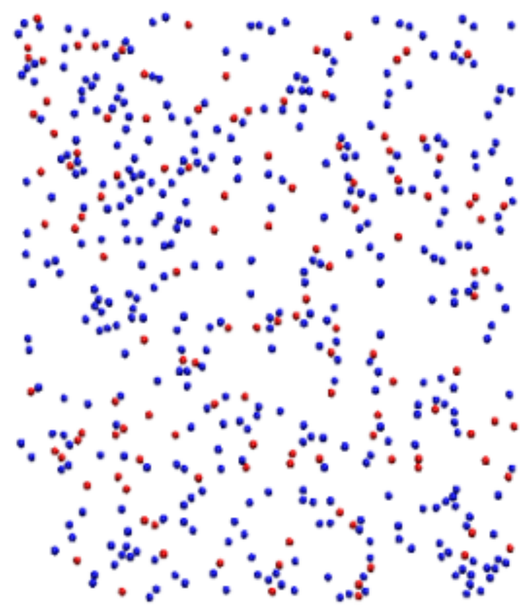

Dispersed particles

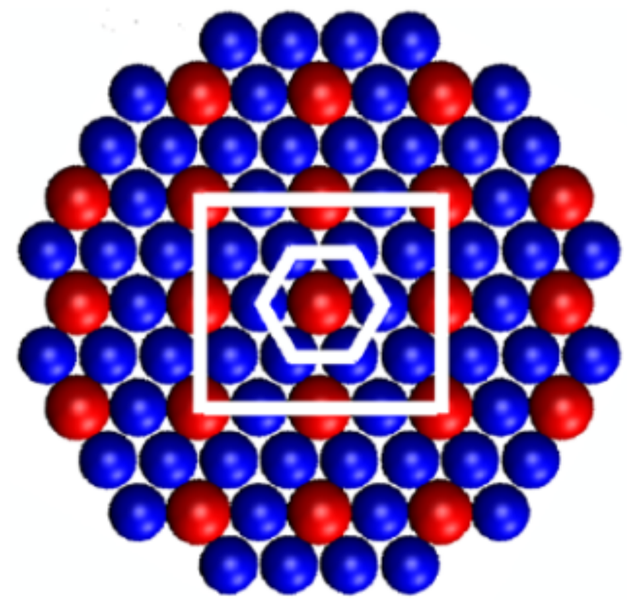

Ordered structure

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

Average number of hexagonally close packed neighboring particles in 2D assembly ⇝ 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 $t = T \equiv 200\,\mathrm{s}$

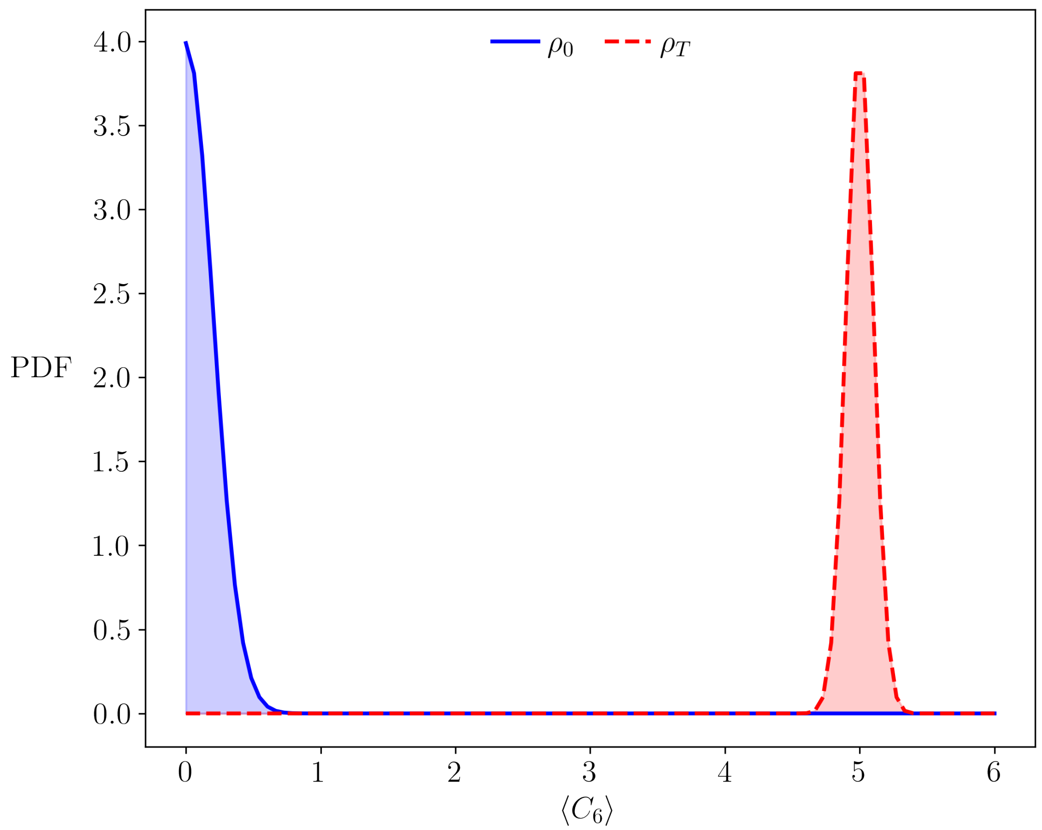

Typical prescribed finite horizon for controlled self-assembly

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:

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

drift landscape | diffusion landscape | free energy landscape

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

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

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

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

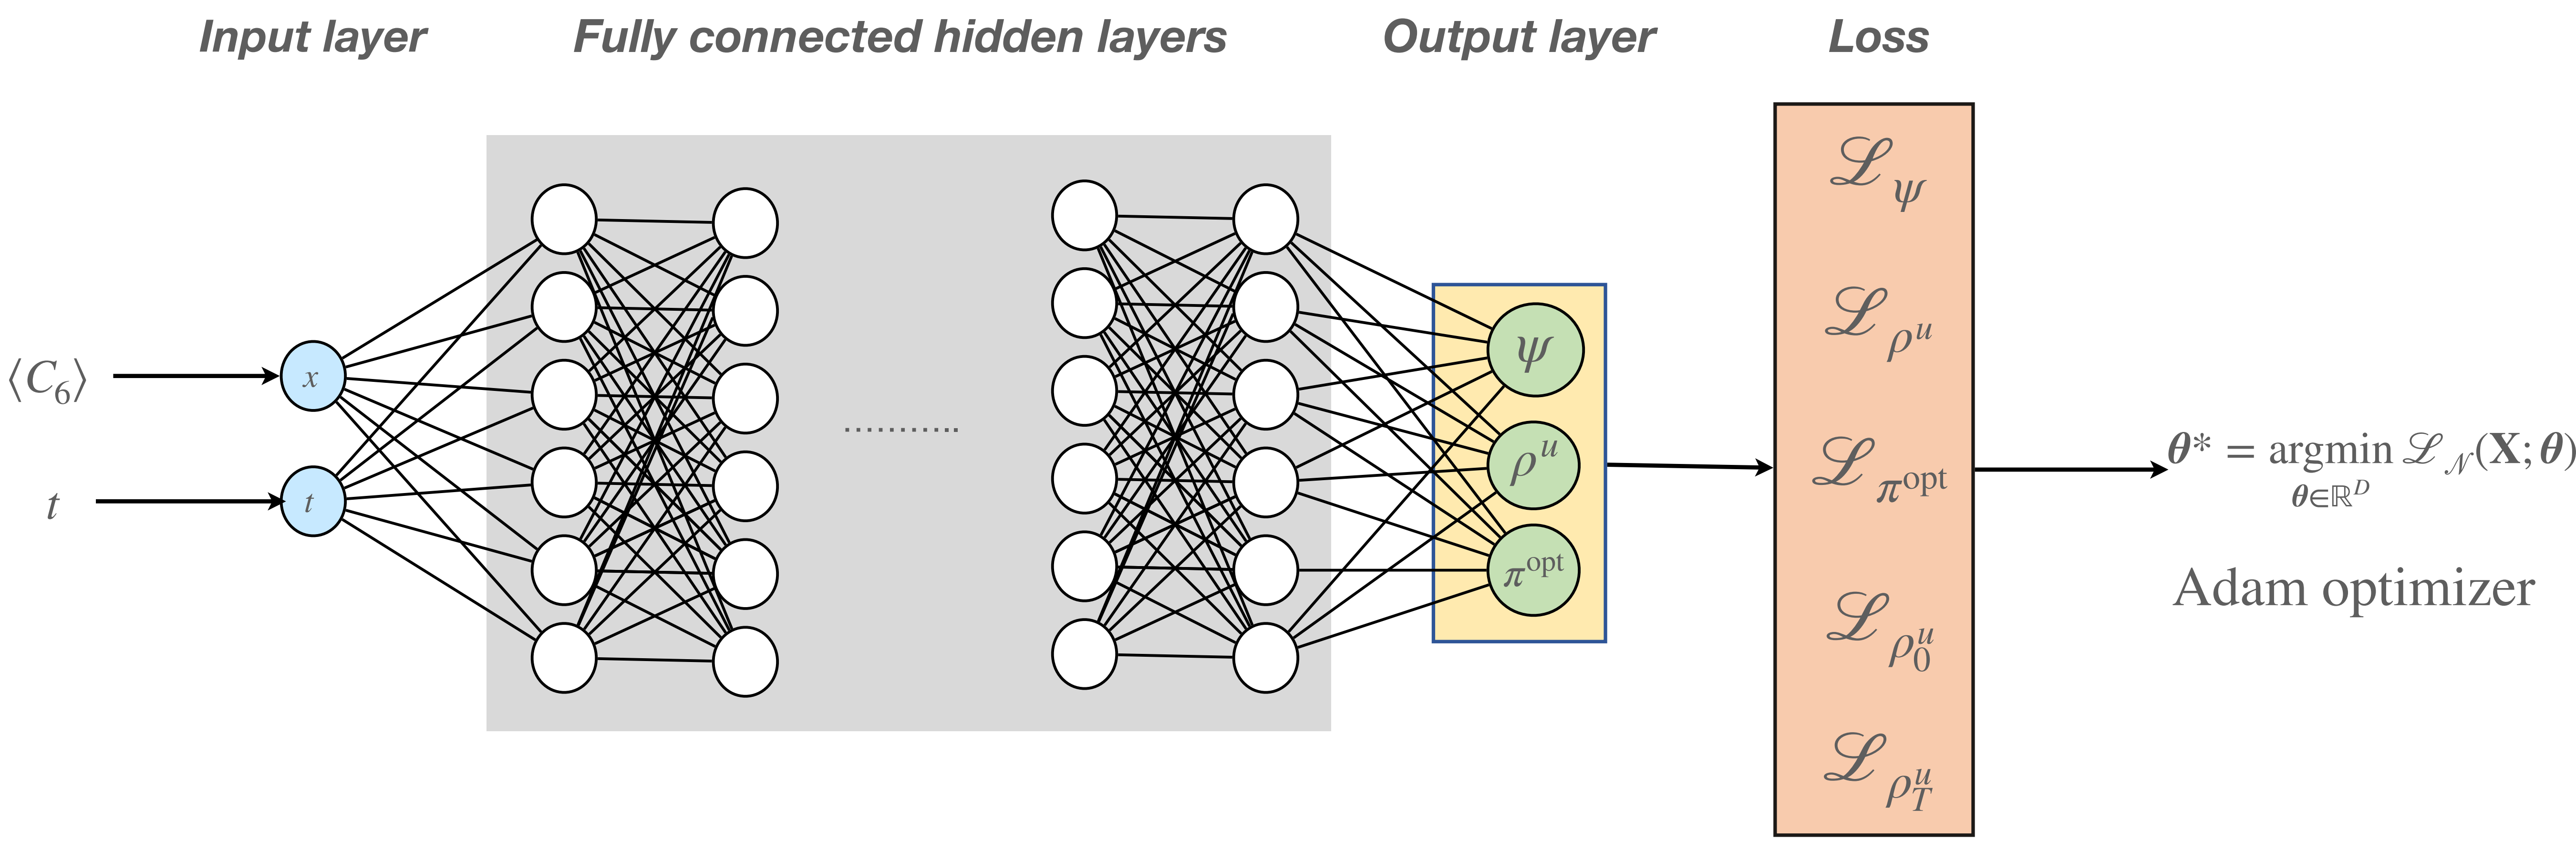

$$\mathcal{L}_{\mathcal{N}} = \mathcal{L}_\psi + \mathcal{L}_{\rho^u} + \mathcal{L}_{\pi^{\text{opt}}} + \mathcal{L}_{\rho^u_0} + \mathcal{L}_{\rho^u_T}$$

[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]

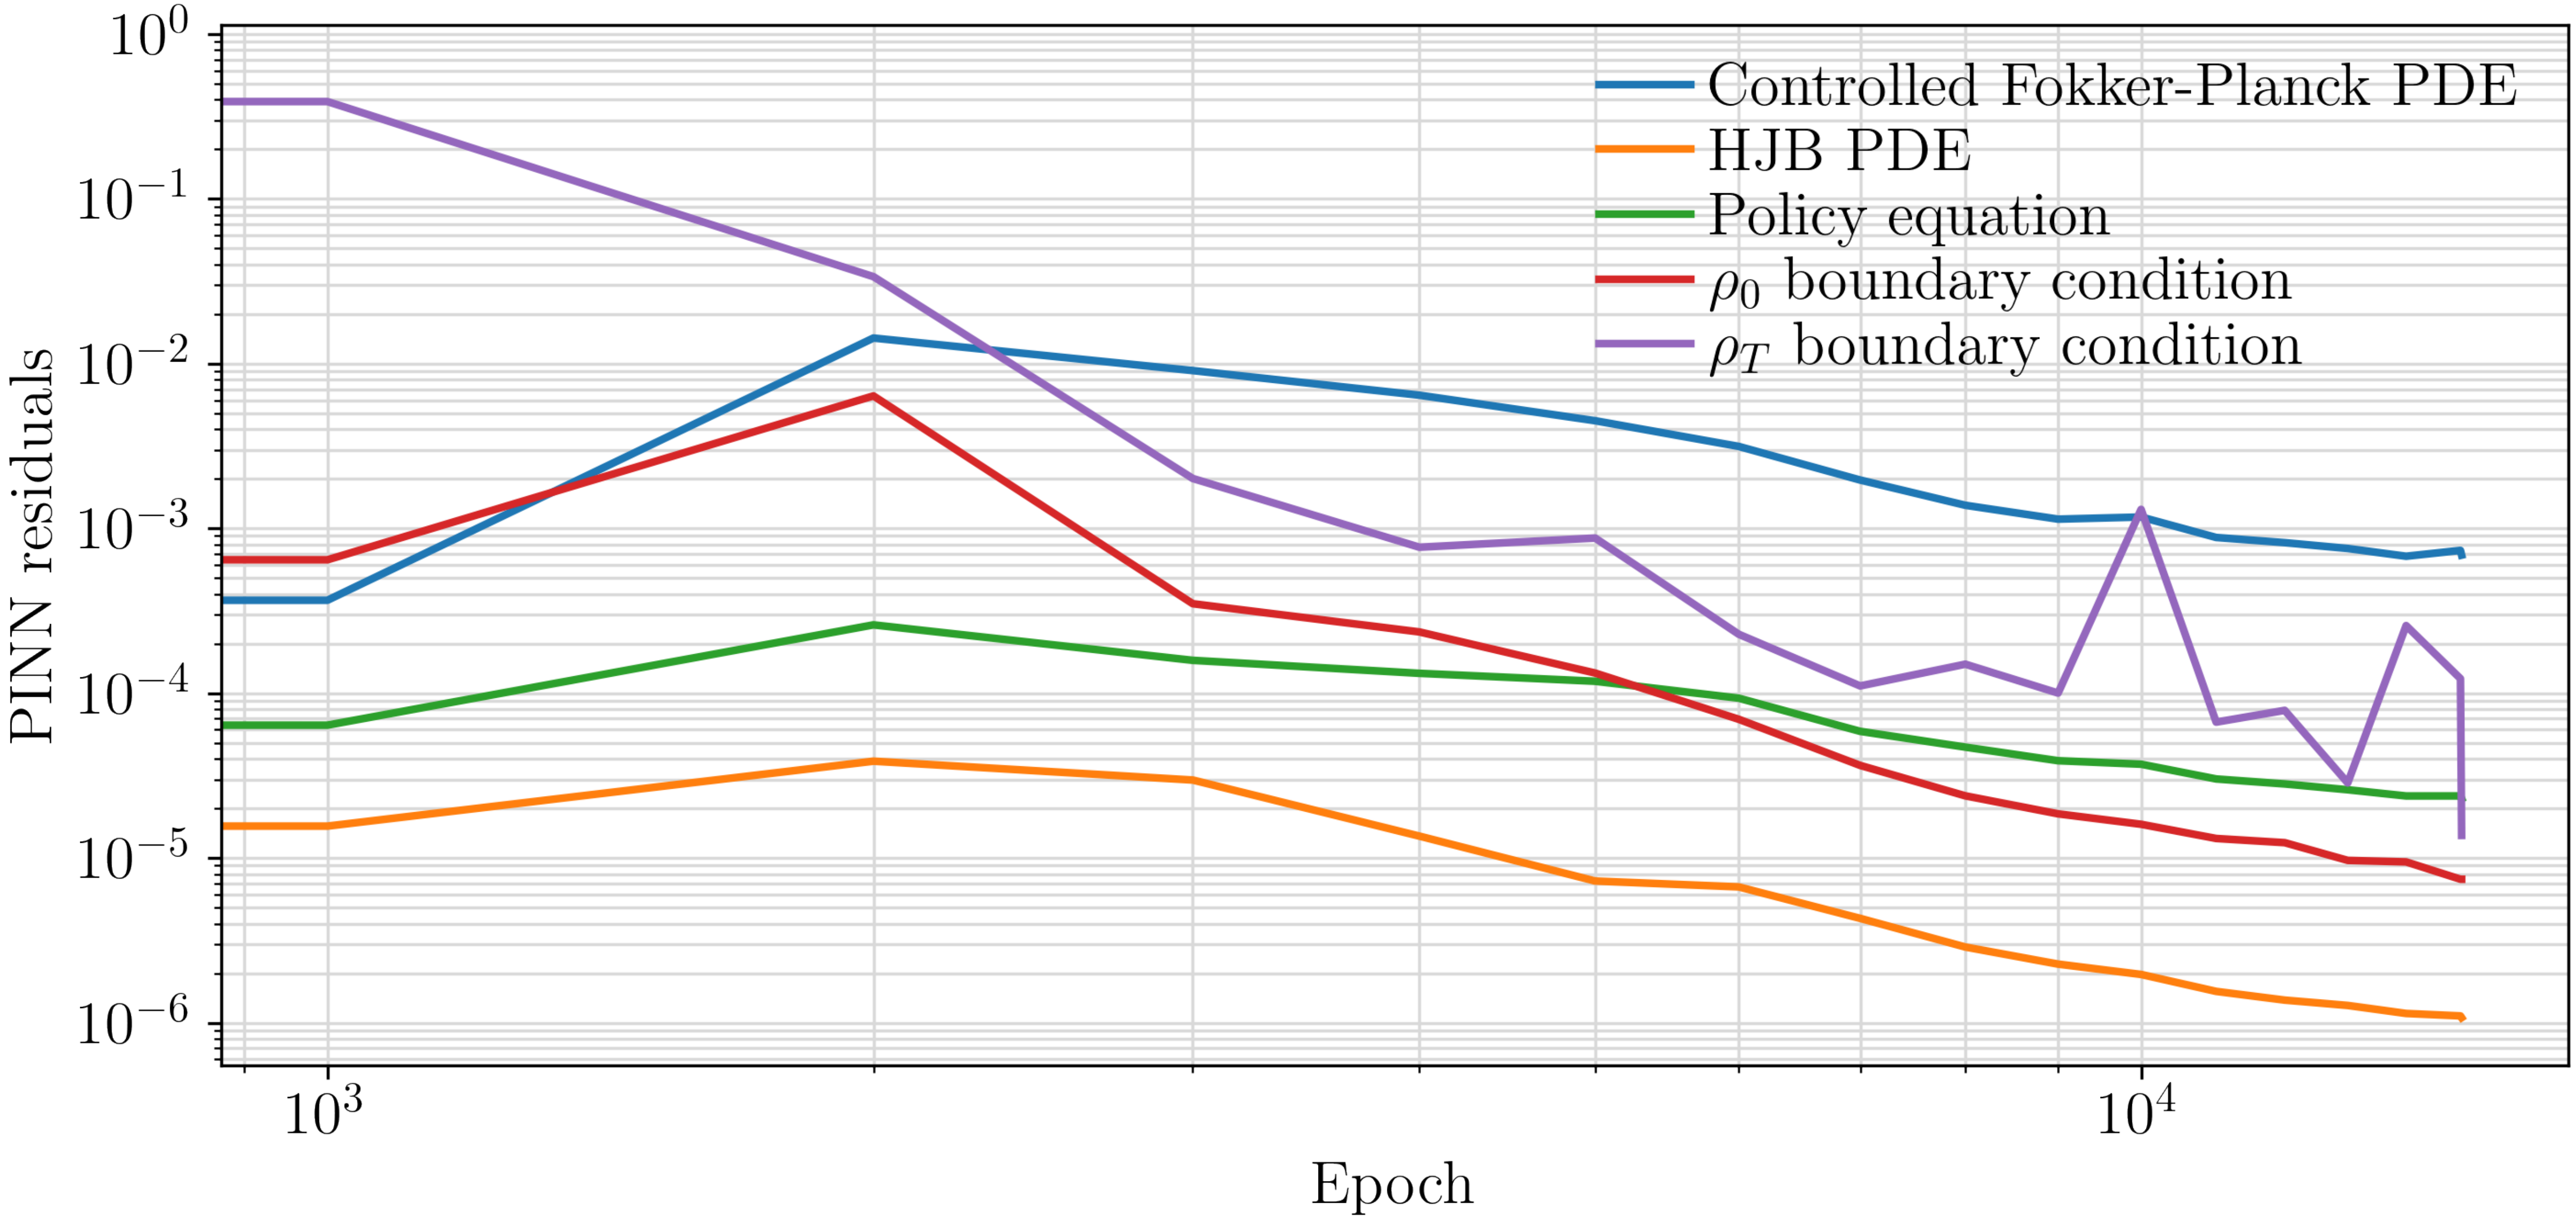

Optimally Controlled State PDFs

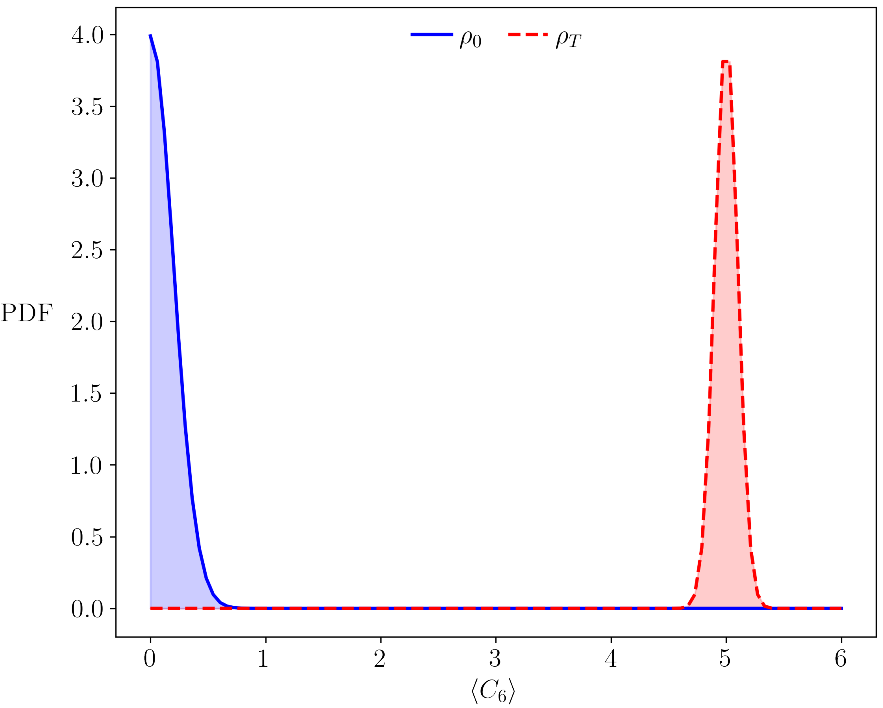

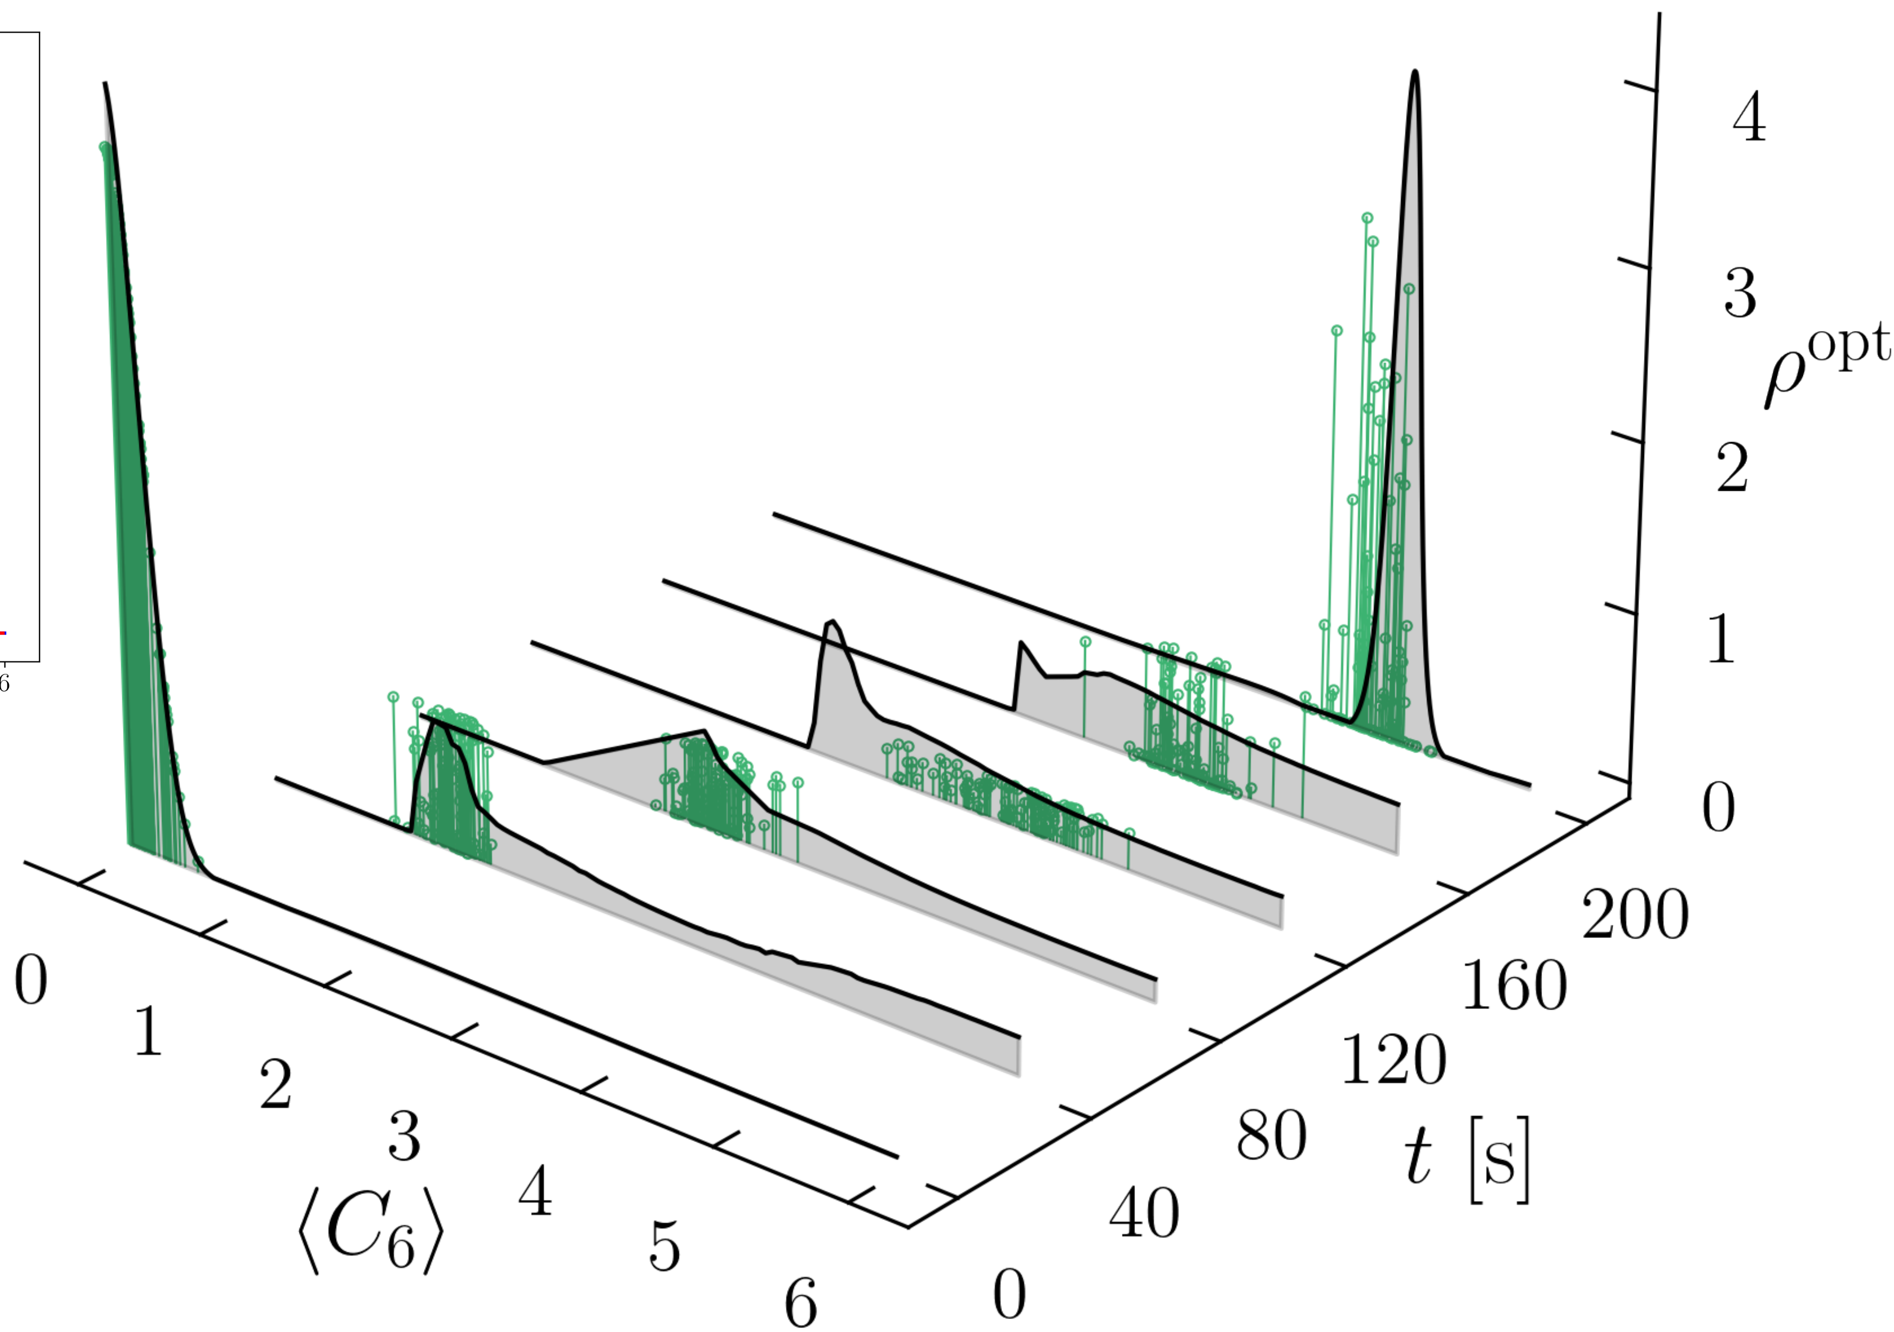

Optimal State and Optimal Control Sample Paths

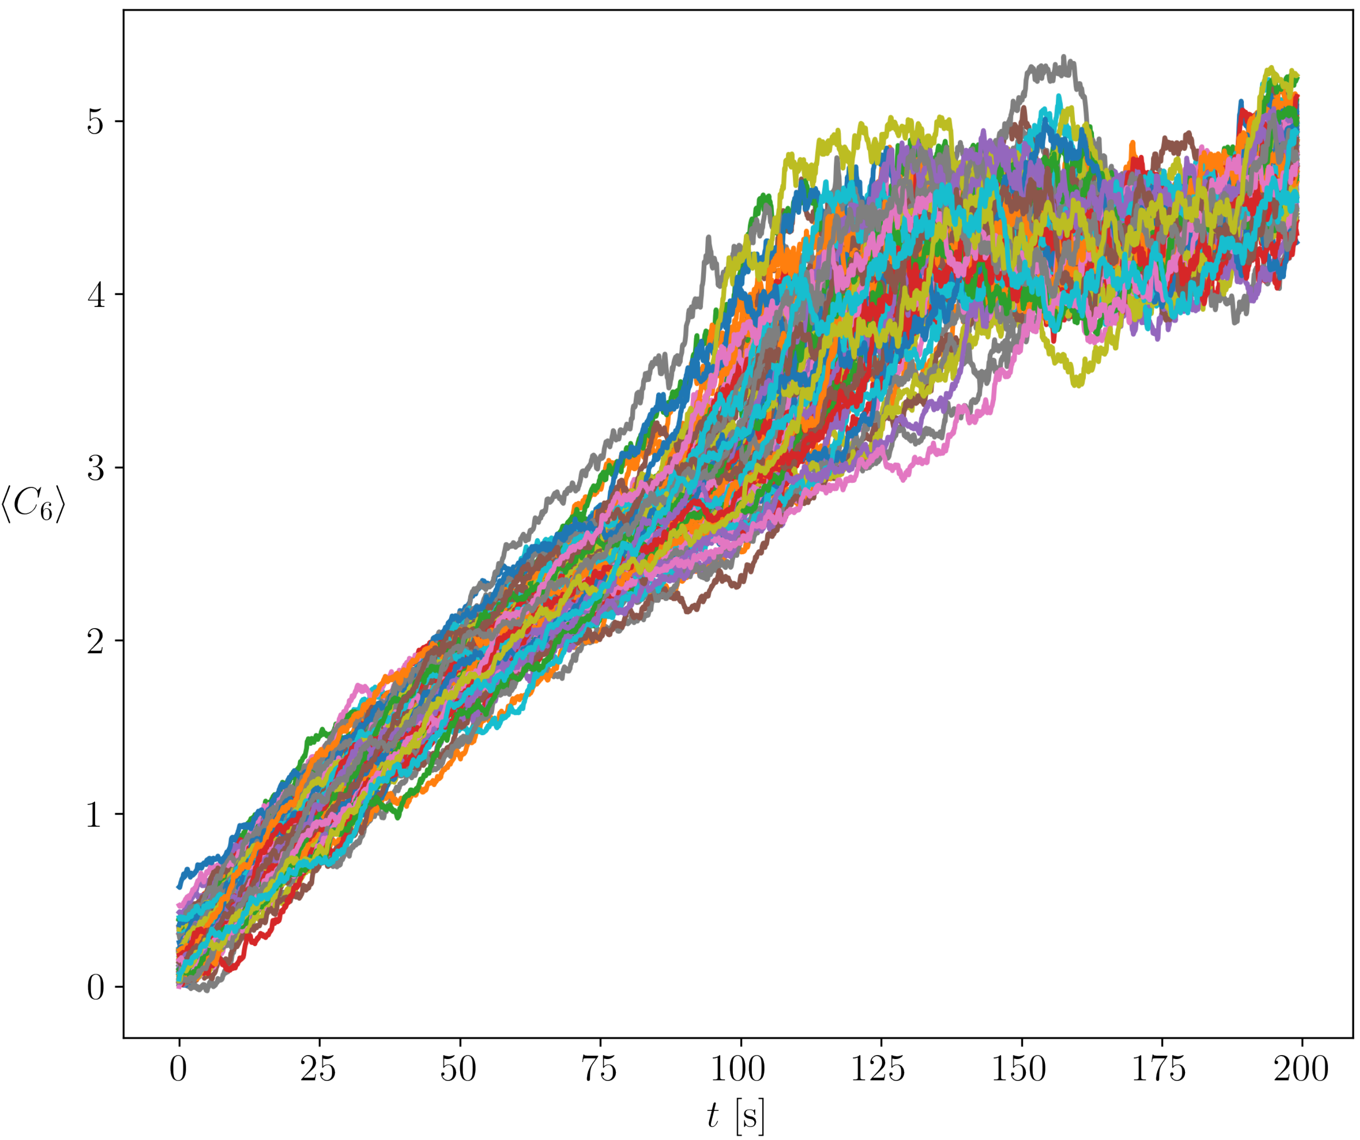

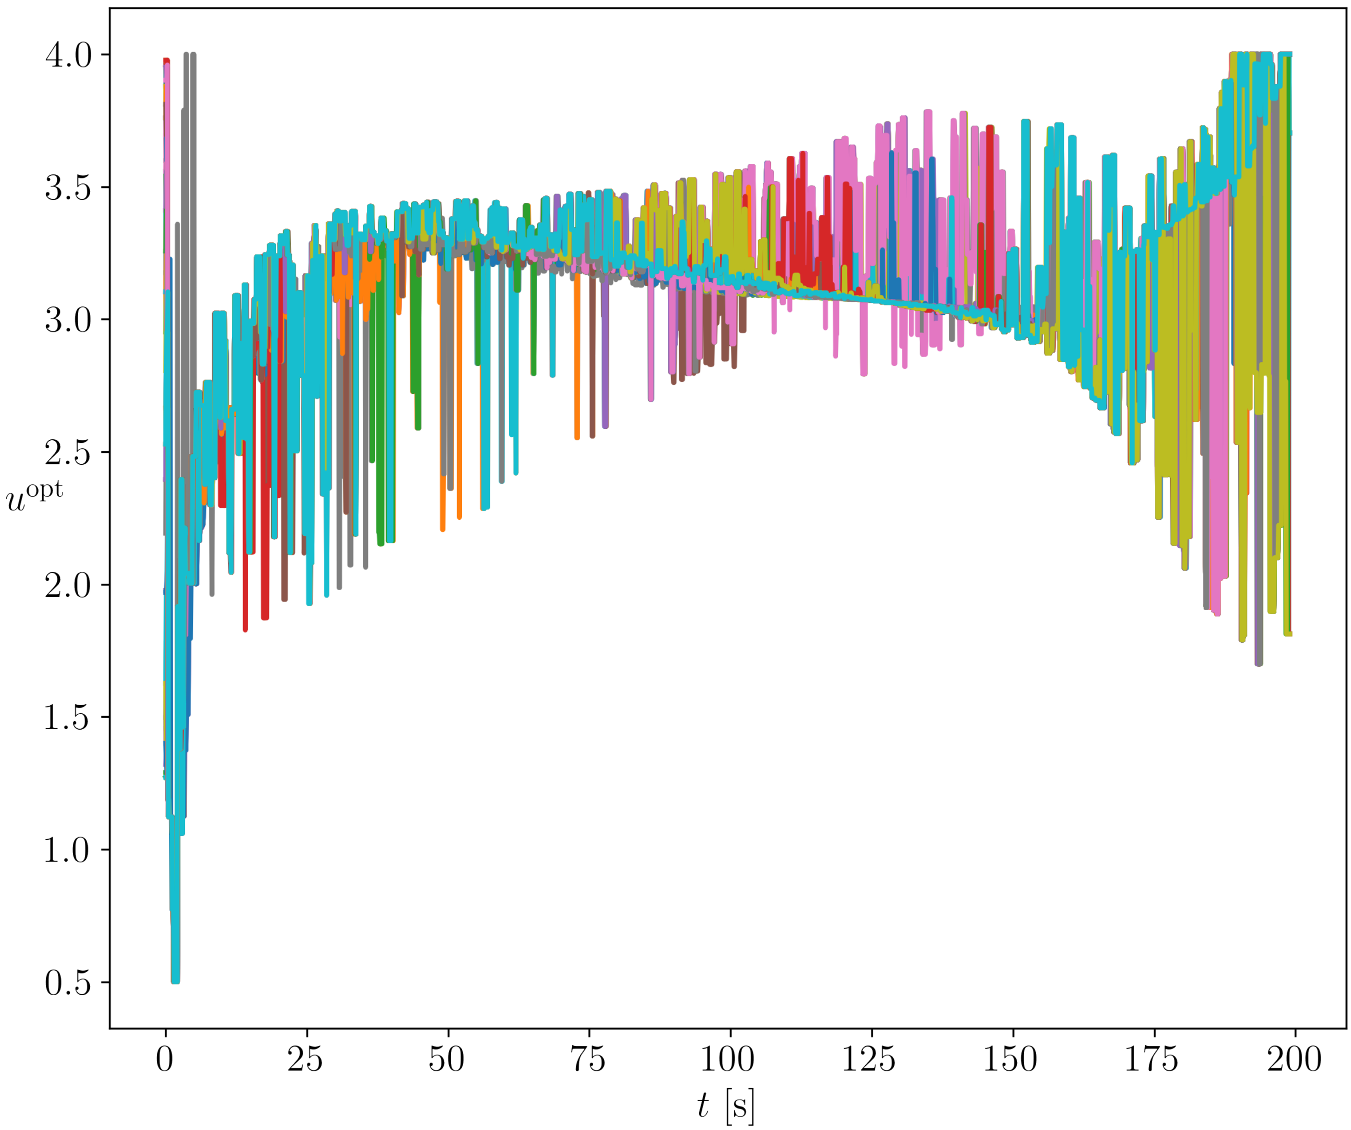

GSBP Conditions for Optimality with m Inputs

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

Drift coefficient   Diffusion tensor

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

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

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

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

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

Data-driven GSBP for Colloidal SA

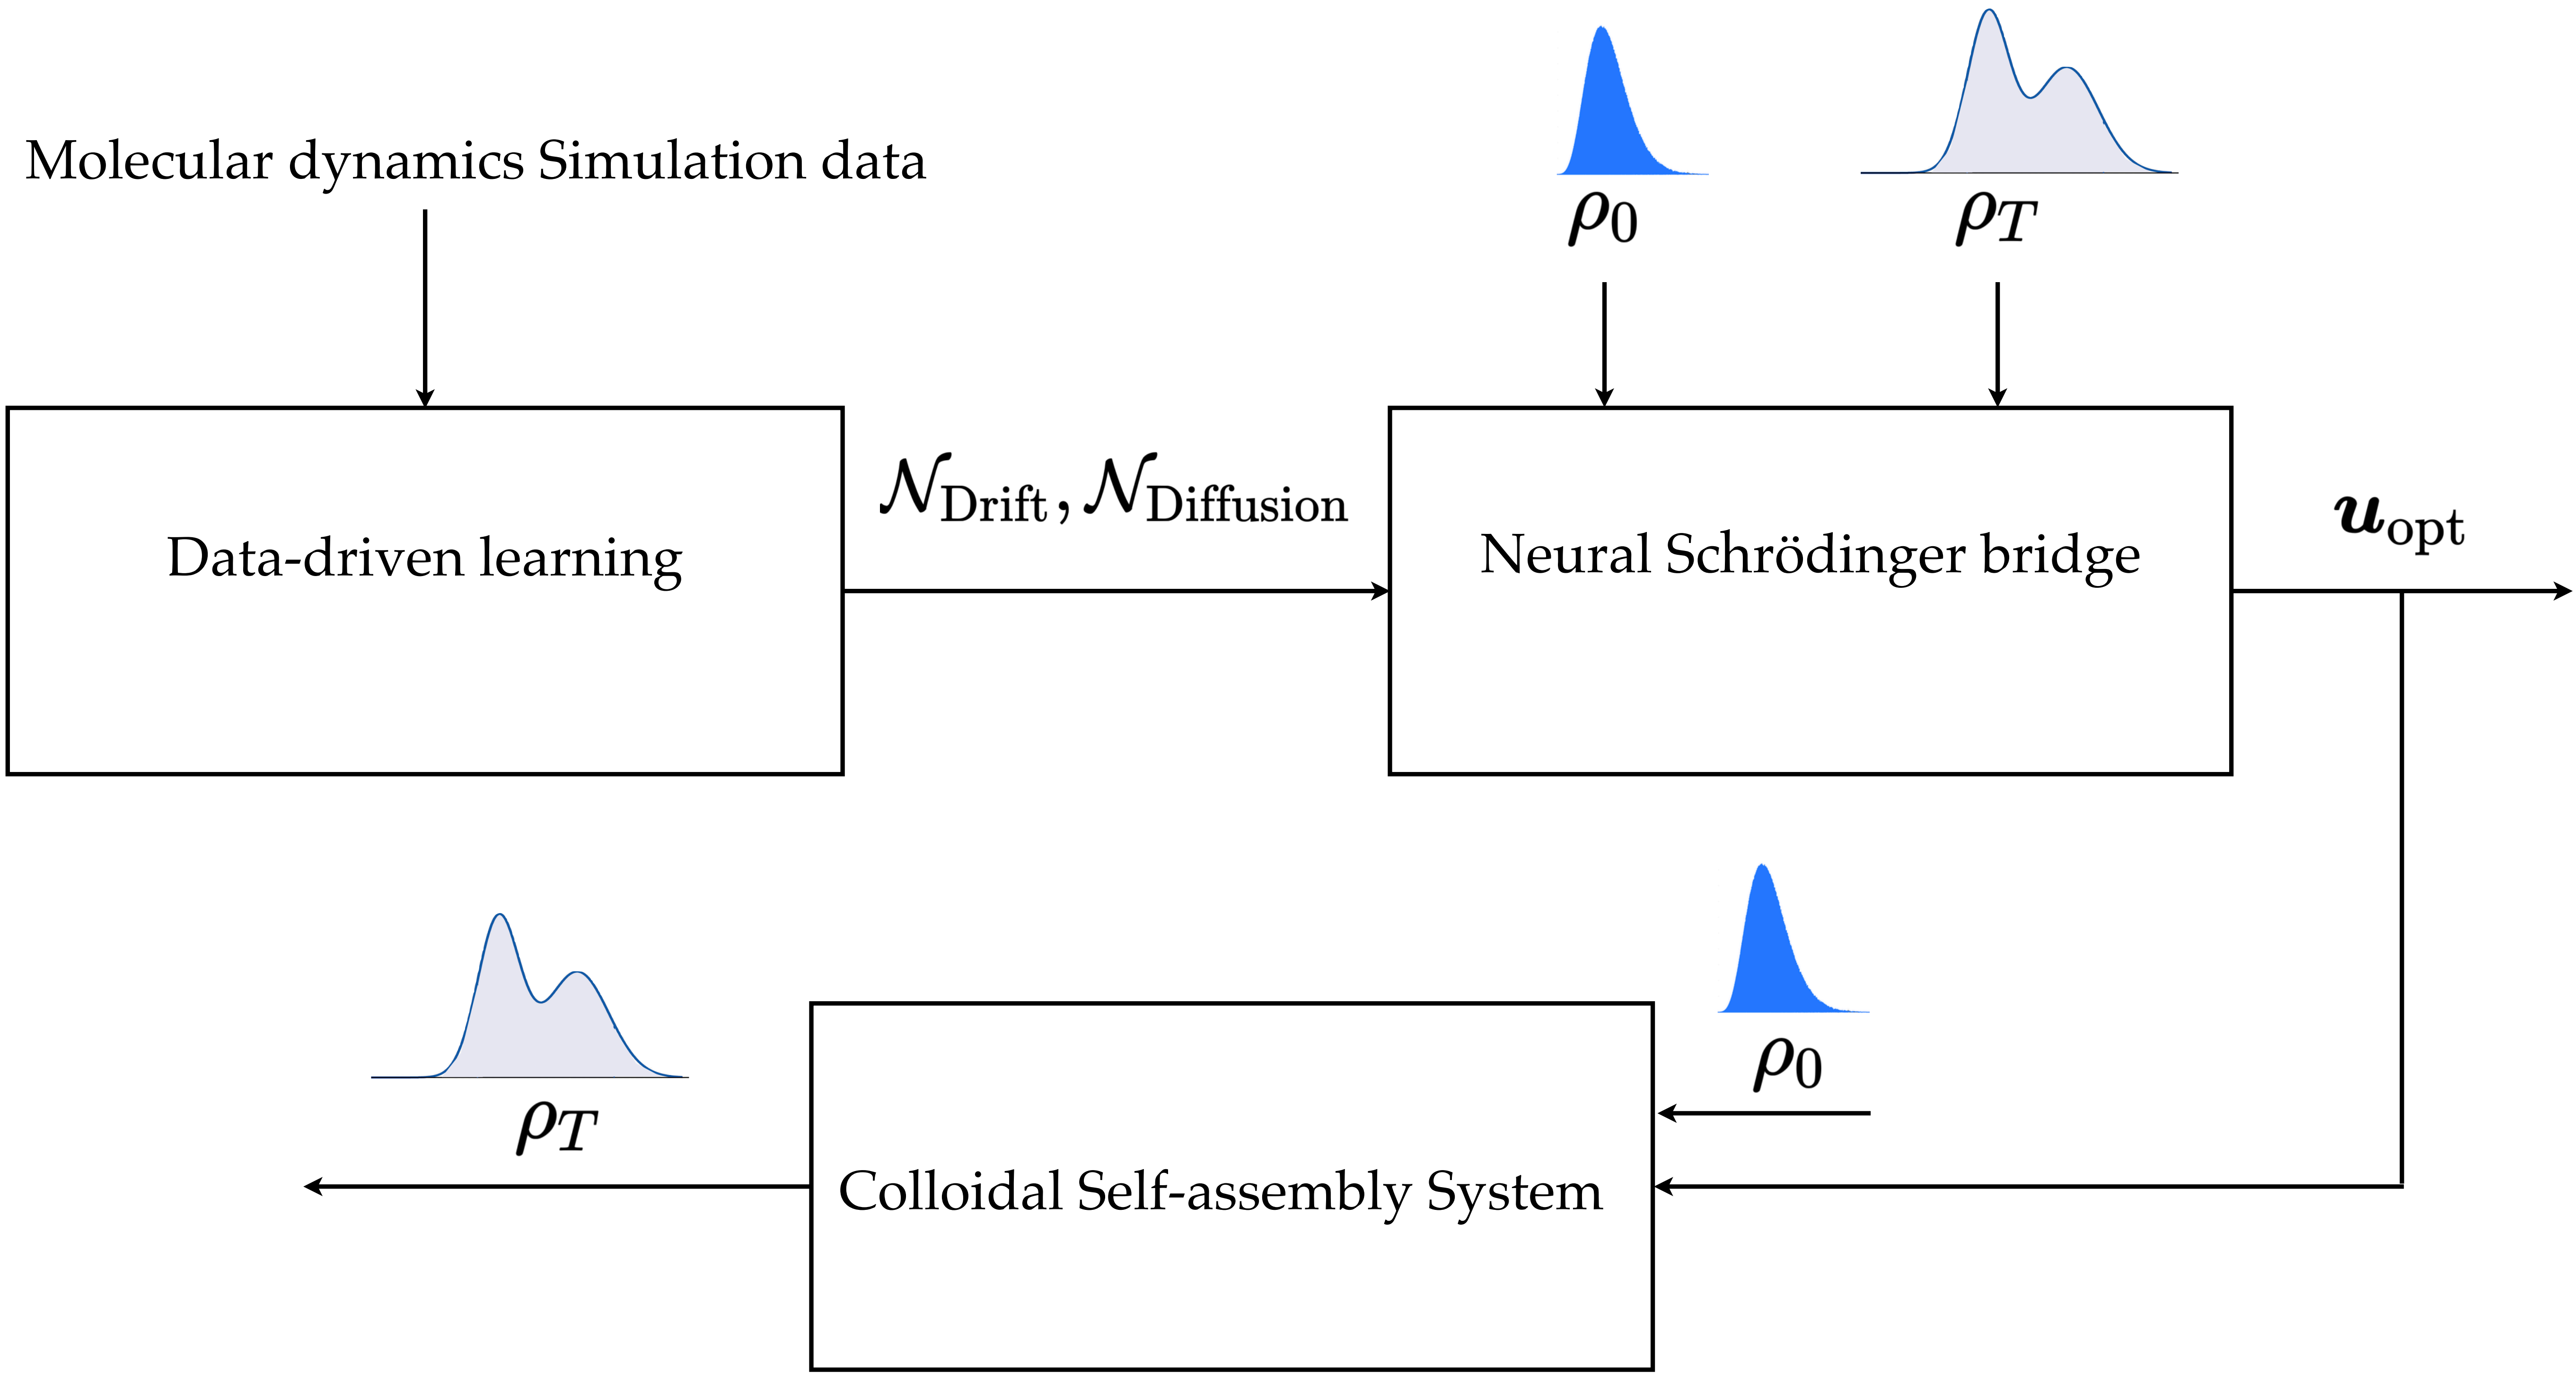

Architecture for Data-driven GSBP

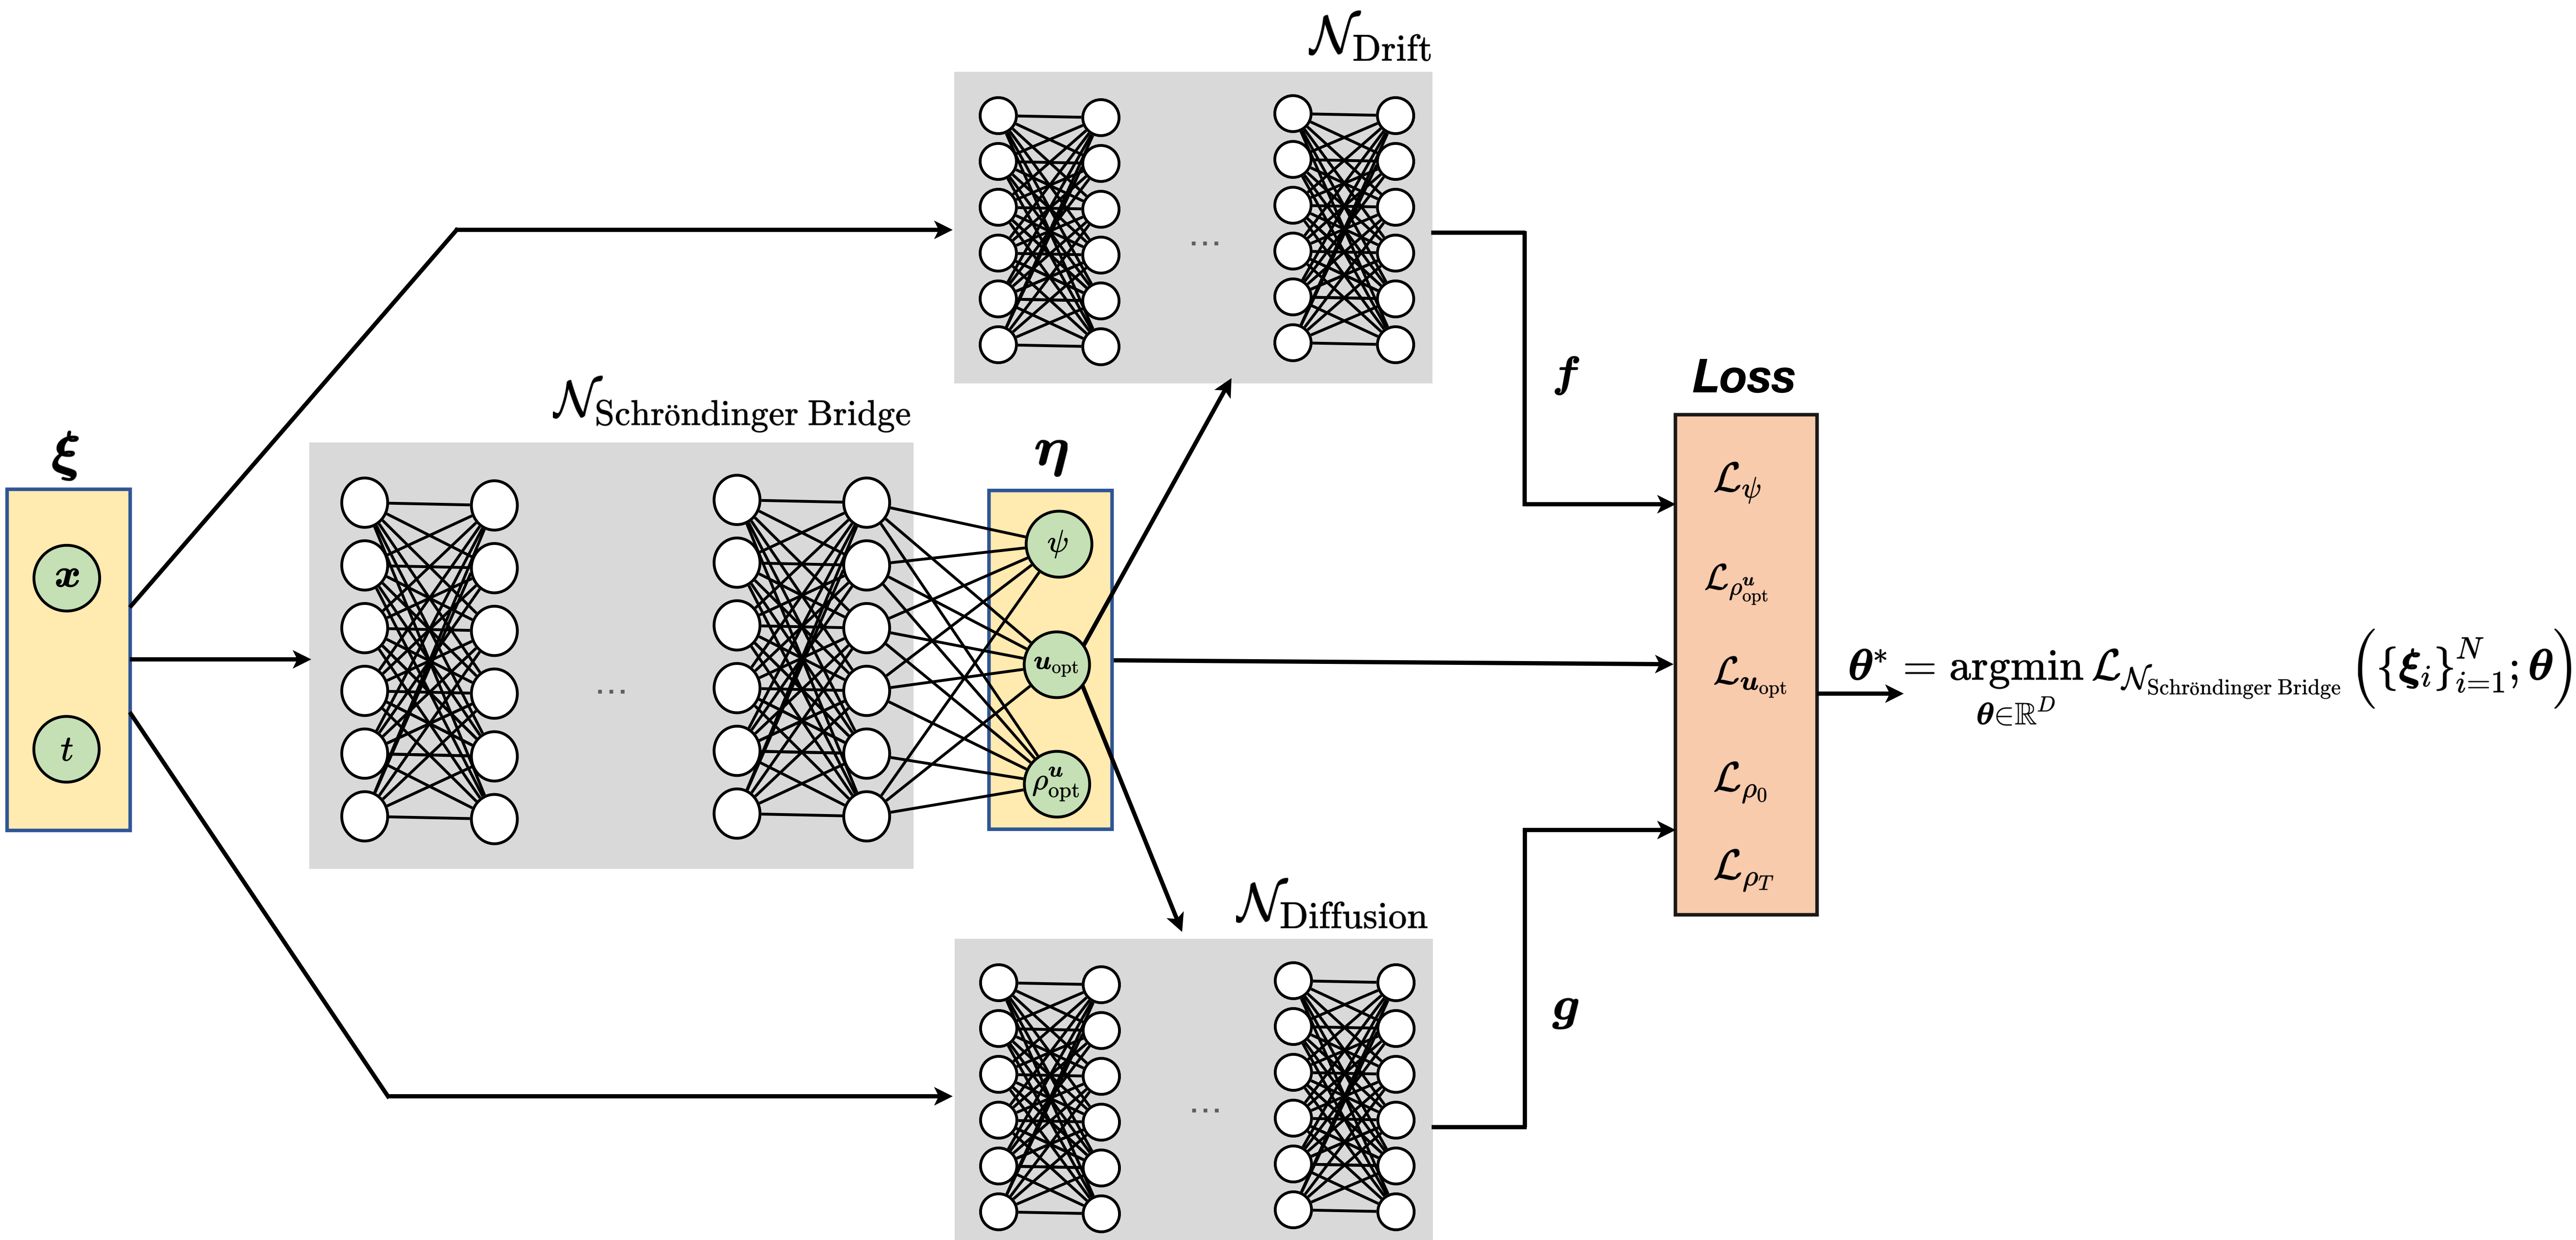

Sinkhorn Losses for Boundary Conditions

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

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

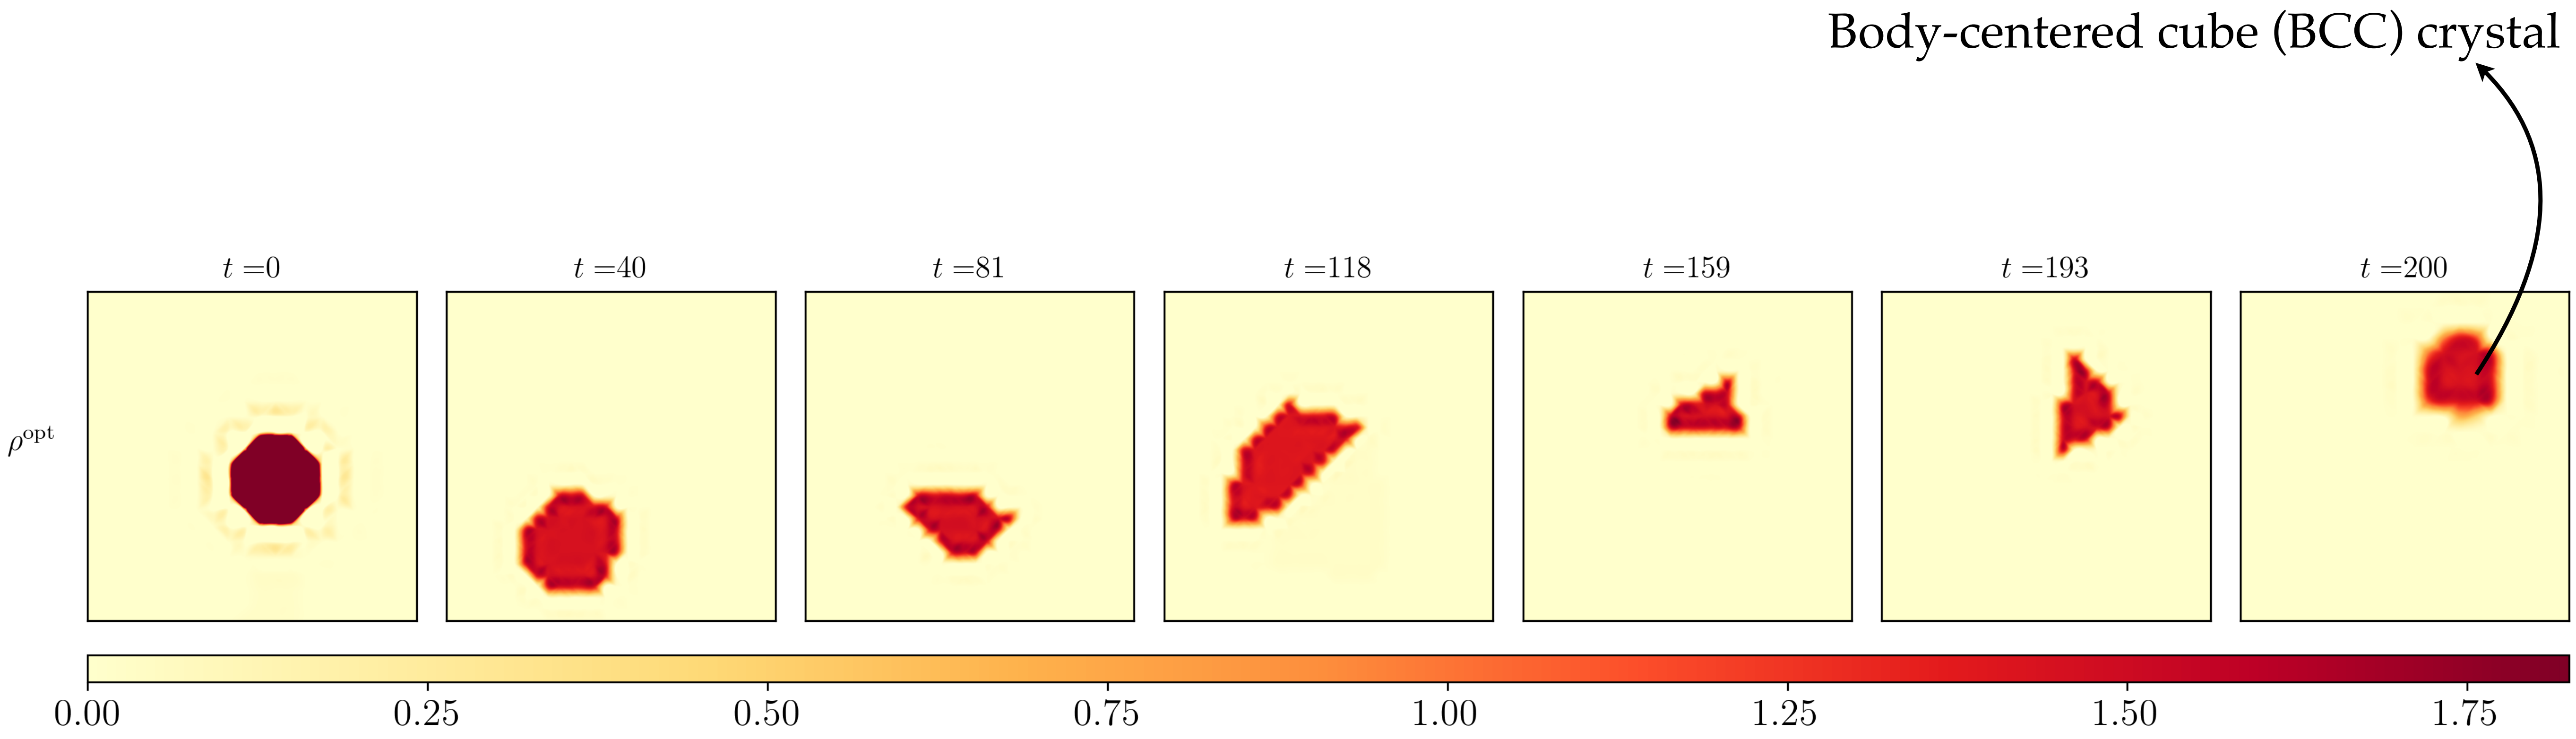

Data-driven:

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

Take Home Message

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

- Excellent intersections with related communities: ML, robotics, systems biology, smart manufacturing

Hiring: students and postdocs to work in Schrödinger bridge, stochastic control, stochastic ML

Thank You

Support: 1923278, 2112755, 2111688