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

Abhishek Halder

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

Joint work with S. Haddad, K.F. Caluya (UC Santa Cruz)
B. Singh (Ford), T.T. Georgiou (UC Irvine), W. Krichene (Google)

Yahoo! Research
March 30, 2022

Overarching Theme

Systems-control theory and algorithms for densities

What is density?

Probability Density Fn.

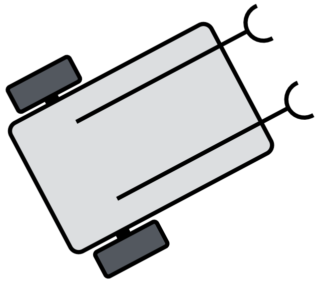

$$\boldsymbol{x}(t) \in \begin{pmatrix} x \\ y \\ \theta \end{pmatrix} \in \mathcal{X} \equiv \mathbb{R}^2 \times \mathbb{S}^1$$

Probability Density Fn.

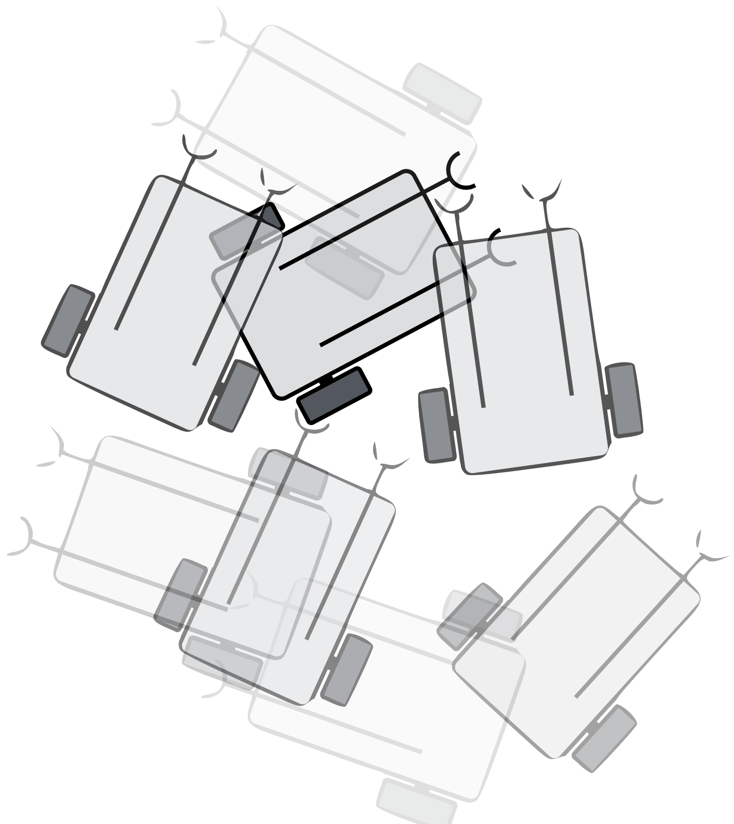

$$\boldsymbol{x}(t) \in \begin{pmatrix} x \\ y \\ \theta \end{pmatrix} \in \mathcal{X} \equiv \mathbb{R}^2 \times \mathbb{S}^1$$

$$\rho\left(\boldsymbol{x}, t\right) : \mathcal{X} \times [0, \infty) \mapsto \mathbb{R}_{\geq 0}$$

$$\int_{\mathcal{X}} \rho \, \mathrm{d}\boldsymbol{x} = 1 \quad \text{for all } t \in [0, \infty)$$

Probability Density Fn.

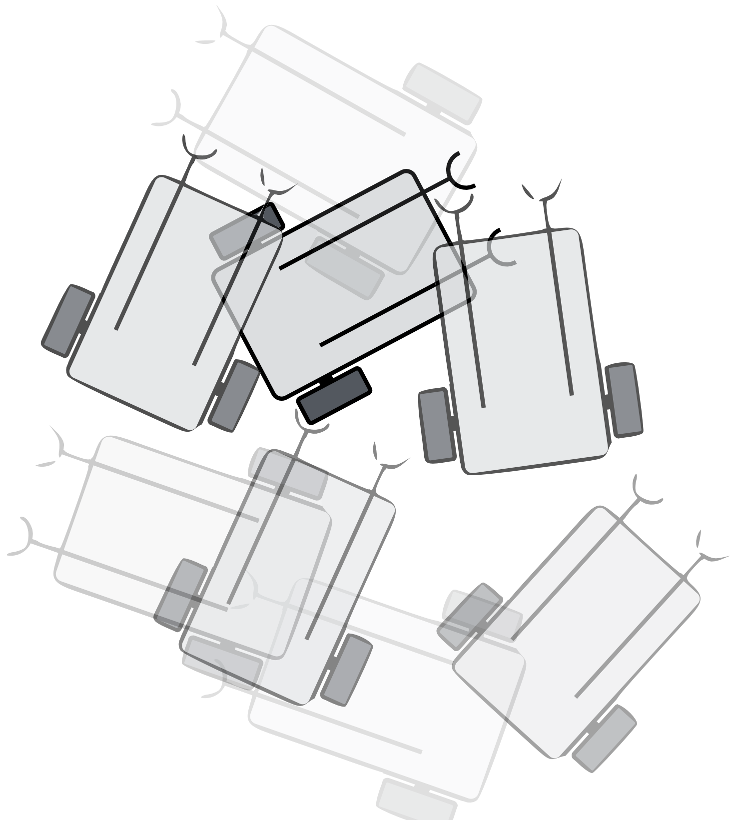

Population Density Fn.

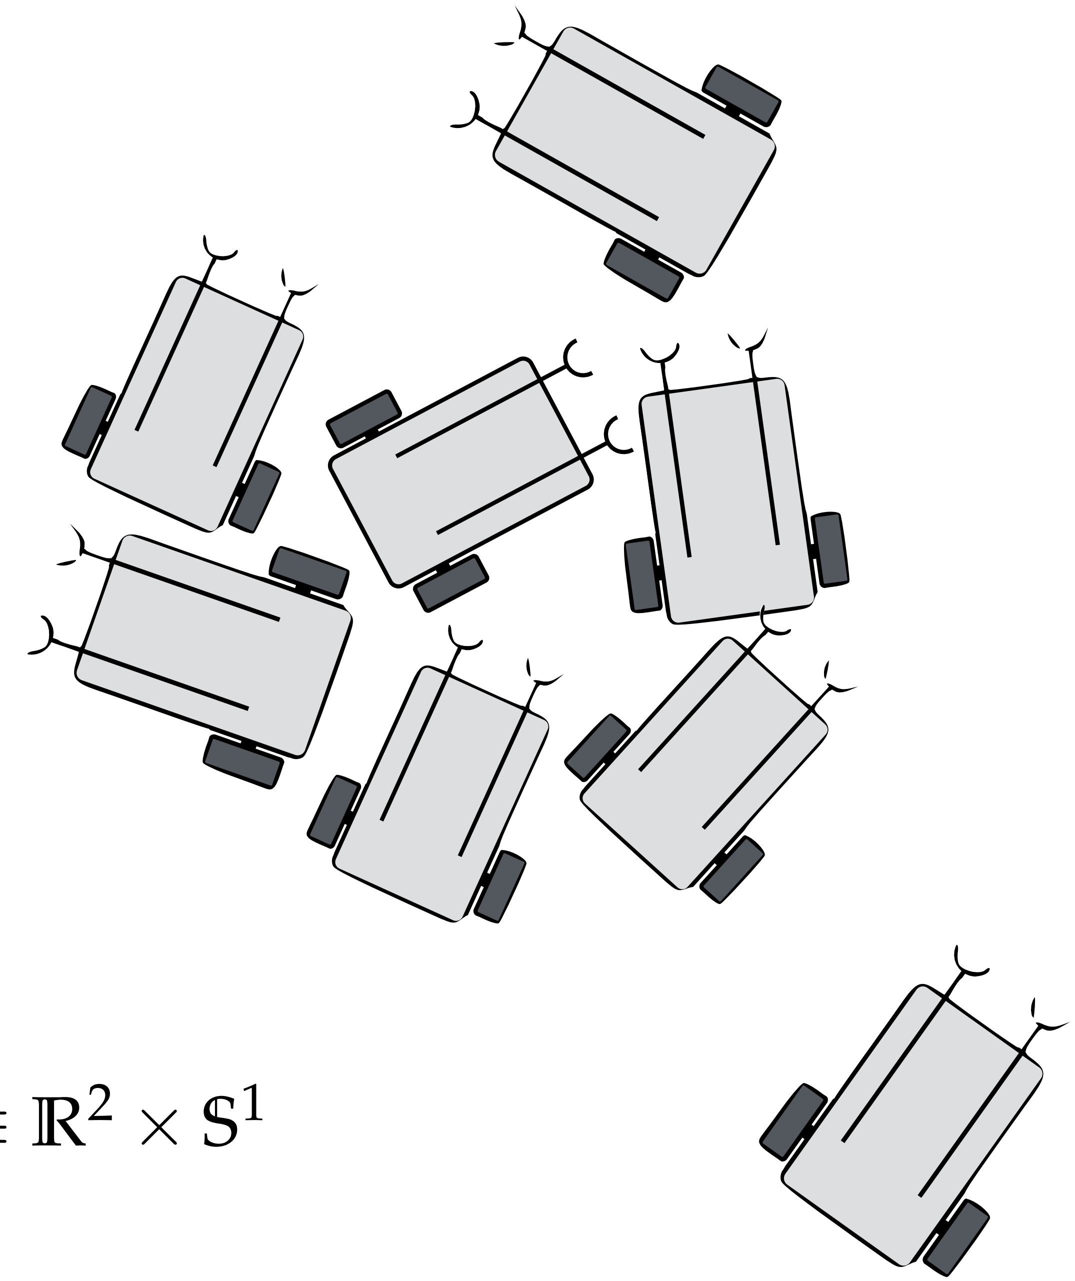

$$\boldsymbol{x}(t) \in \begin{pmatrix} x \\ y \\ \theta \end{pmatrix} \in \mathcal{X} \equiv \mathbb{R}^2 \times \mathbb{S}^1$$

$$\rho\left(\boldsymbol{x}, t\right) : \mathcal{X} \times [0, \infty) \mapsto \mathbb{R}_{\geq 0}$$

$$\int_{\mathcal{X}} \rho \, \mathrm{d}\boldsymbol{x} = 1 \quad \text{for all } t \in [0, \infty)$$

Why do we care about densities in systems-control problems?

Prediction Problem

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

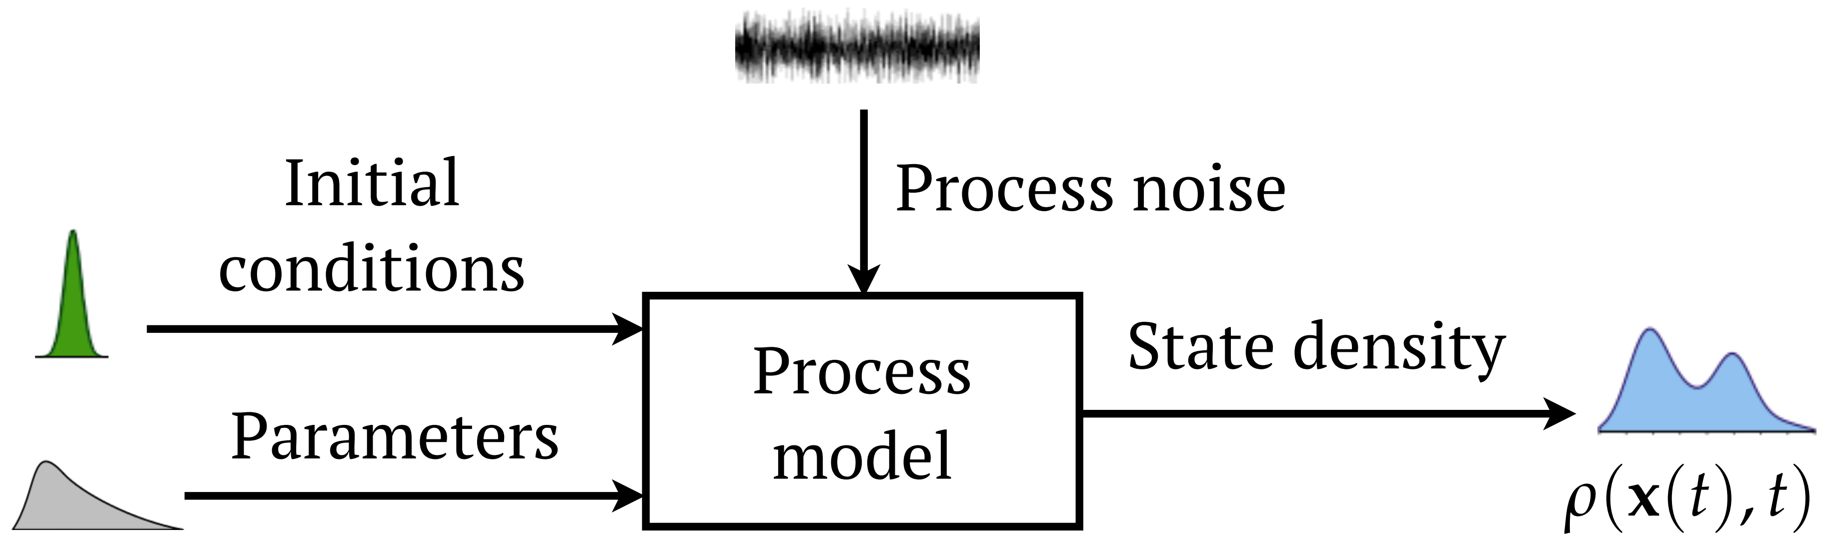

Trajectory flow:

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

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

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

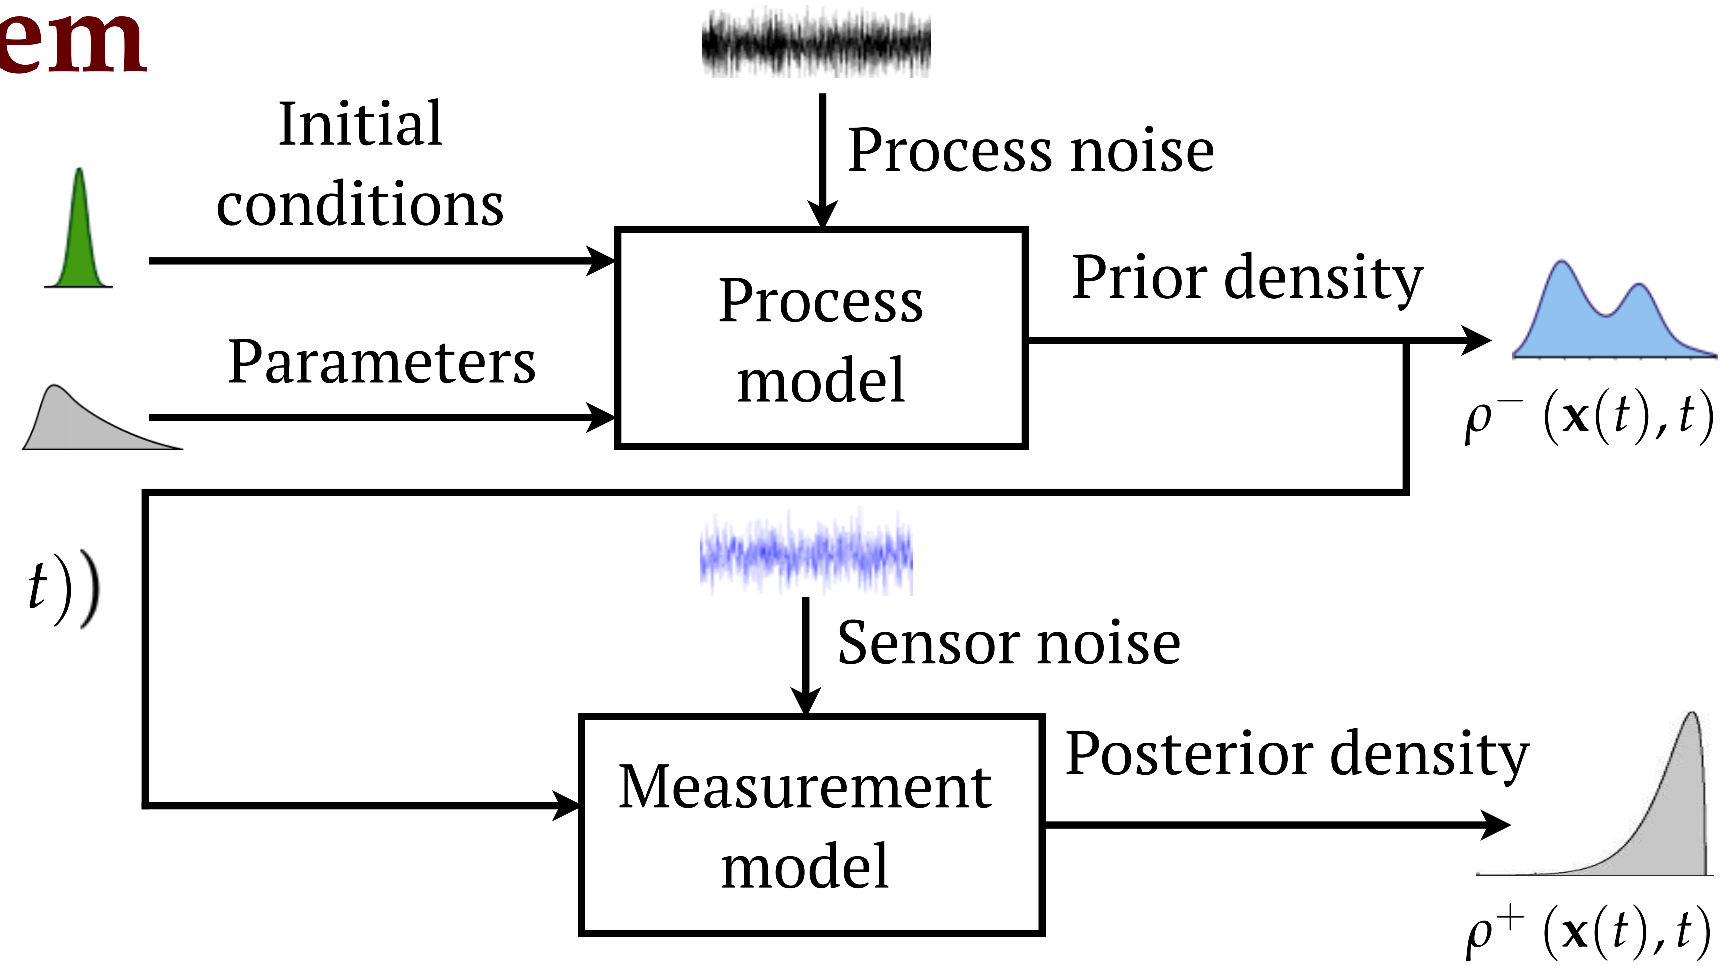

Trajectory flow:

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

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

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

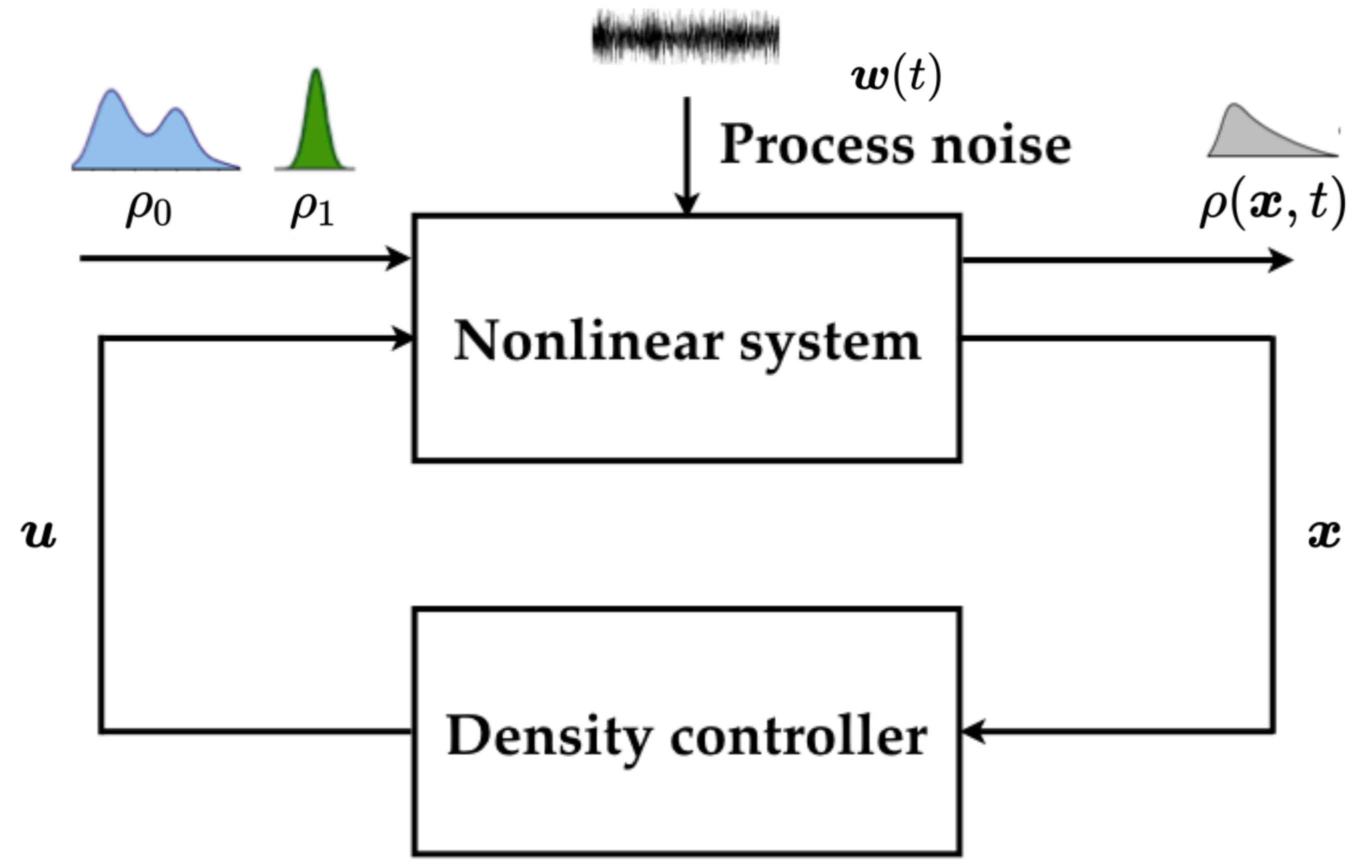

$$\underset{u \in \mathcal{U}}{\text{minimize}} \quad \mathbb{E}\left[\int_0^1 \|\boldsymbol{u}\|_2^2 \, \mathrm{d}t\right]$$

subject to

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

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

Neural Network Learning Problem

Consider fully connected NN

Think "layers" as interacting population of neurons

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

PDF dynamics:

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

Solving prediction problem as generalized 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:

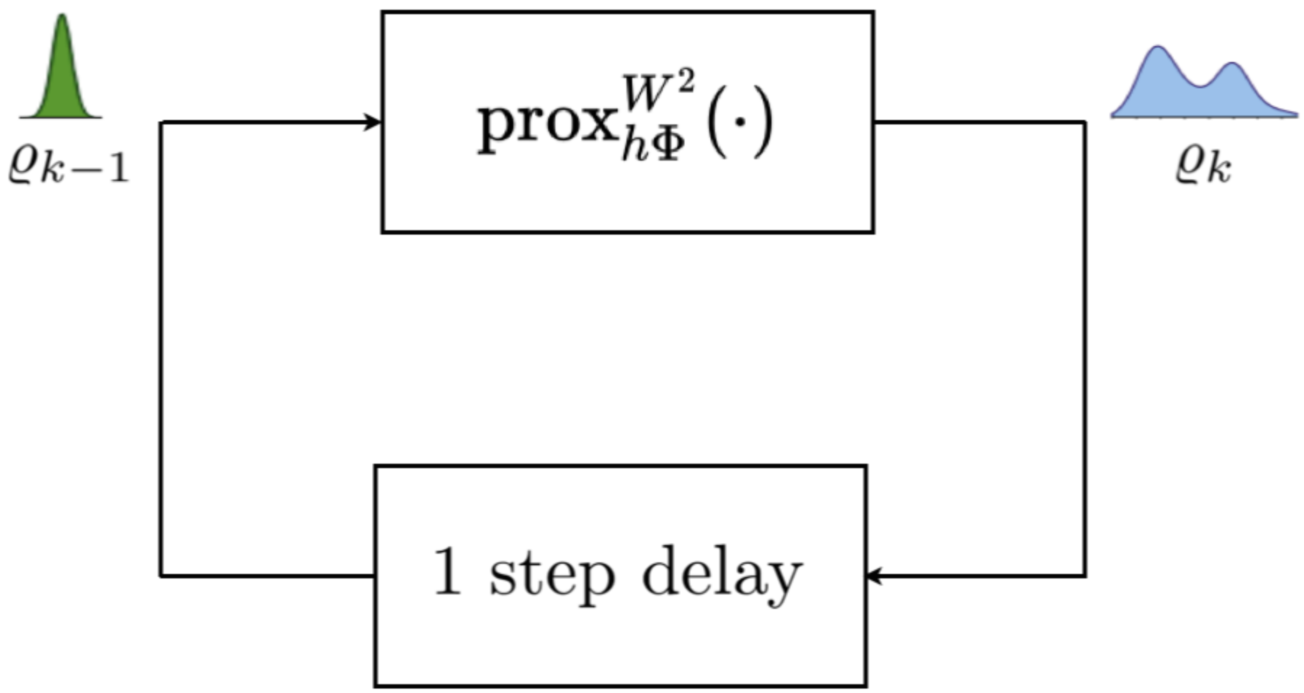

Proximal operator: $\varrho_k = \mathrm{prox}_{h\Phi}^{W^2}(\varrho_{k-1}) := \underset{\varrho \in \mathcal{P}_2(\mathcal{X})}{\arg\inf} \left\{ \frac{1}{2} W^2(\varrho, \varrho_{k-1}) + h\Phi(\varrho) \right\}$

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

Geometric Meaning of Gradient Flow

Gradient Flow in $\mathcal{X}$

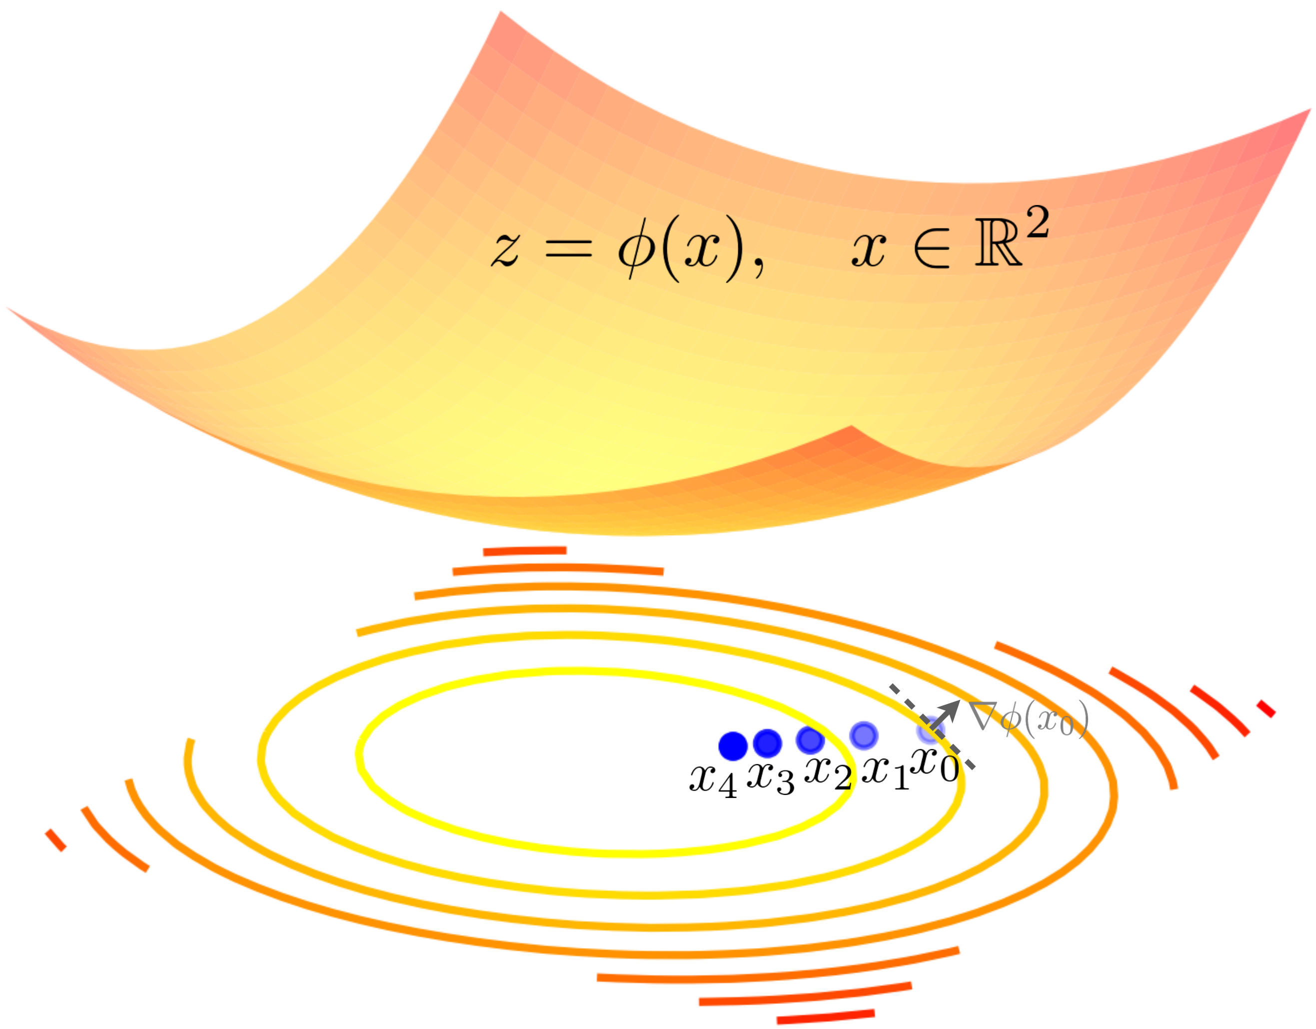

Gradient Flow in $\mathcal{P}_2(\mathcal{X})$

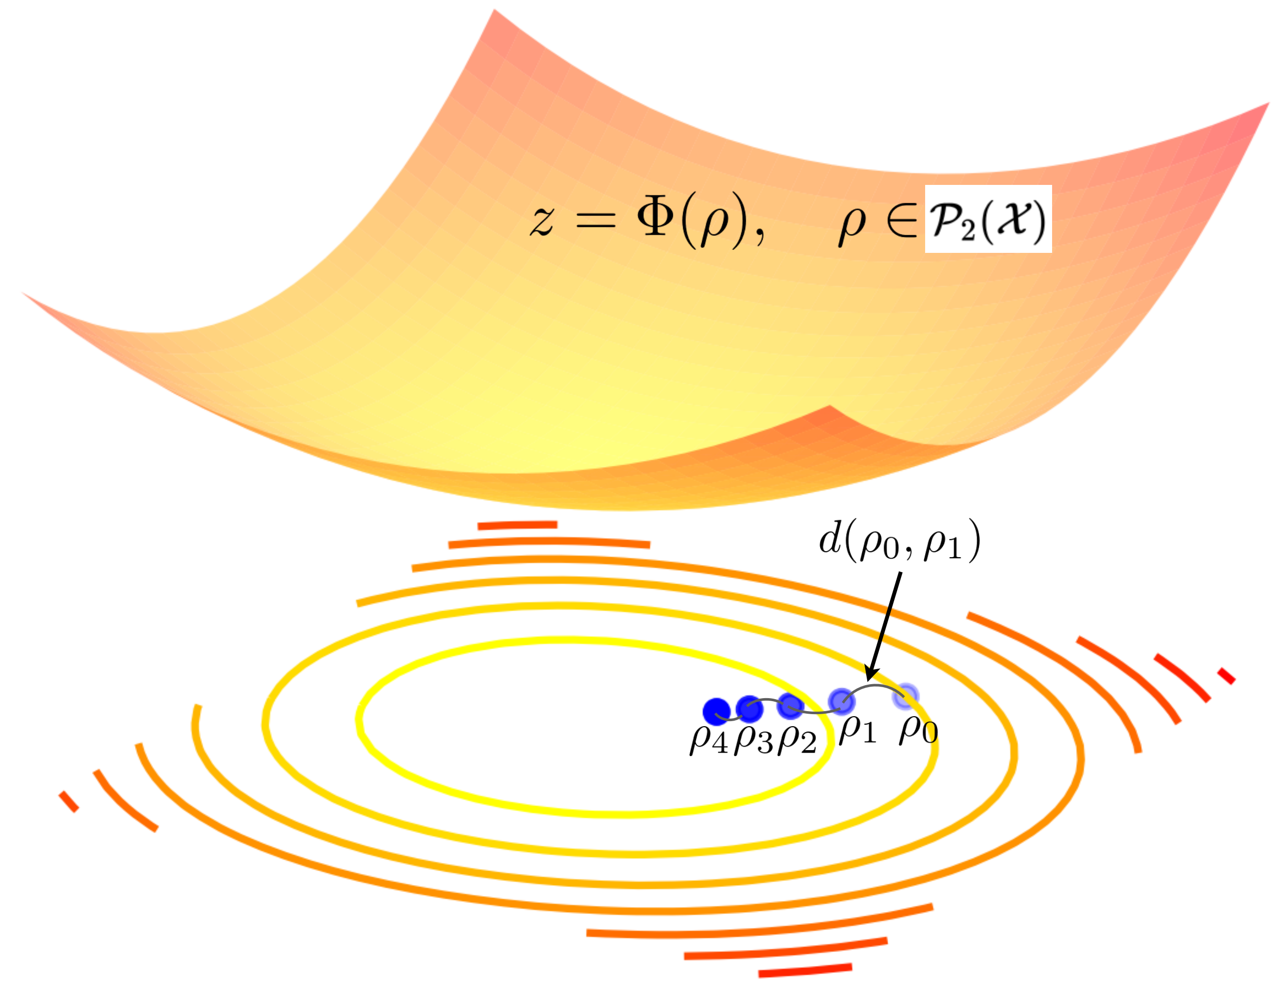

Algorithm: Gradient Ascent on the Dual Space

Uncertainty propagation via point clouds

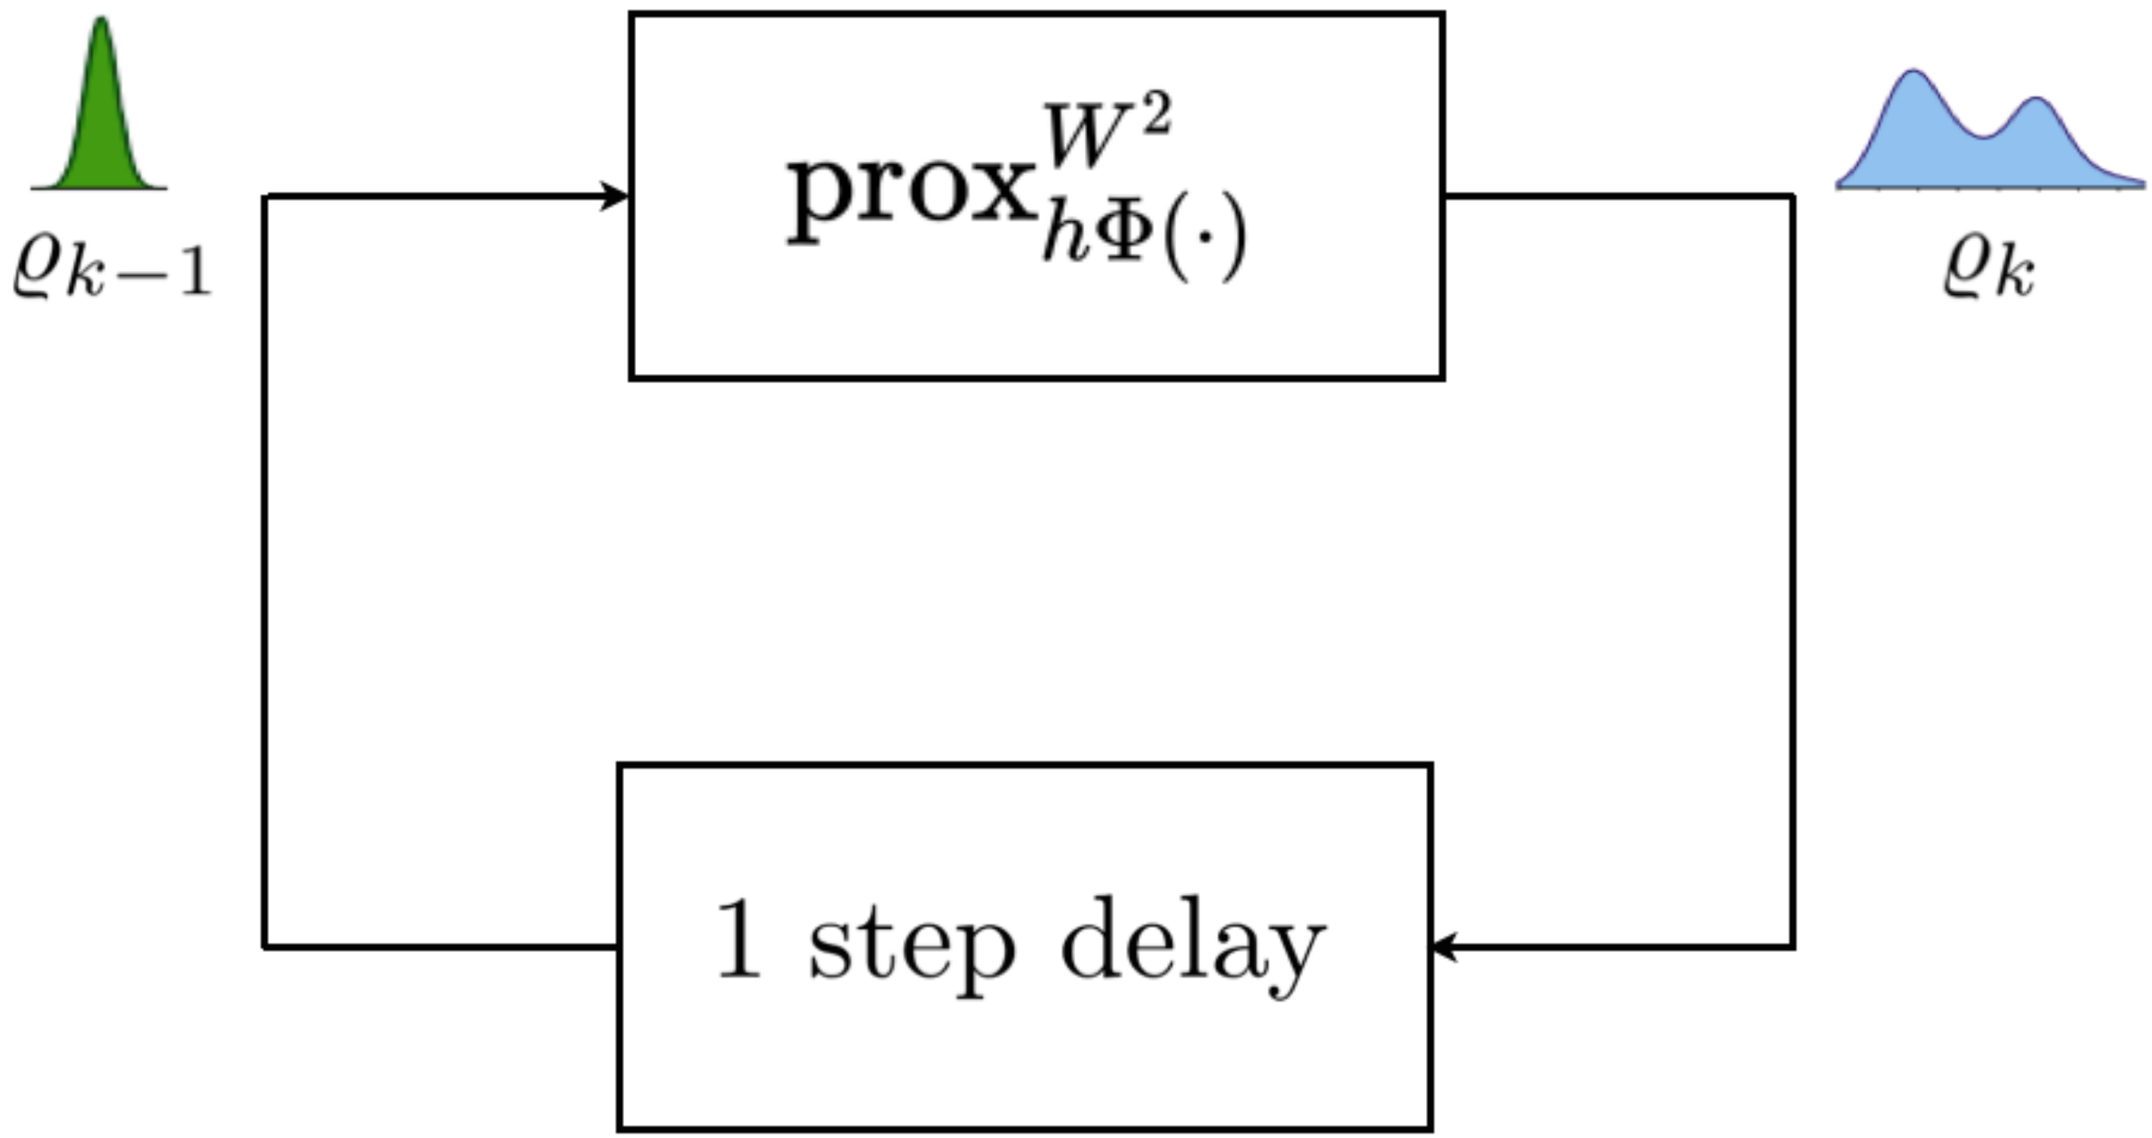

No spatial discretization or function approximation

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**

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

Recursion on the Cone

$$\boldsymbol{y} = e^{\frac{\lambda_0^*}{\epsilon}h} \qquad \boldsymbol{z} = e^{\frac{\lambda_1^*}{\epsilon}h}$$

Coupled Transcendental Equations in $\boldsymbol{y}$ and $\boldsymbol{z}$

$$\boldsymbol{\Gamma}_k = e^{\frac{-\boldsymbol{C}_k}{2\epsilon}}, \quad \varrho_{k-1}, \quad \boldsymbol{\xi}_{k-1} = \frac{e^{-\beta\boldsymbol{\psi}_{k-1}}}{e} \longrightarrow$$

$$\boldsymbol{y} \odot \boldsymbol{\Gamma}_k \boldsymbol{z} = \varrho_{k-1}$$

$$\boldsymbol{z} \odot \boldsymbol{\Gamma}_k^\top \boldsymbol{y} = \boldsymbol{\xi}_{k-1} \odot \boldsymbol{z}^{-\beta\epsilon/2h}$$

$$\longrightarrow \varrho_k = \boldsymbol{z} \odot \boldsymbol{\Gamma}_k^\top \boldsymbol{y}$$

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* 2020.

Algorithmic Setup

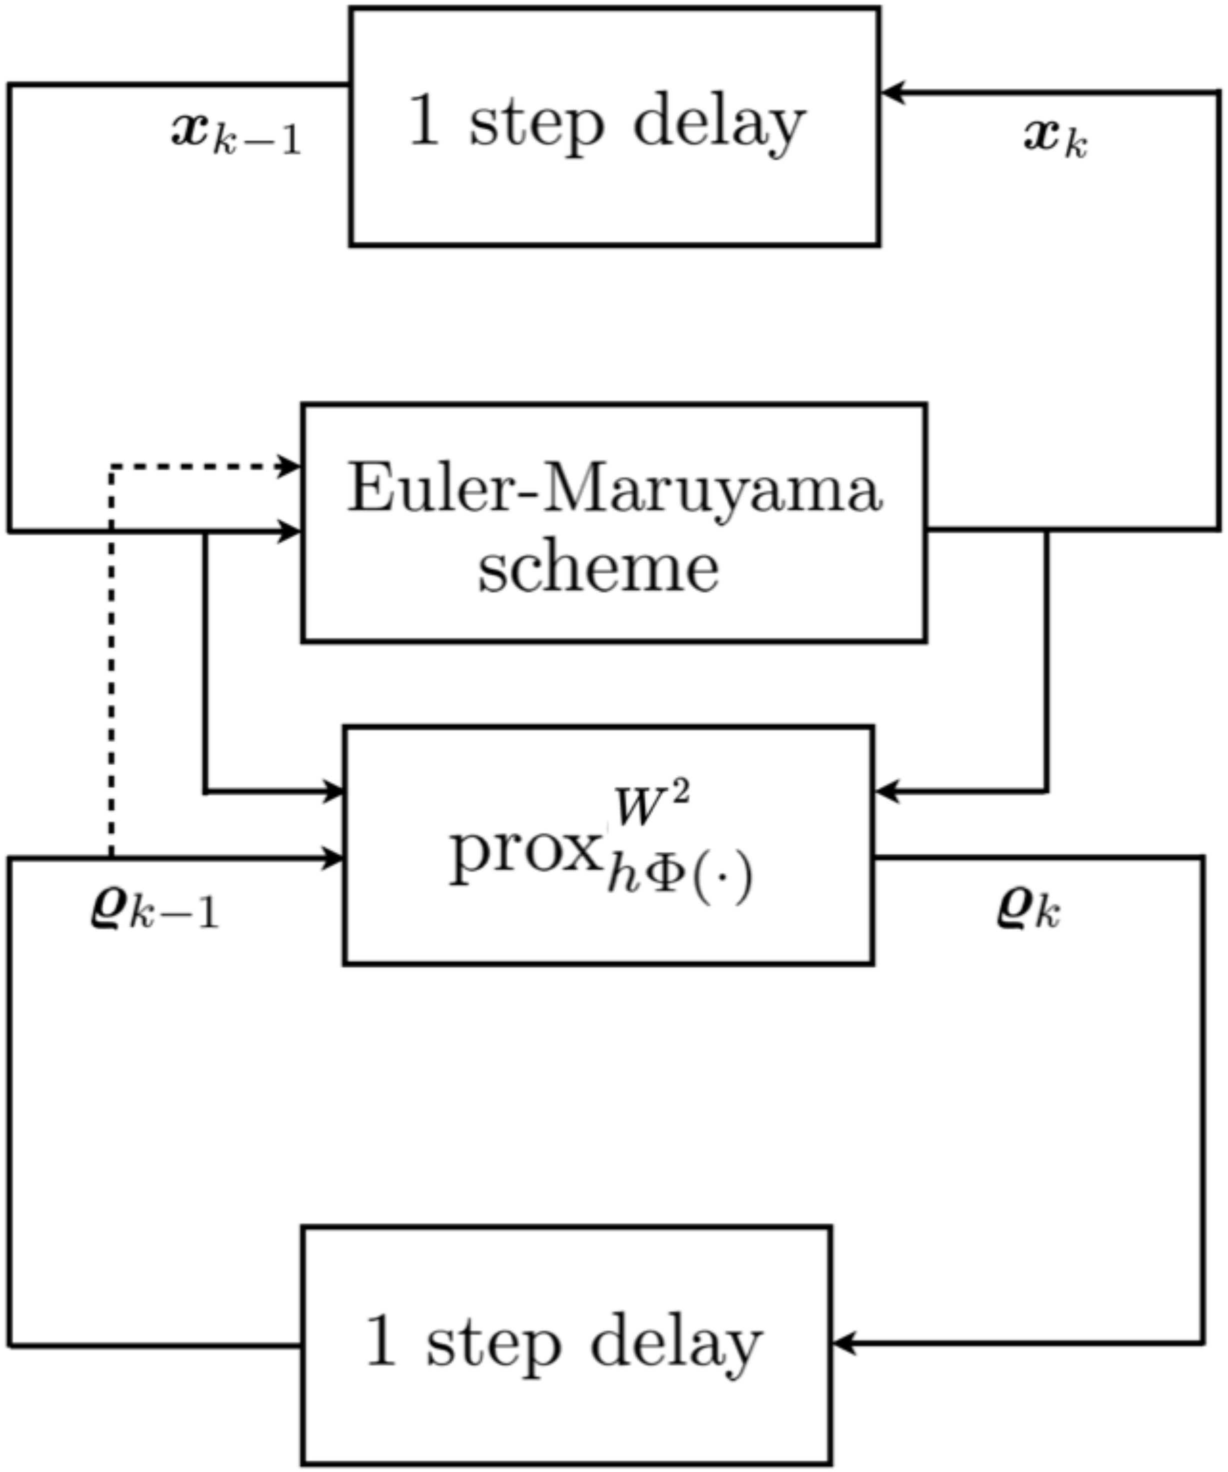

Theorem: Block co-ordinate iteration of $(\mathbf{y}, \mathbf{z})$ recursion is contractive on $\mathbb{R}^n_{>0} \times \mathbb{R}^n_{>0}$.

Proximal Prediction: 2D Linear Gaussian

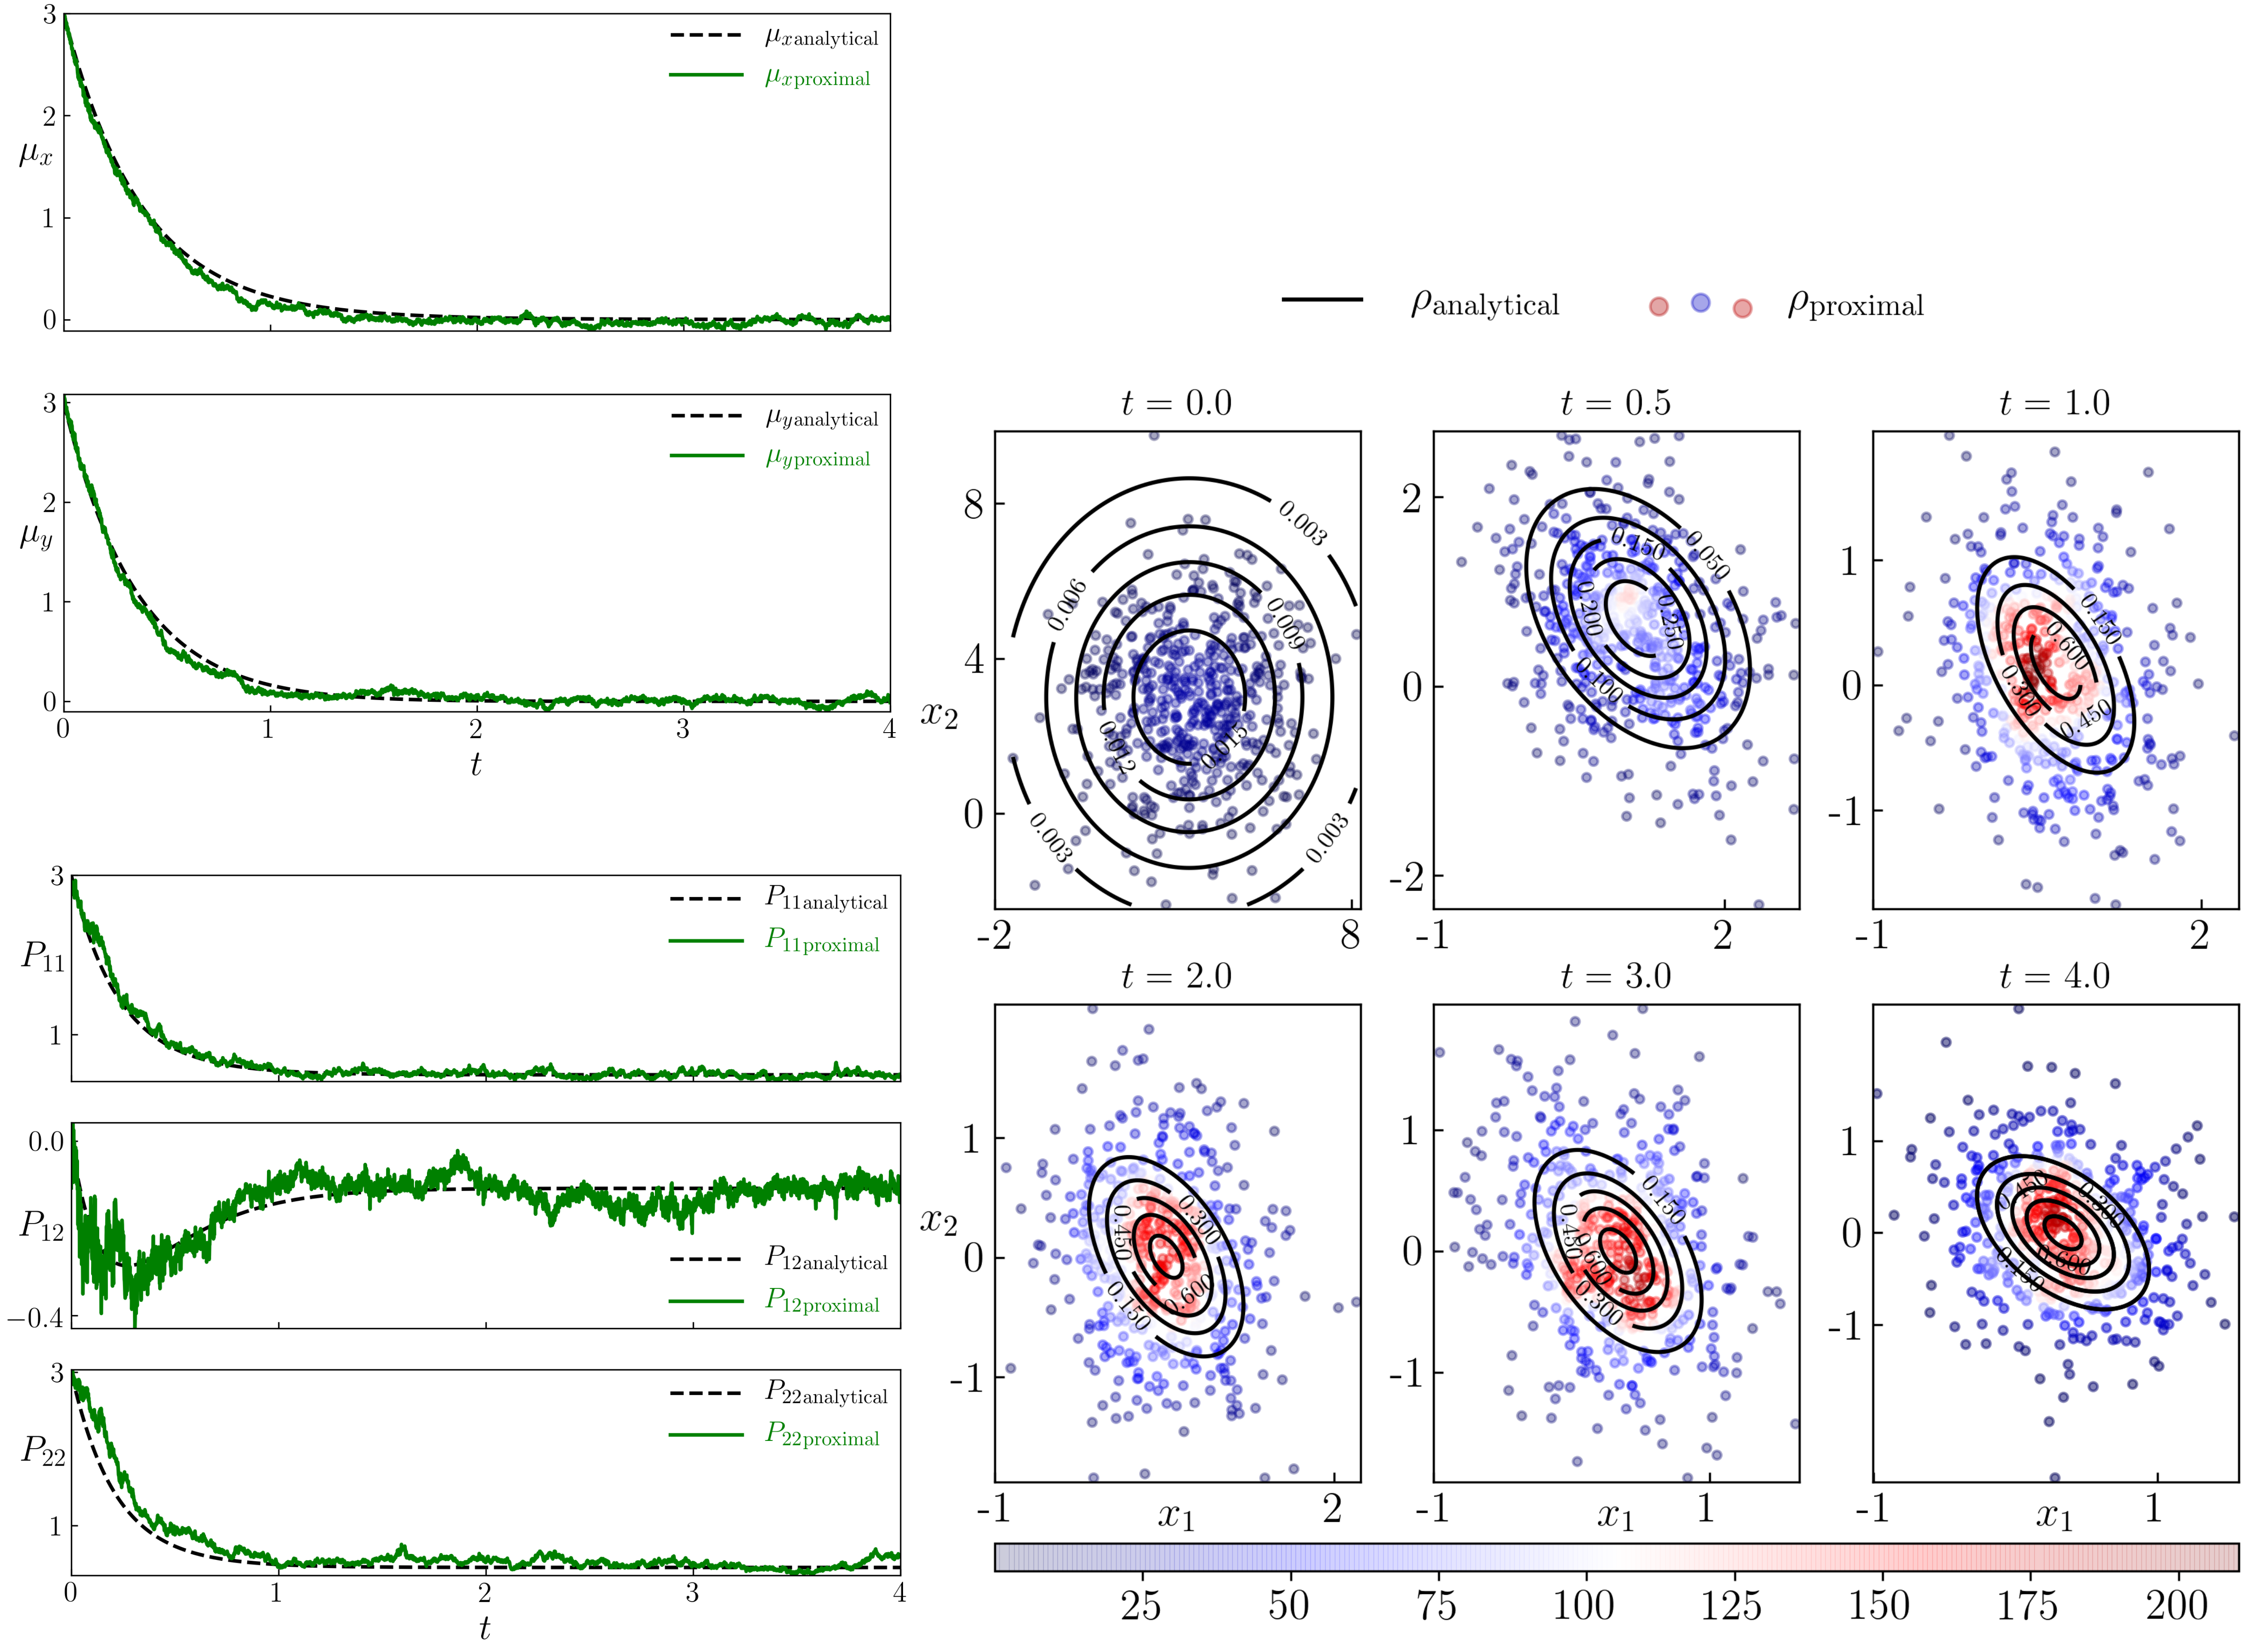

Proximal Prediction: Nonlinear Non-Gaussian

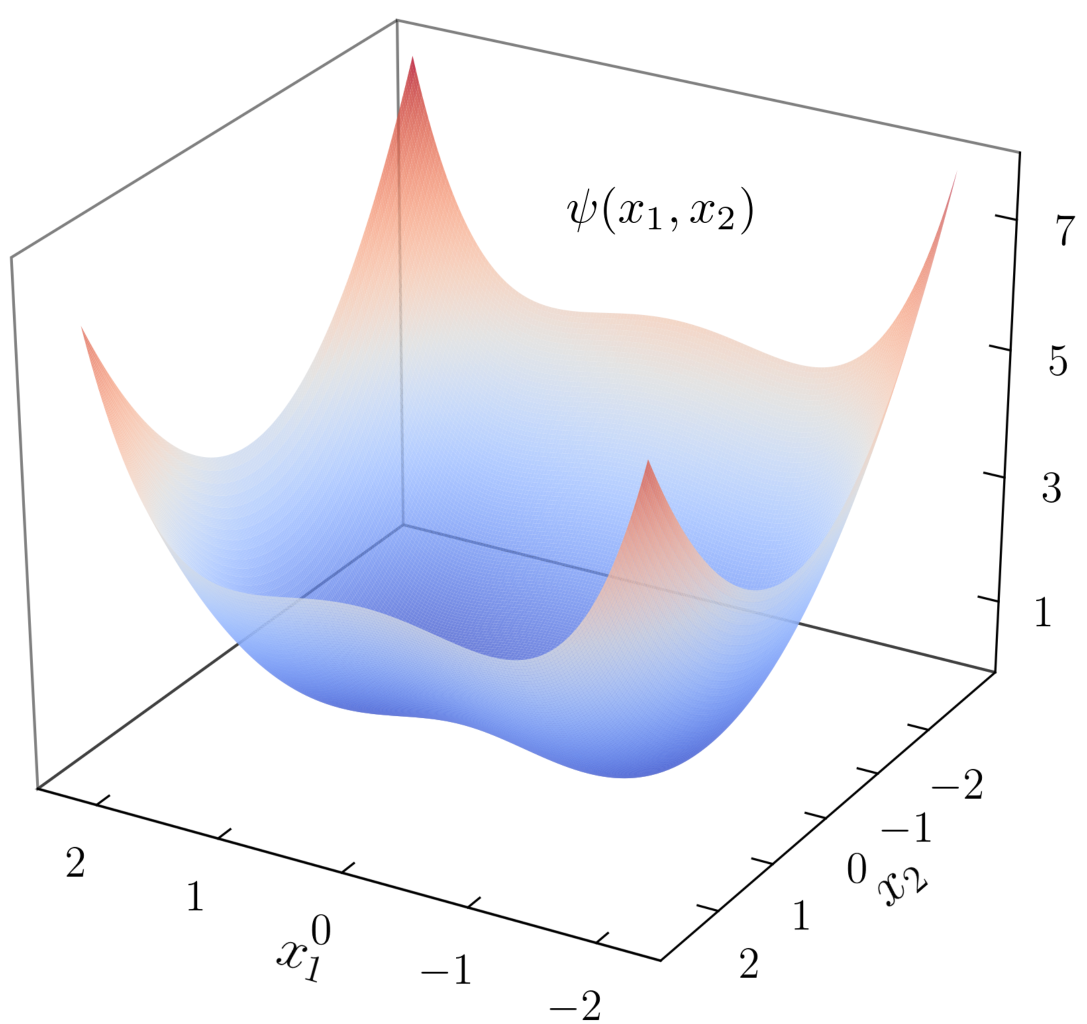

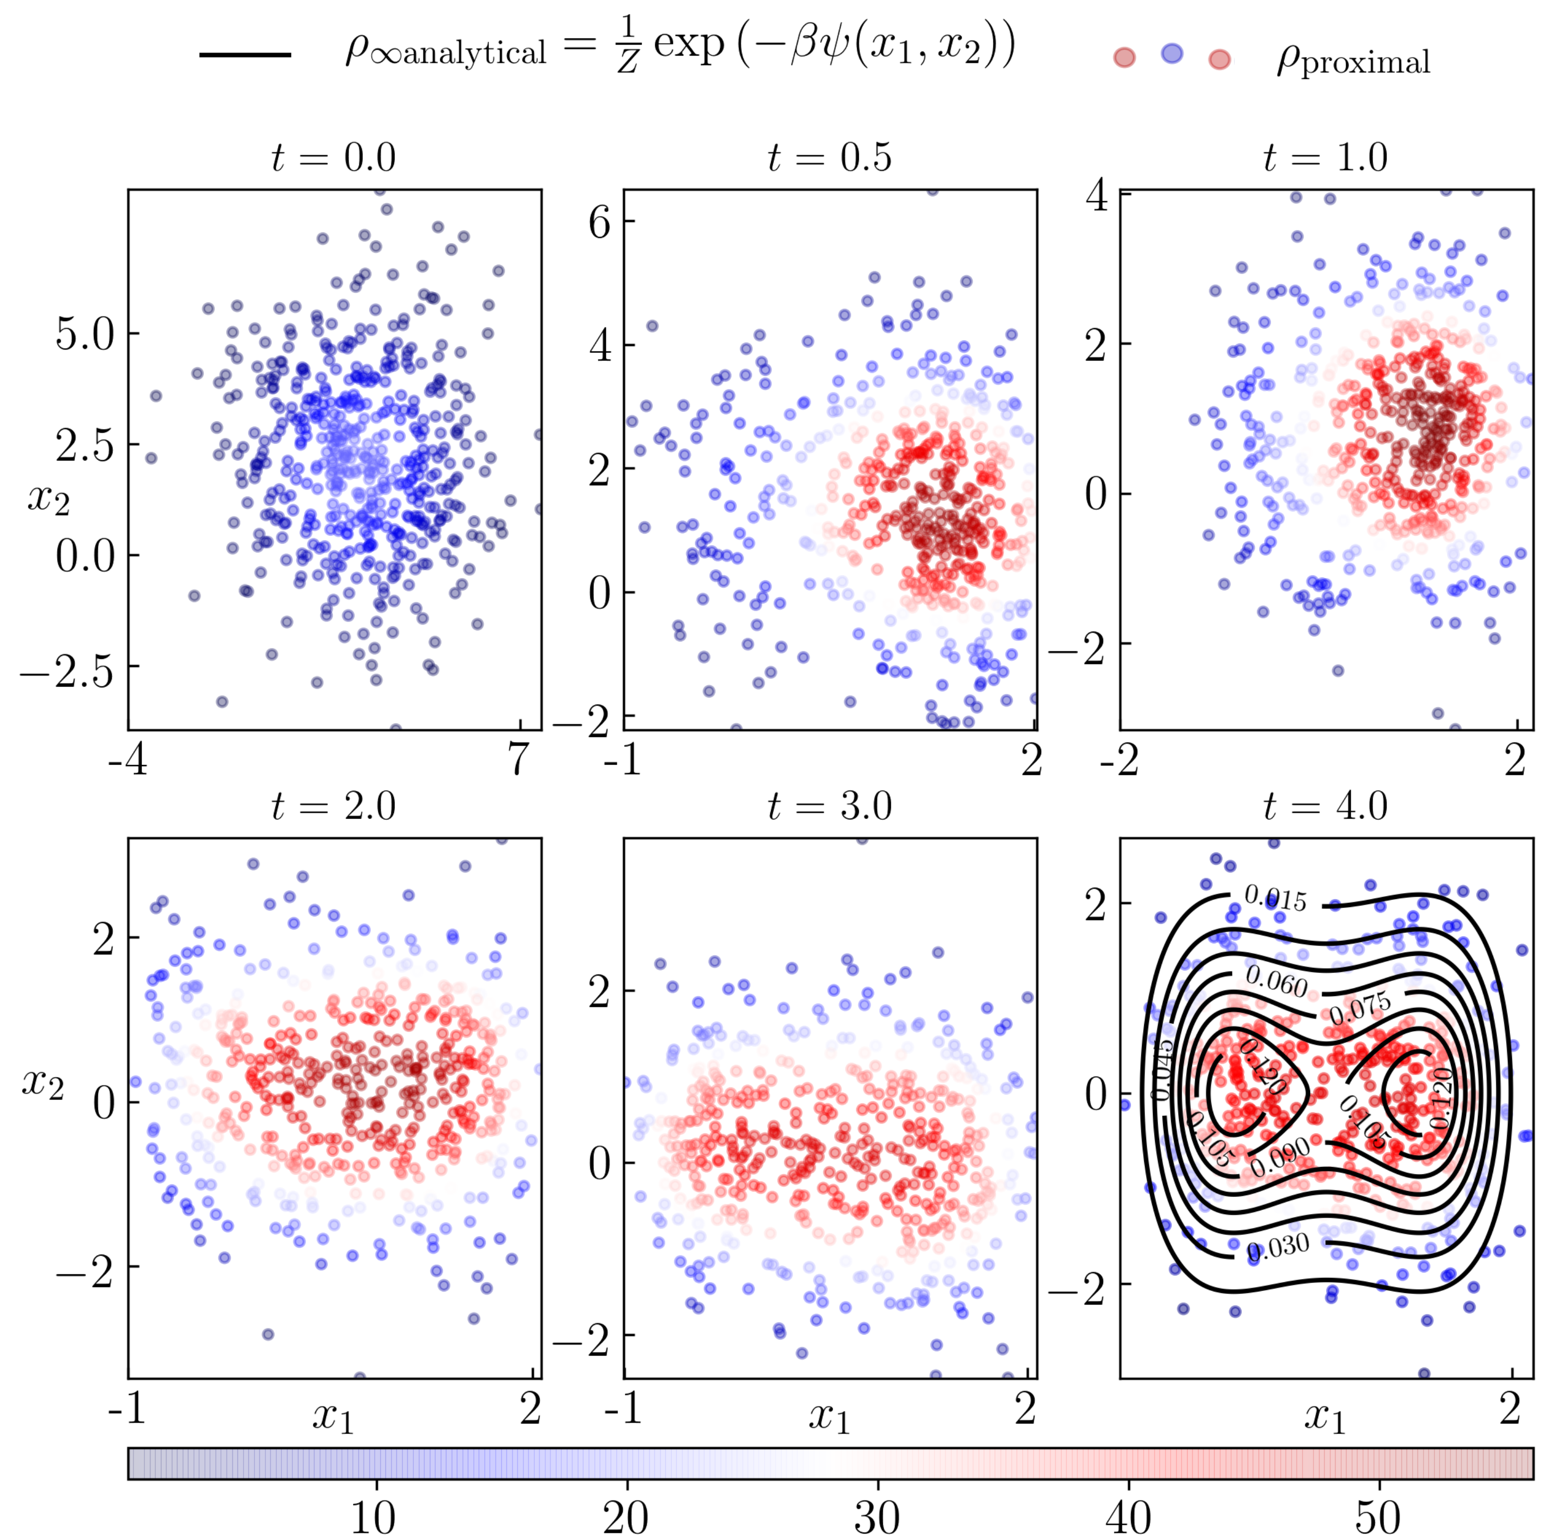

Computational Time: Nonlinear Non-Gaussian

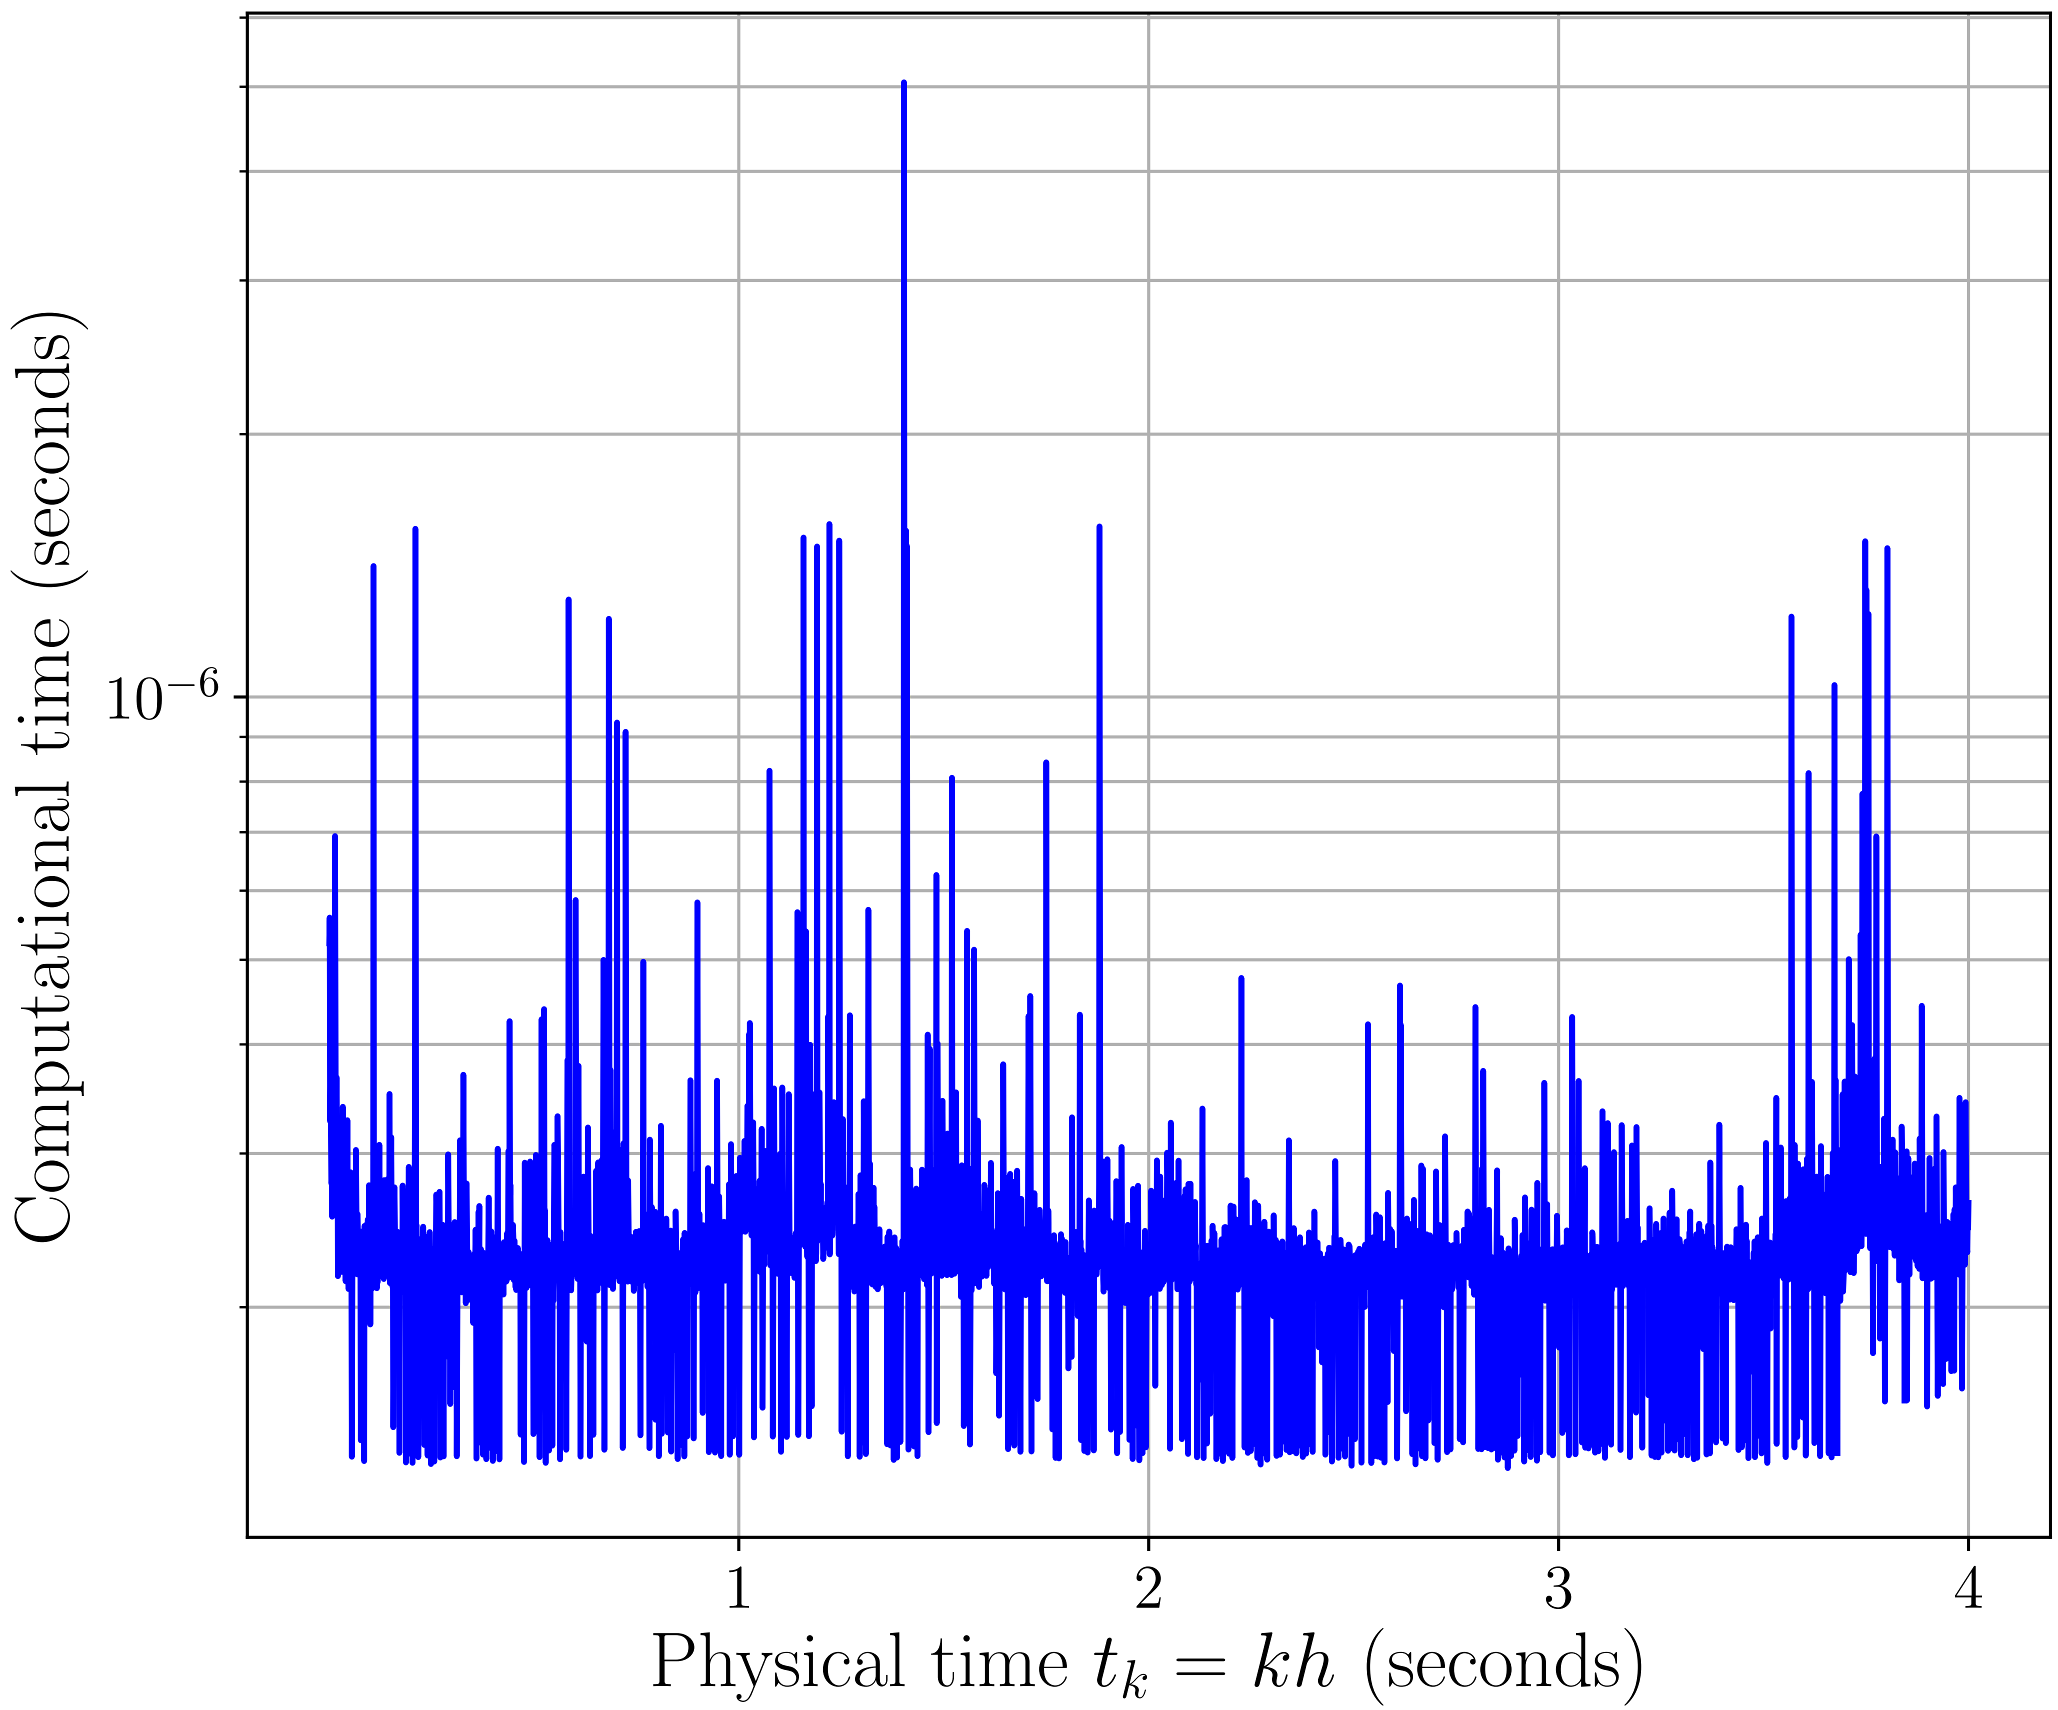

Proximal Prediction: Satellite in Geocentric Orbit

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

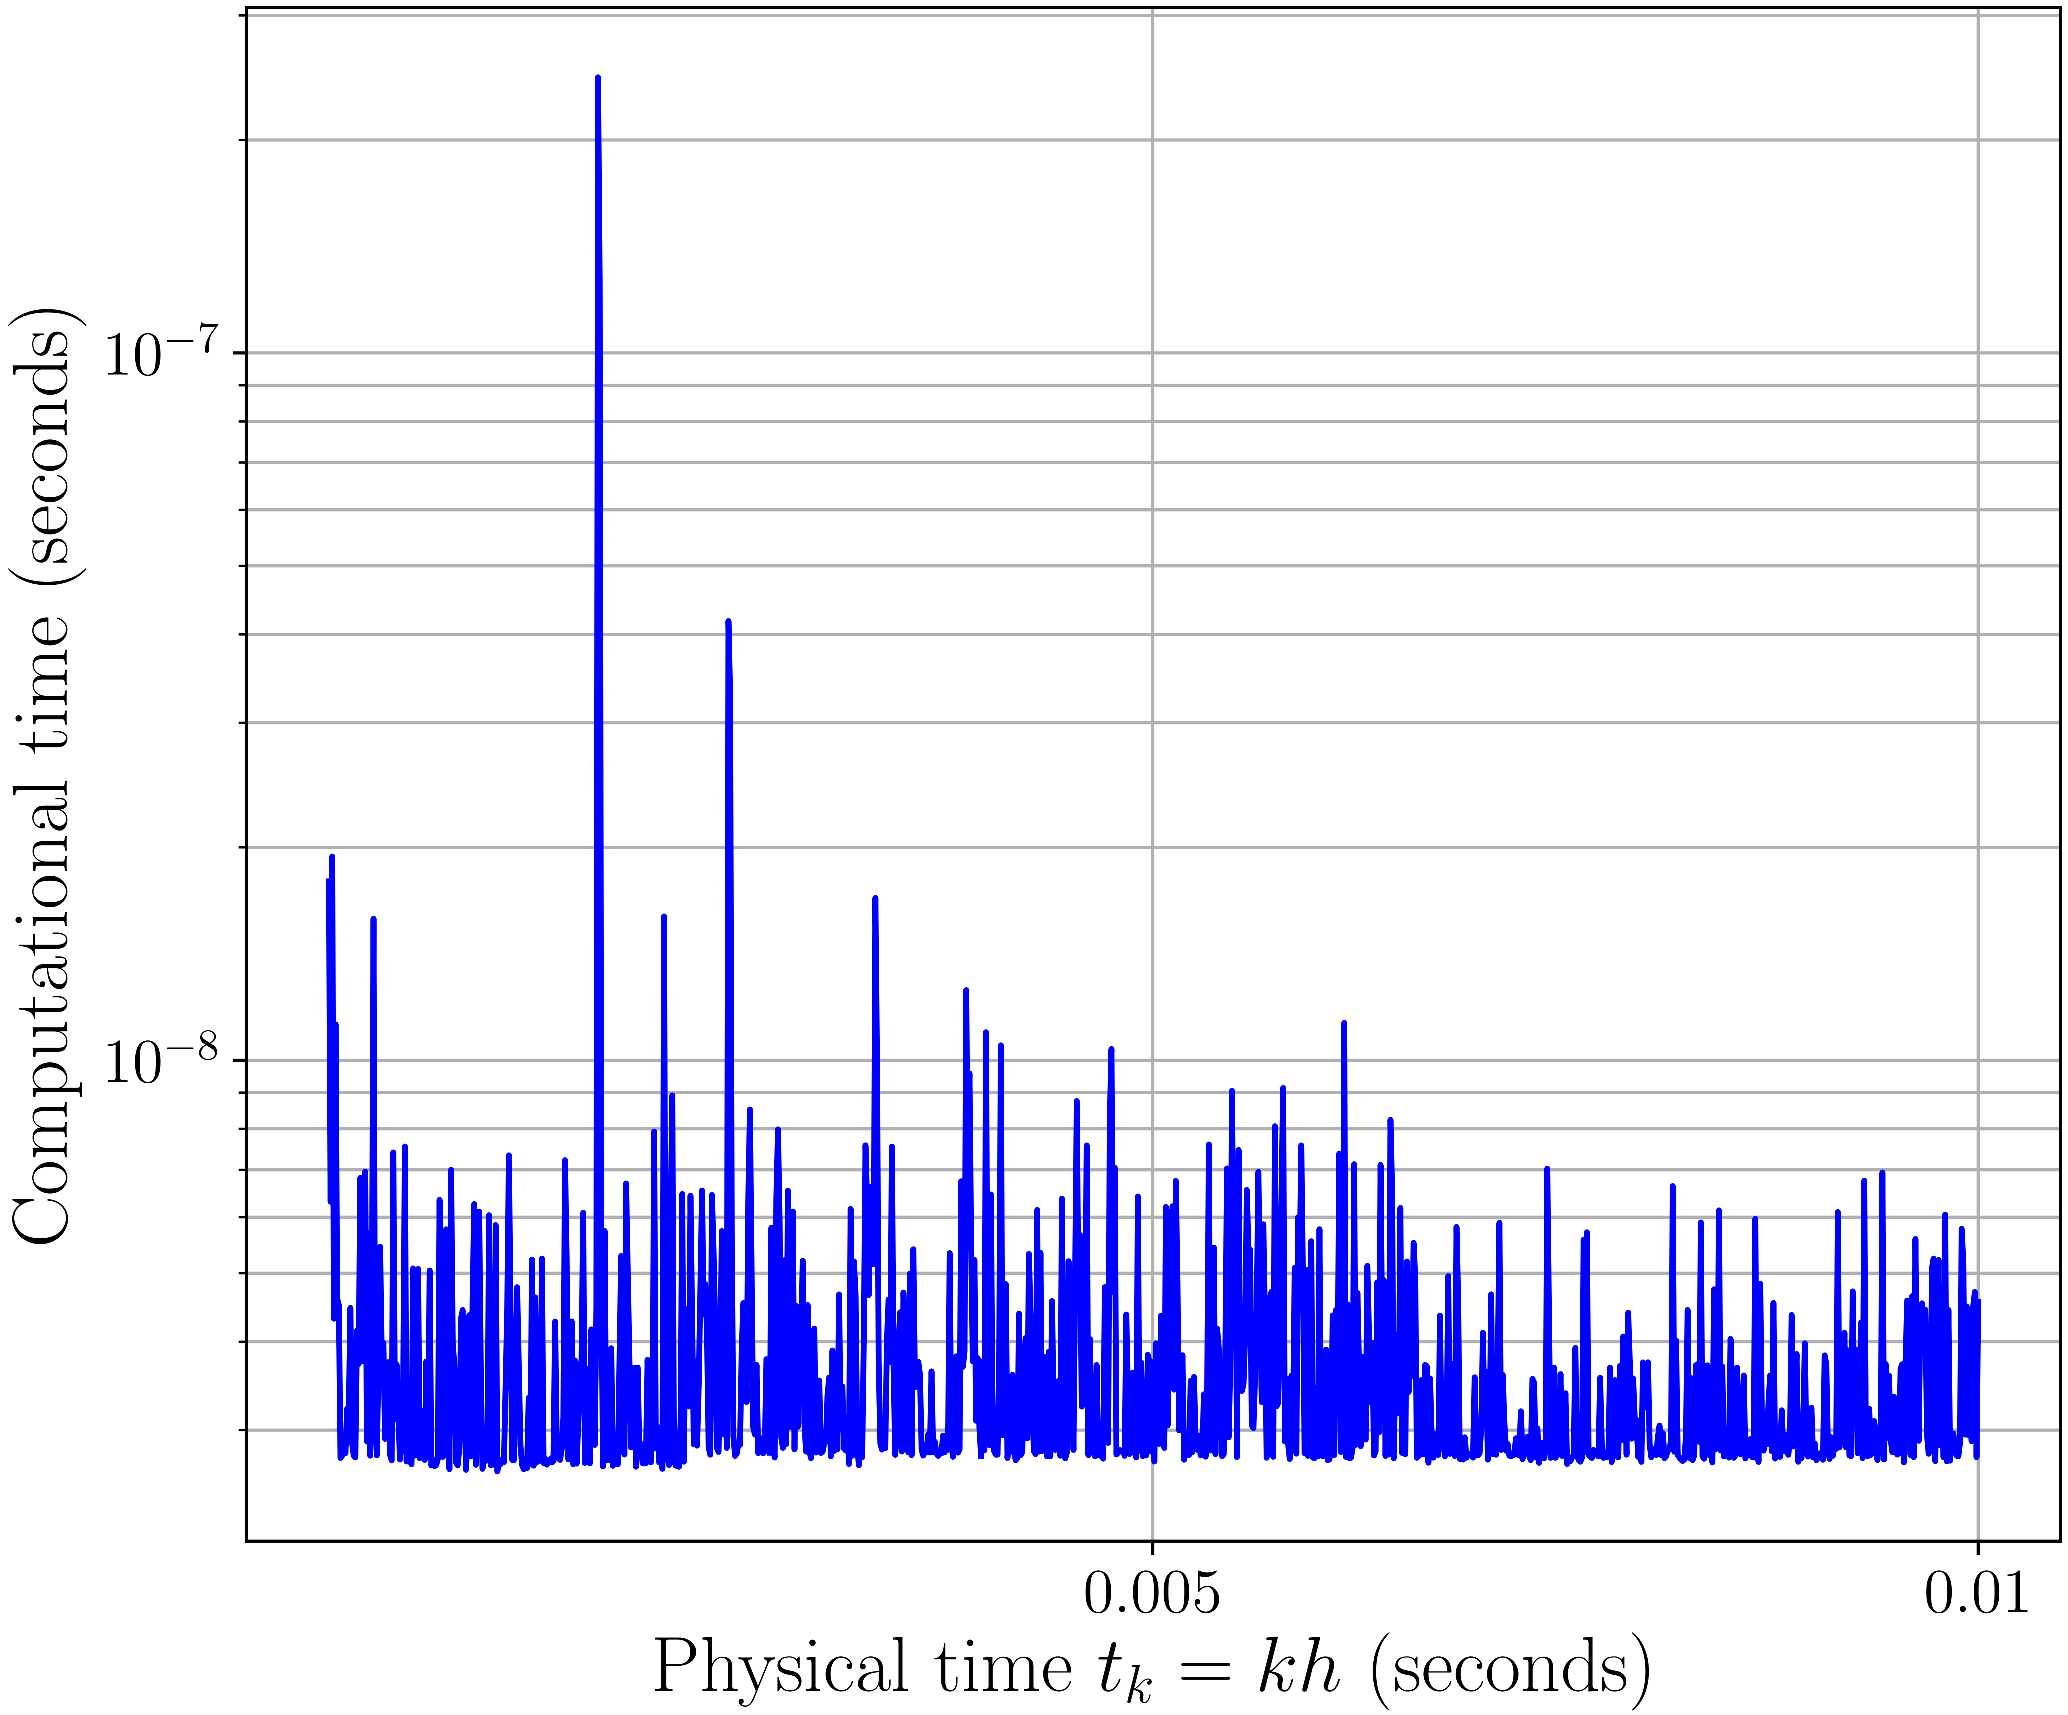

Extensions: Nonlocal Interactions

PDF dependent sample path dynamics:

$$\mathrm{d}\mathbf{x} = -\left(\nabla U\left(\mathbf{x}\right) + \nabla\rho * V\right)\mathrm{d}t + \sqrt{2\beta^{-1}}\,\mathrm{d}\mathbf{w}$$

Mckean-Vlasov-Fokker-Planck-Kolmogorov integro PDE:

$$\frac{\partial\rho}{\partial t} = \nabla\cdot\left(\rho\nabla\left(U + \rho * V\right)\right) + \beta^{-1}\Delta\rho$$

Free energy:

$$F(\rho) := \mathbb{E}_{\rho}\left[U + \beta^{-1}\rho\log\rho + \rho * V\right]$$

Extensions: Nonlocal Interactions

$$U(\cdot) = V(\cdot) = \| \cdot \|_2^2$$

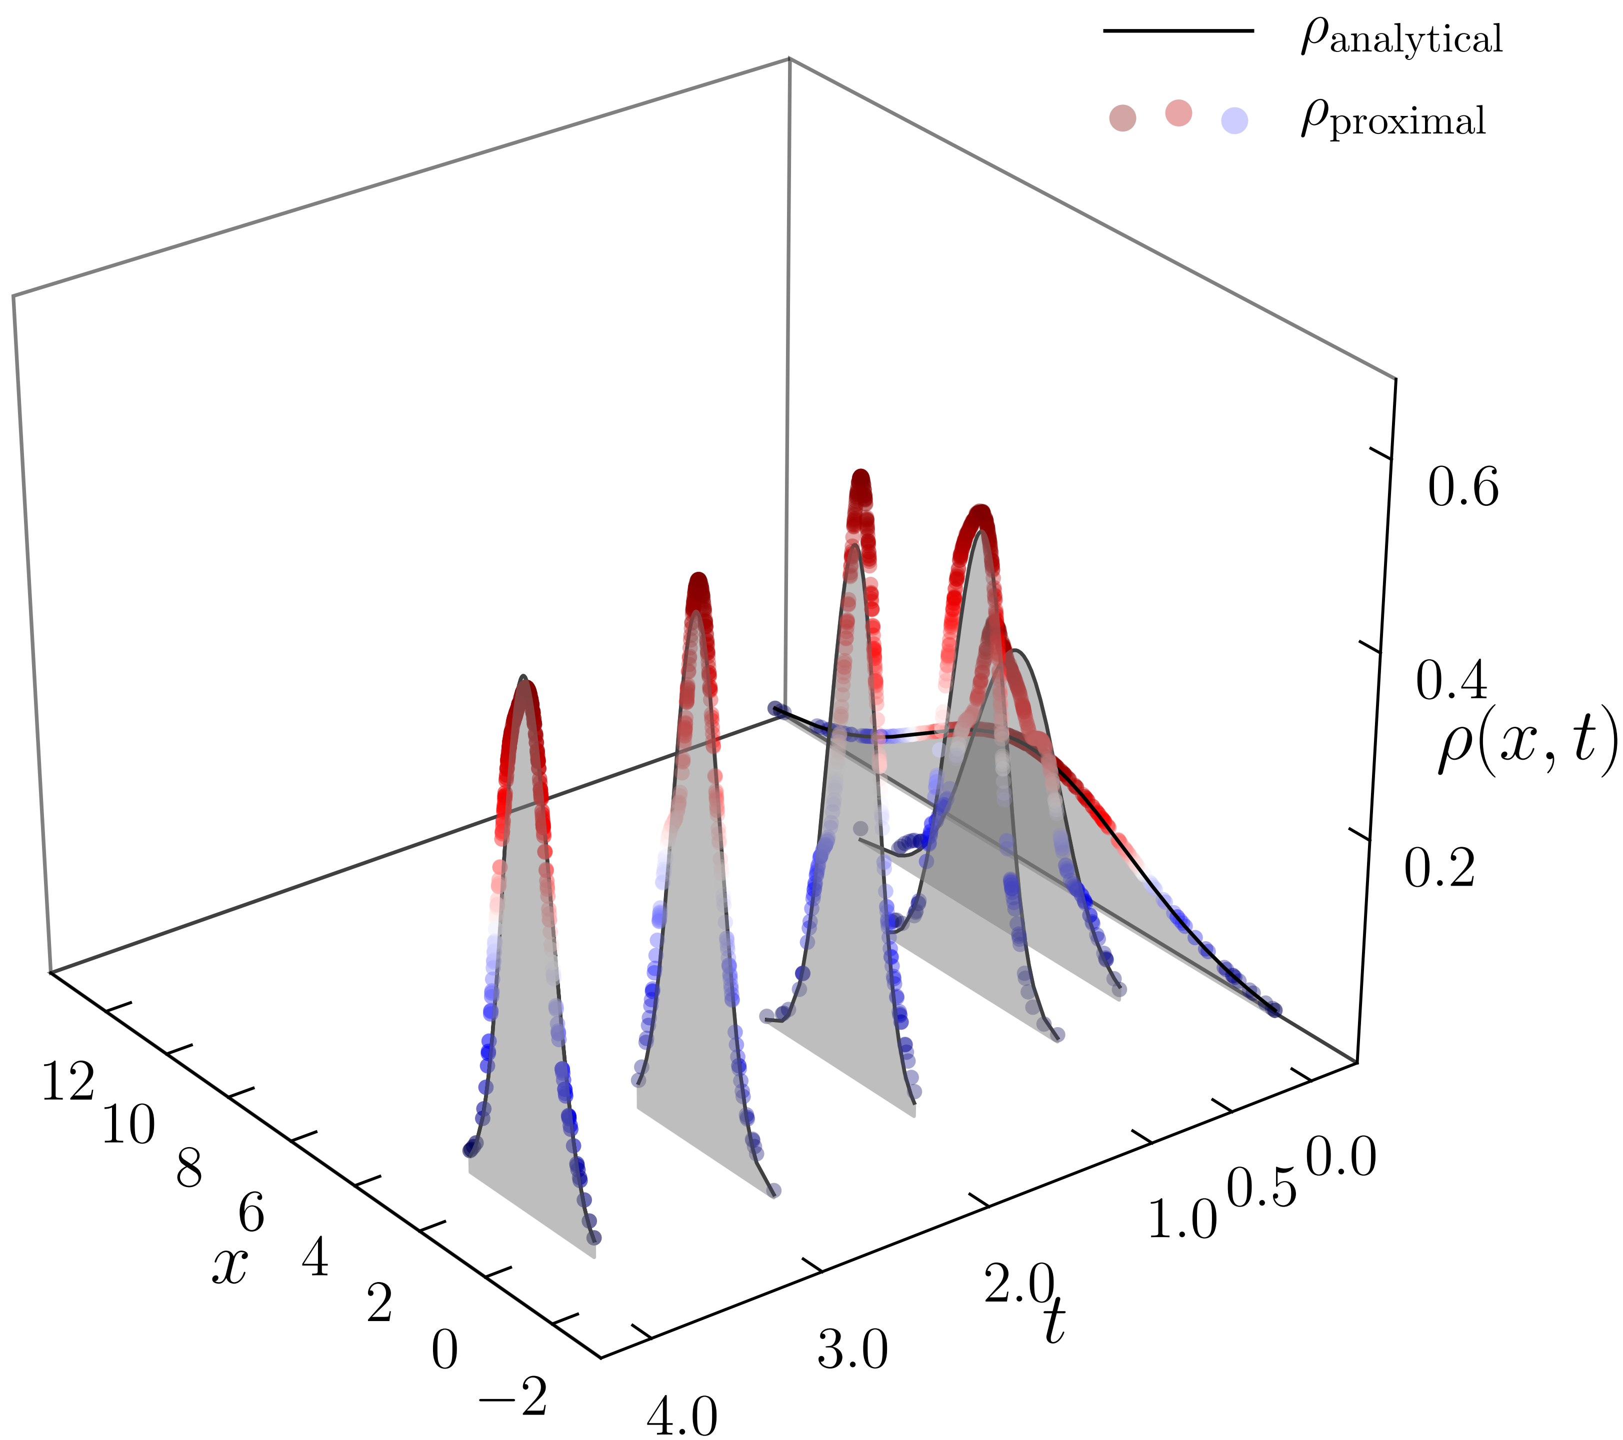

Extensions: Multiplicative Noise

Cox-Ingersoll-Ross: $\mathrm{d}x = a\,(\theta - x)\ \mathrm{d}t + b\sqrt{x}\,\mathrm{d}w,\ 2a > b^2, \theta > 0$

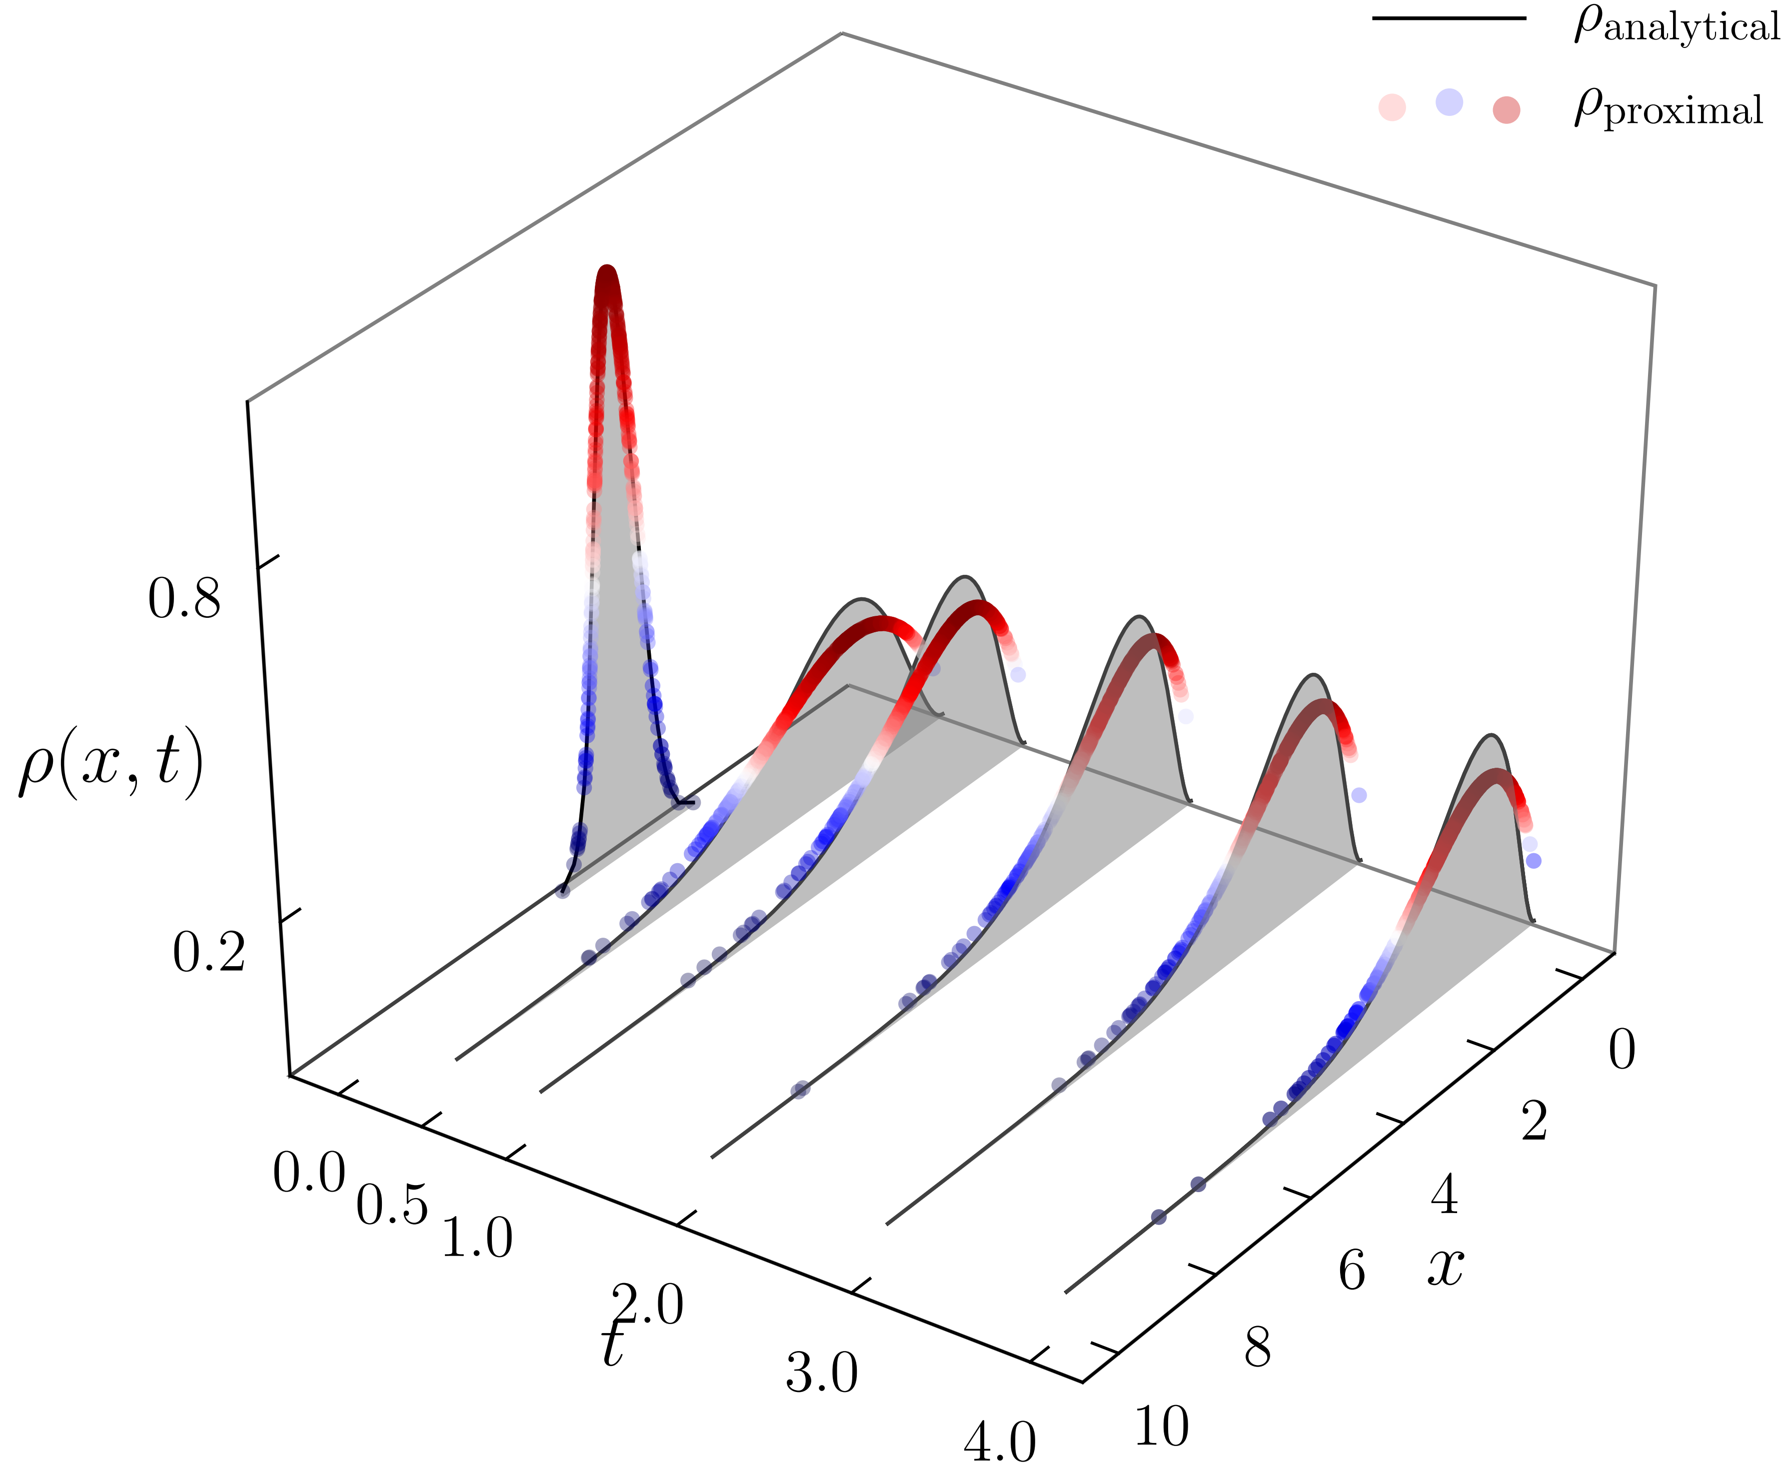

Solving filtering as generalized gradient flow

What's New?

Main idea: Solve the Kushner-Stratonovich SPDE

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

Recursion of {deterministic ∘ stochastic} proximal operators:

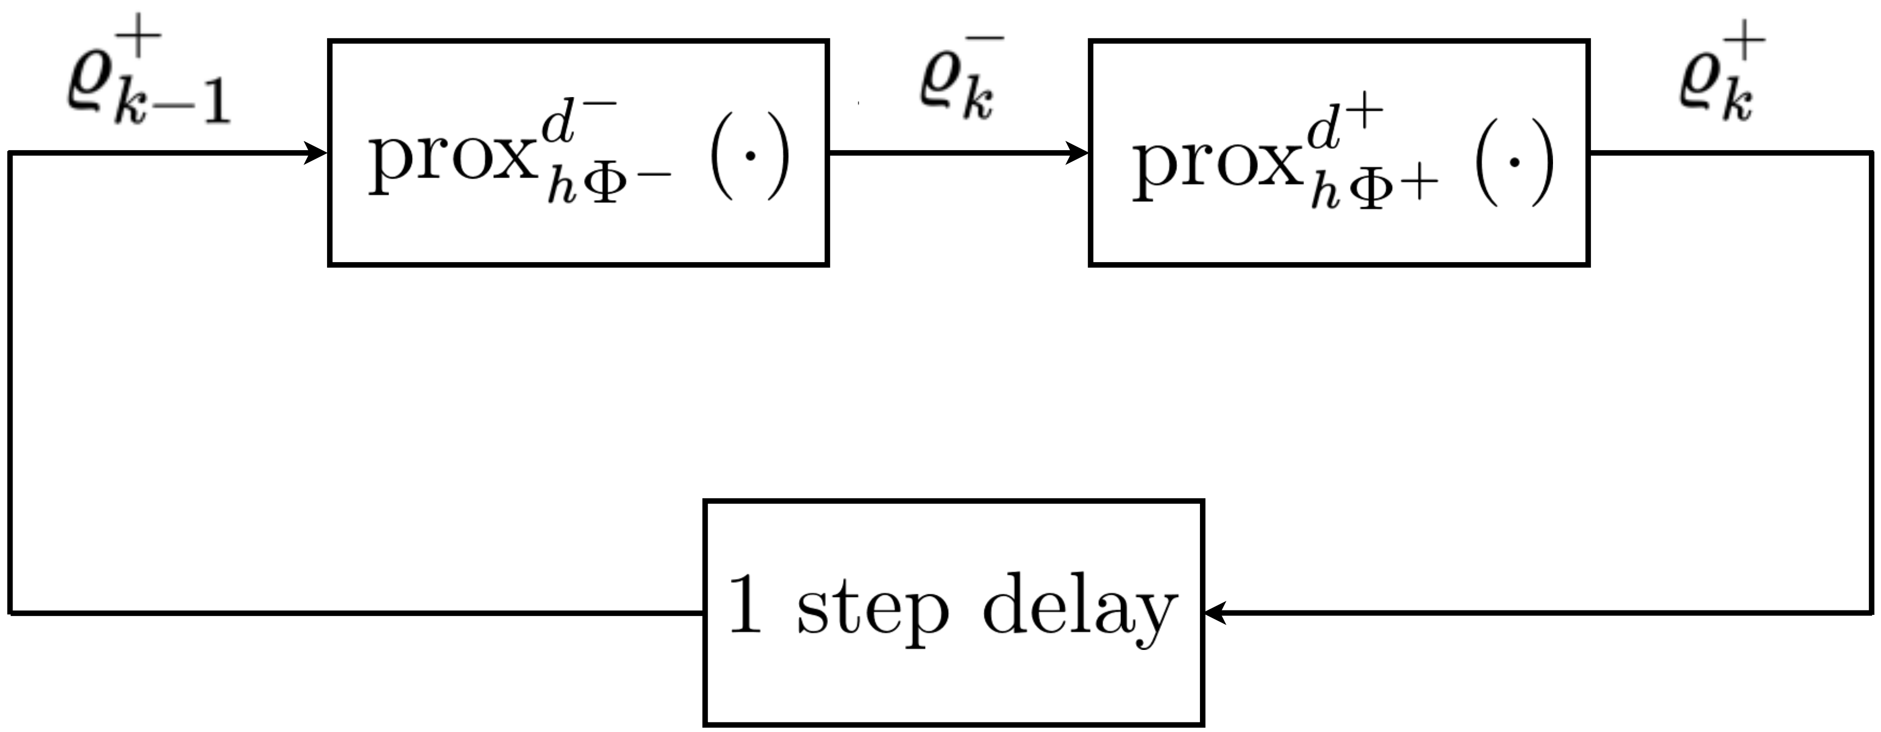

Convergence: $\varrho_k^+(h) \to \rho^+(x, t = kh) \quad \text{as} \quad h \downarrow 0$

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:

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

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

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:

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

> 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

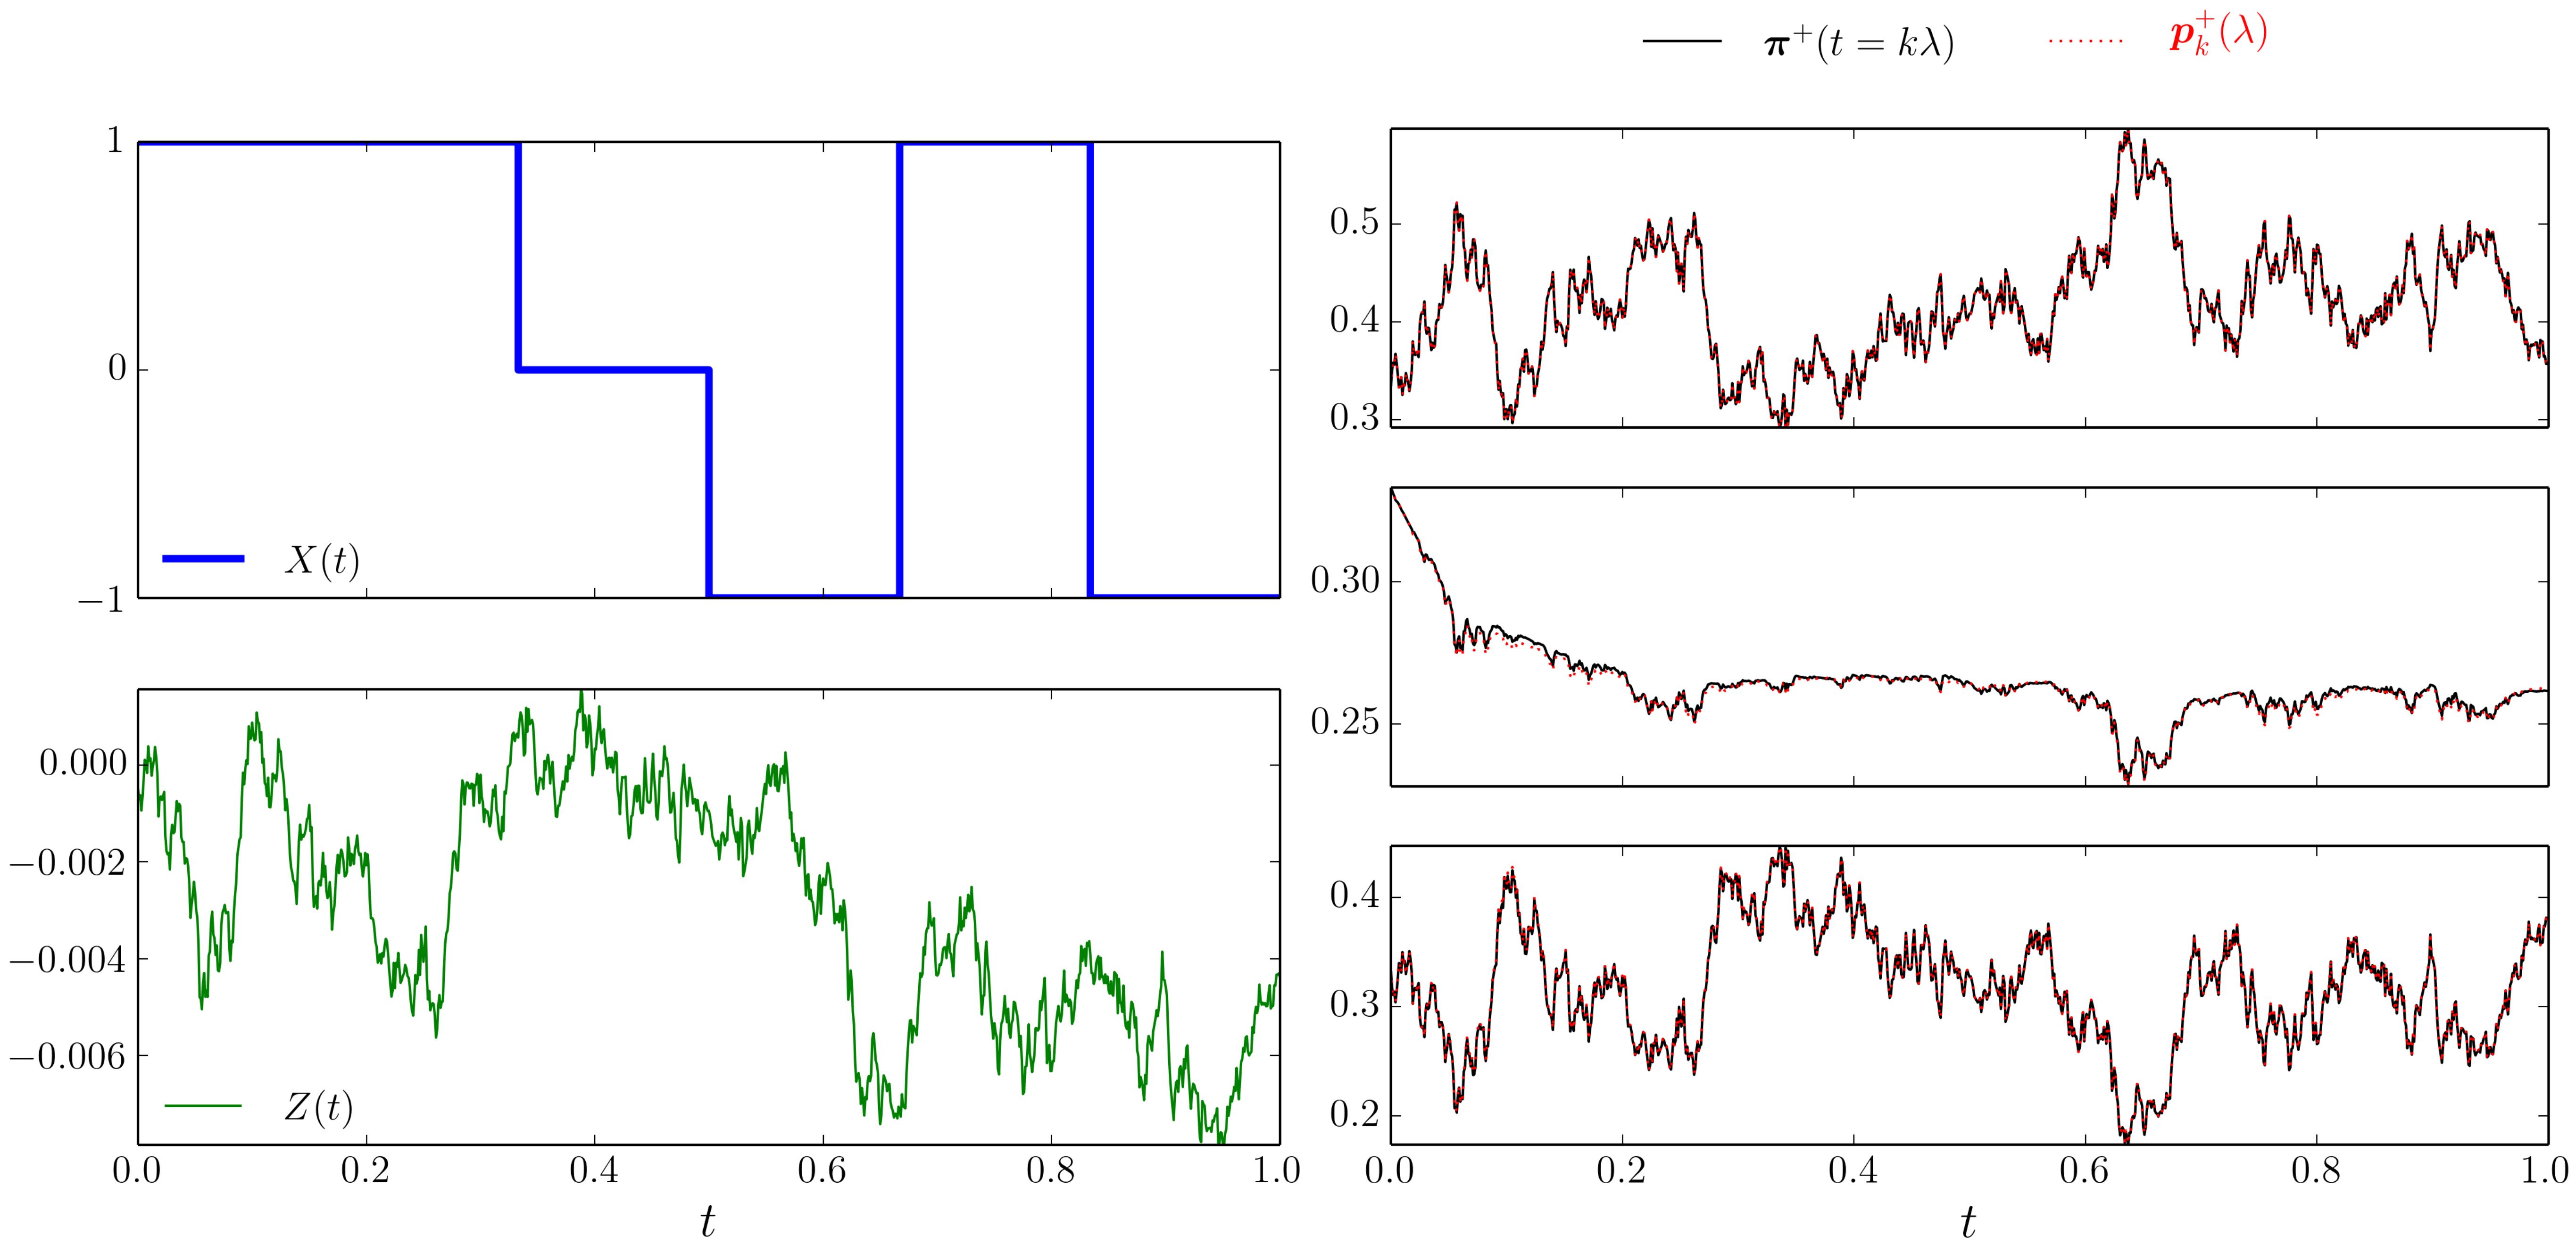

— A.H. and T.T. Georgiou, Proximal Recursion for the Wonham Filter, *CDC 2019*.

Solving density control as generalized gradient flow

State Feedback Density Steering

Steer joint state PDF via feedback control over finite time horizon

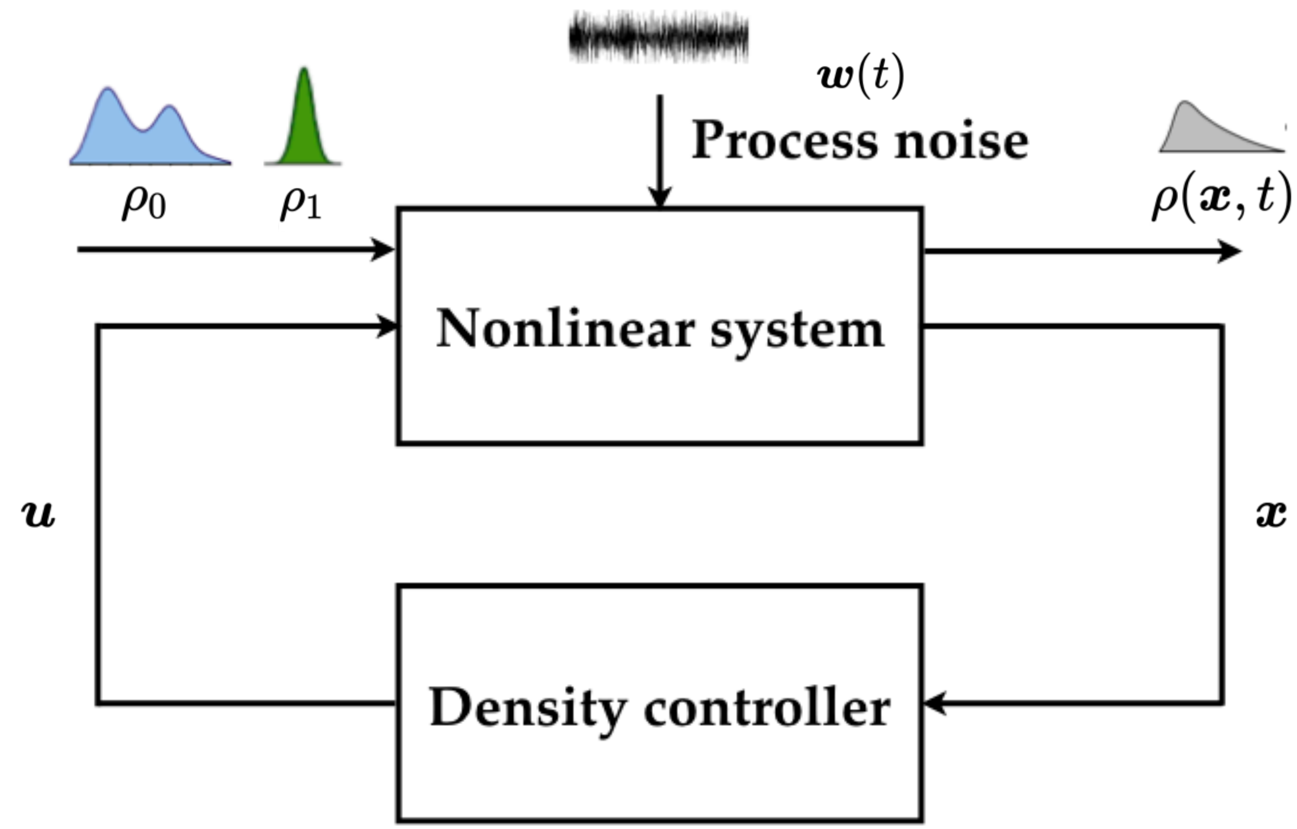

Common scenario: $\boldsymbol{G} \equiv \boldsymbol{B}$

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

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

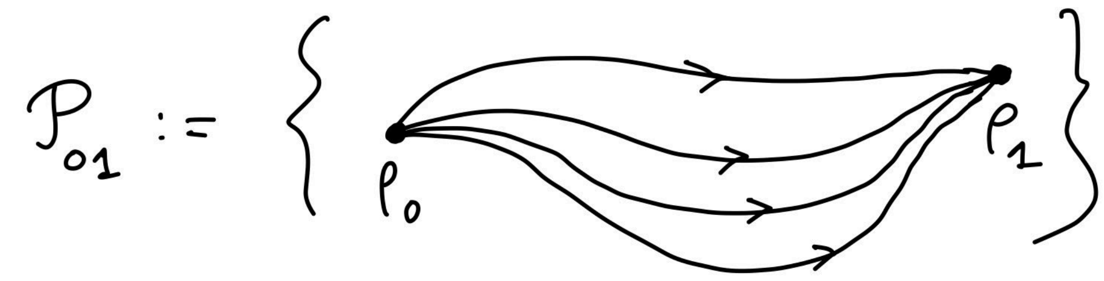

Diffusion tensor: $\boldsymbol{D} := \boldsymbol{G}\boldsymbol{G}^{\top}$

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

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

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

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

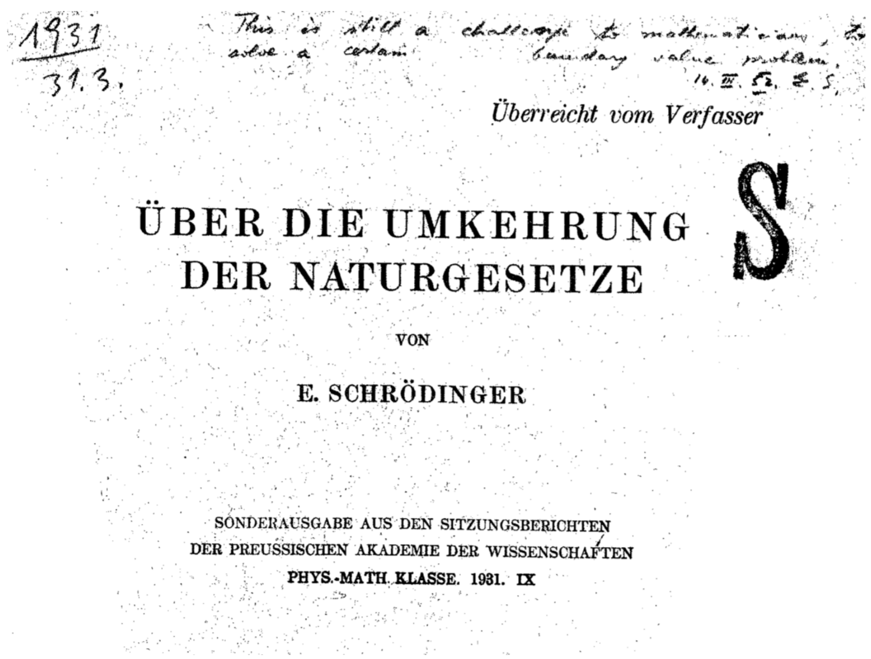

1931
31.3.

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

Überreicht vom Verfasser

ÜBER DIE UMKEHRUNG DER NATURGESETZE

VON

E. SCHRÖDINGER

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

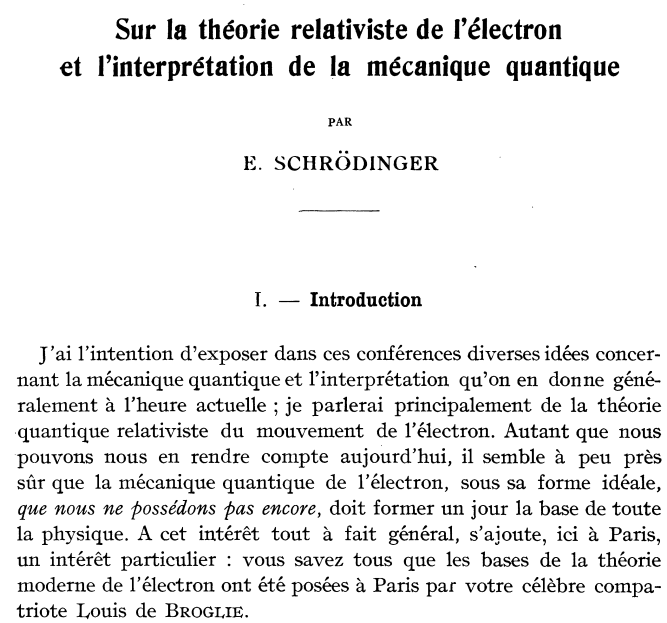

Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique

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

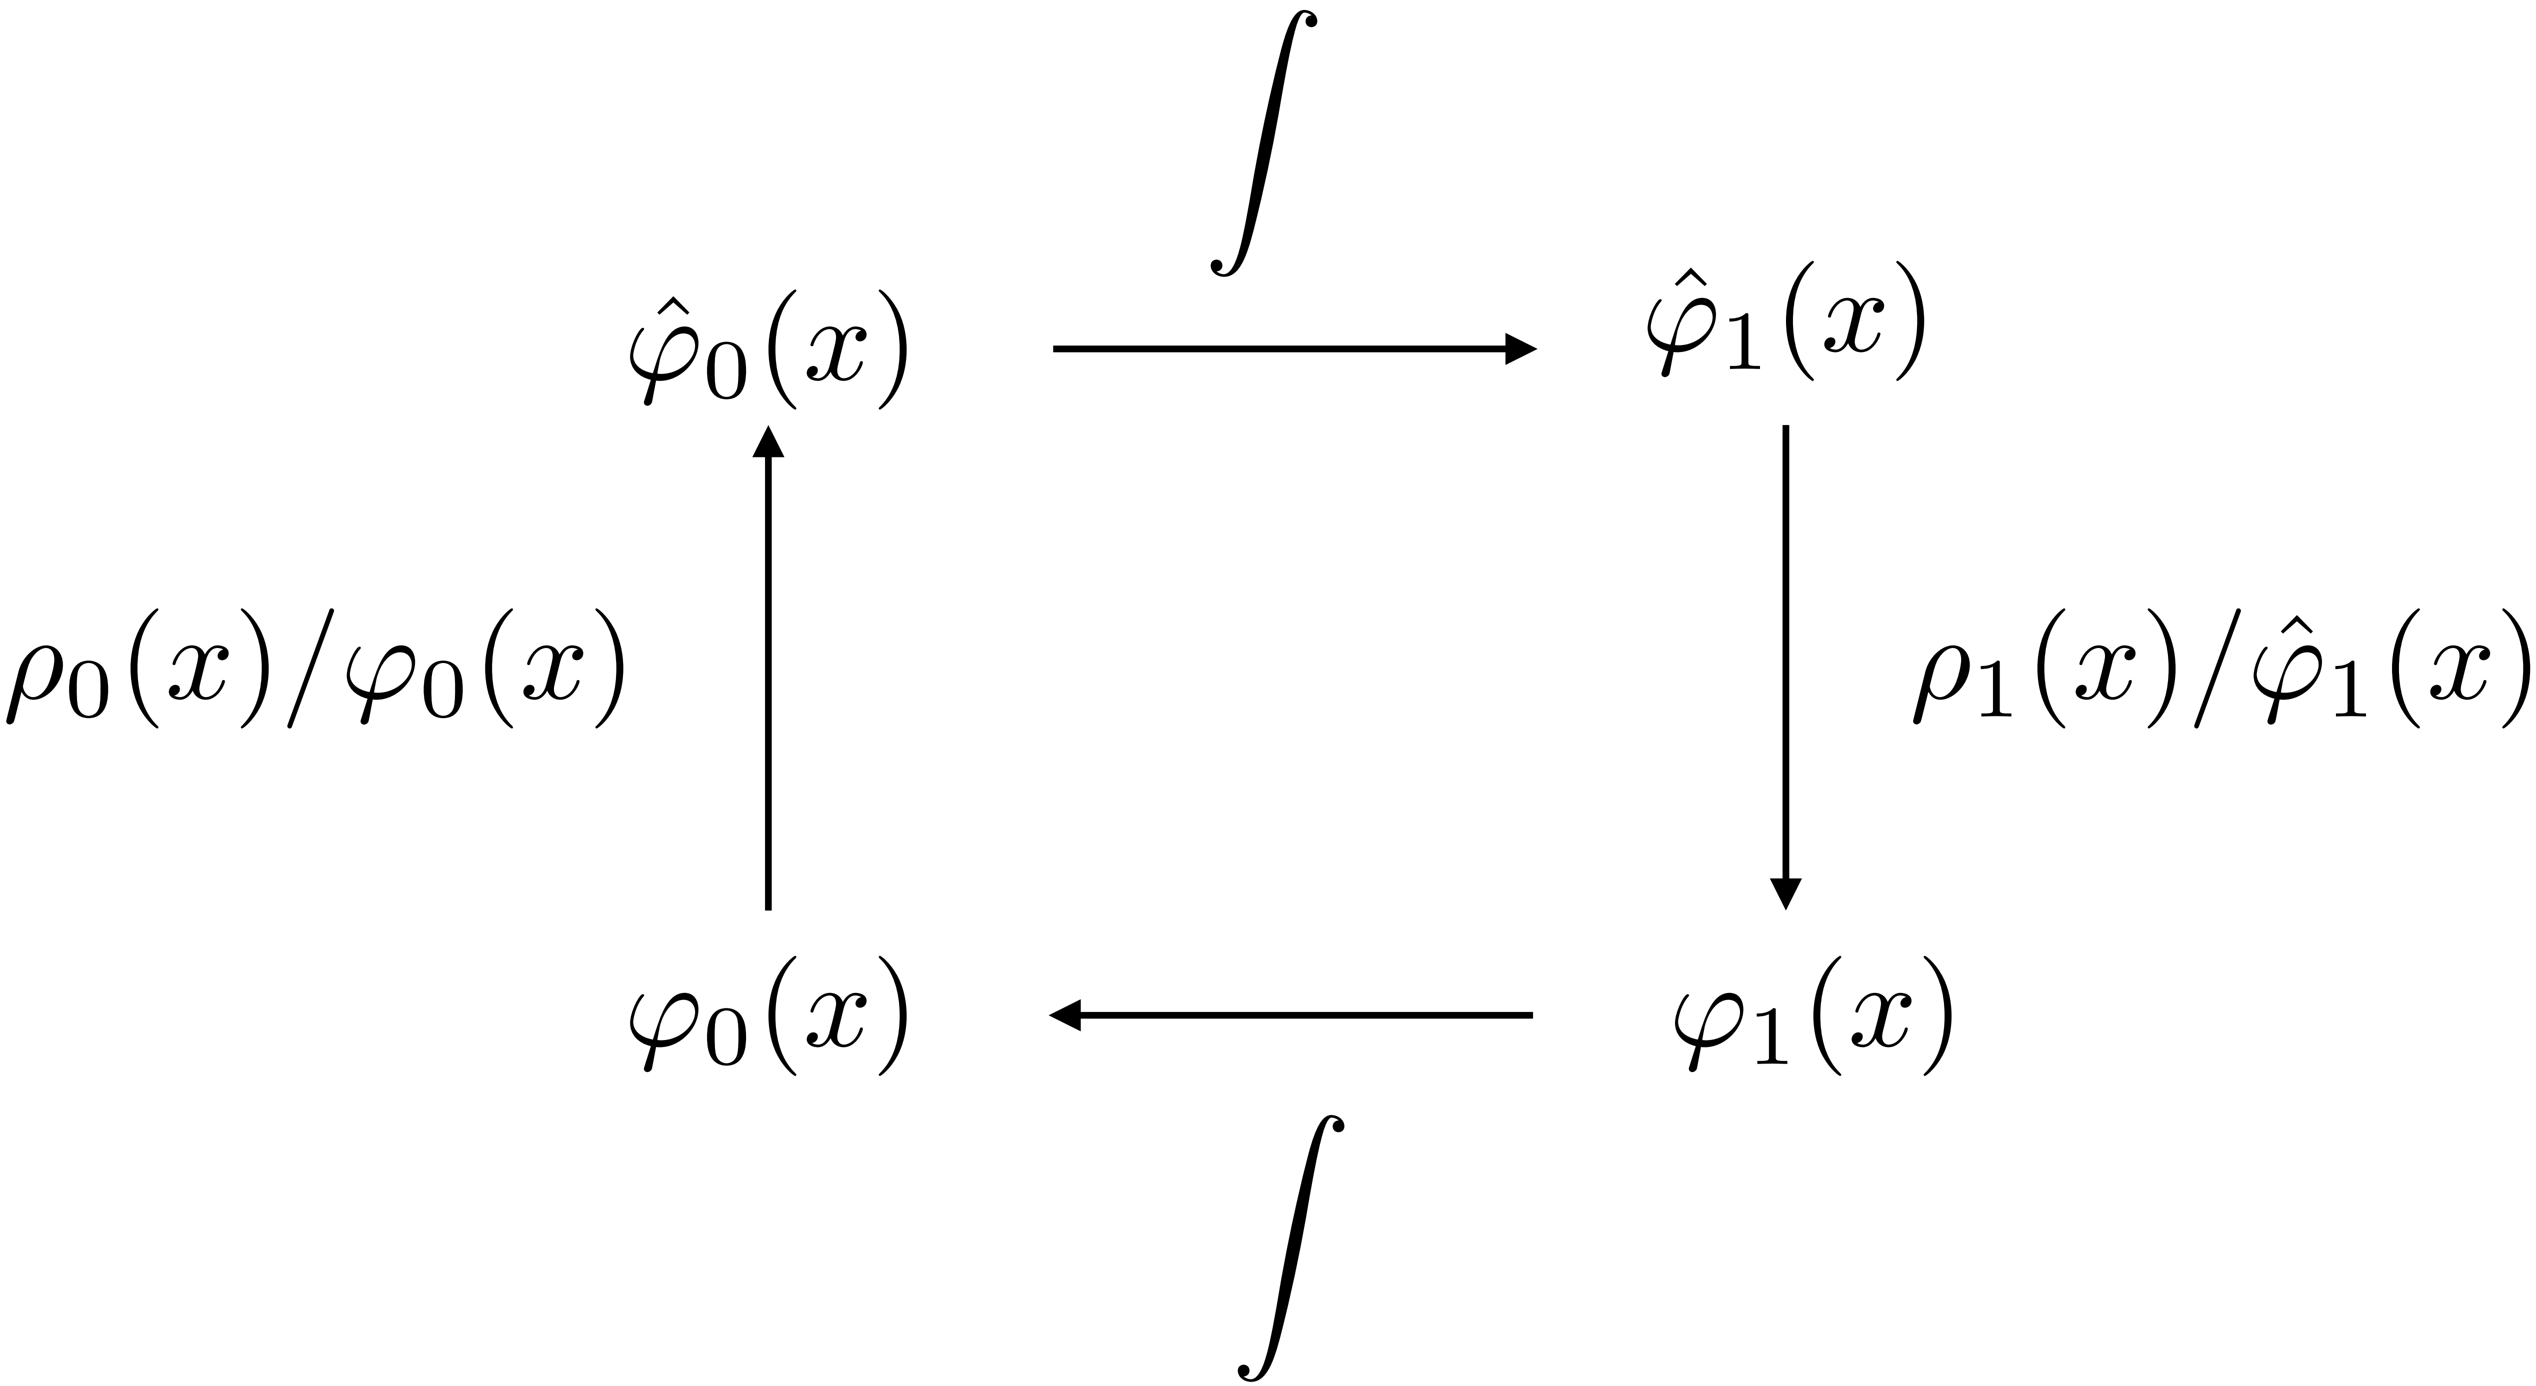

This recursion is contractive in the Hilbert metric!!

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

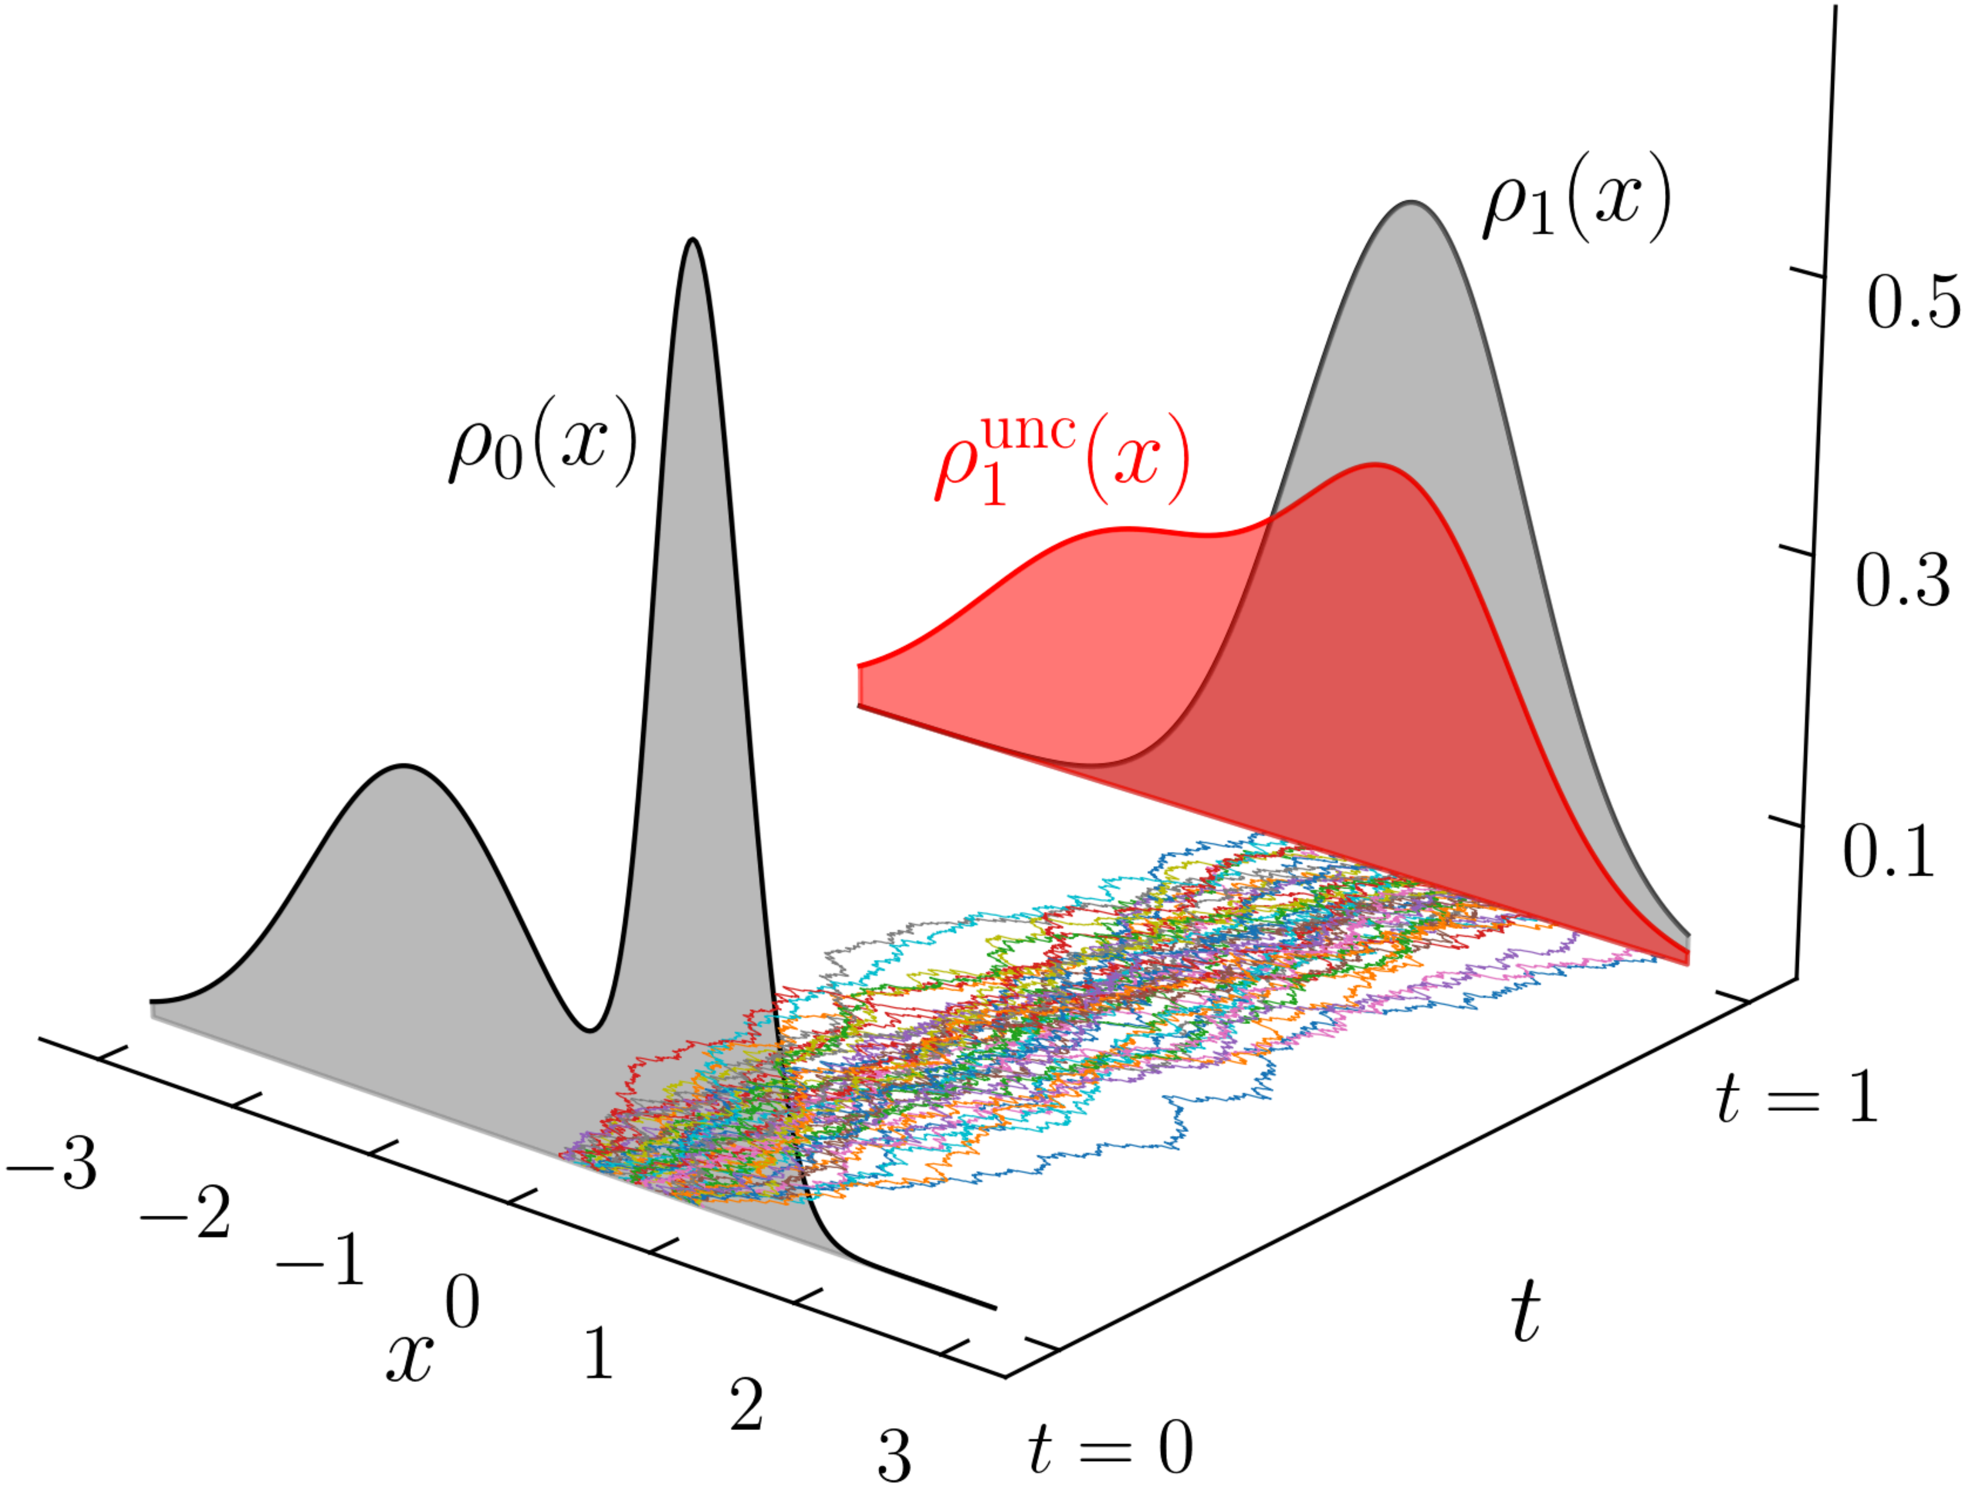

Zero prior dynamics

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

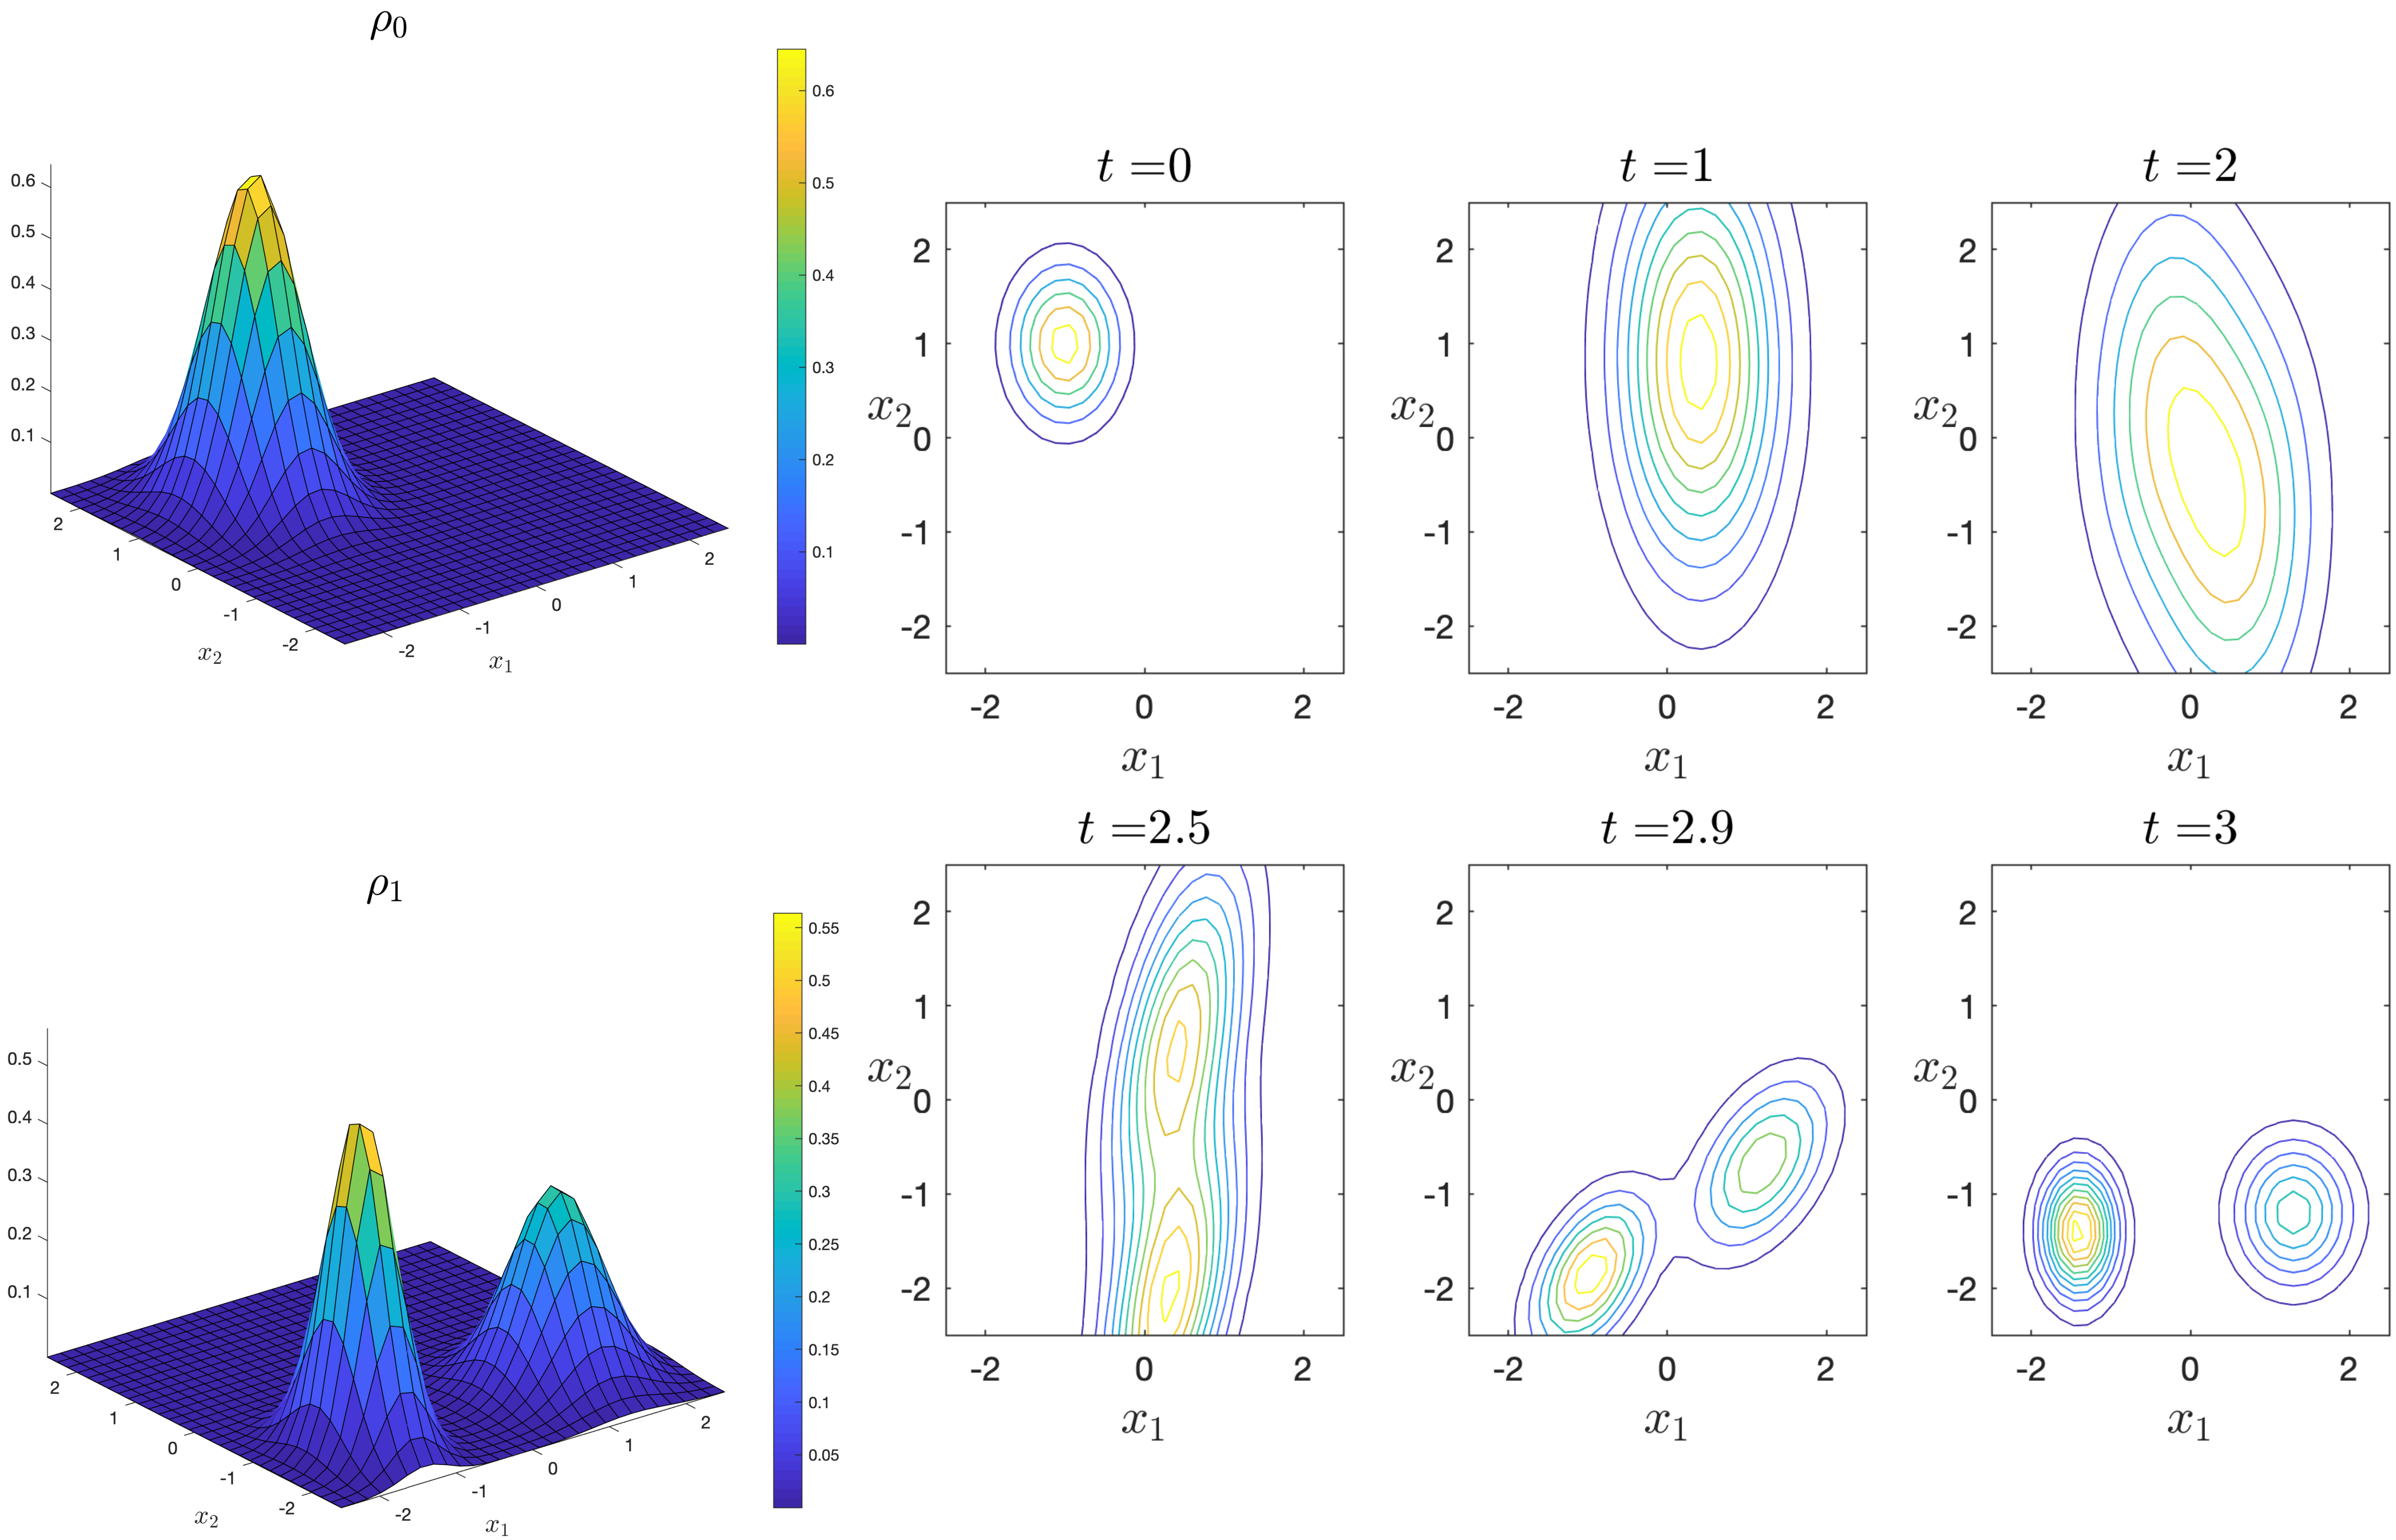

Linear prior dynamics

Feedback Density Control: Nonlinear Grad. Drift

Uncontrolled joint PDF evolution:

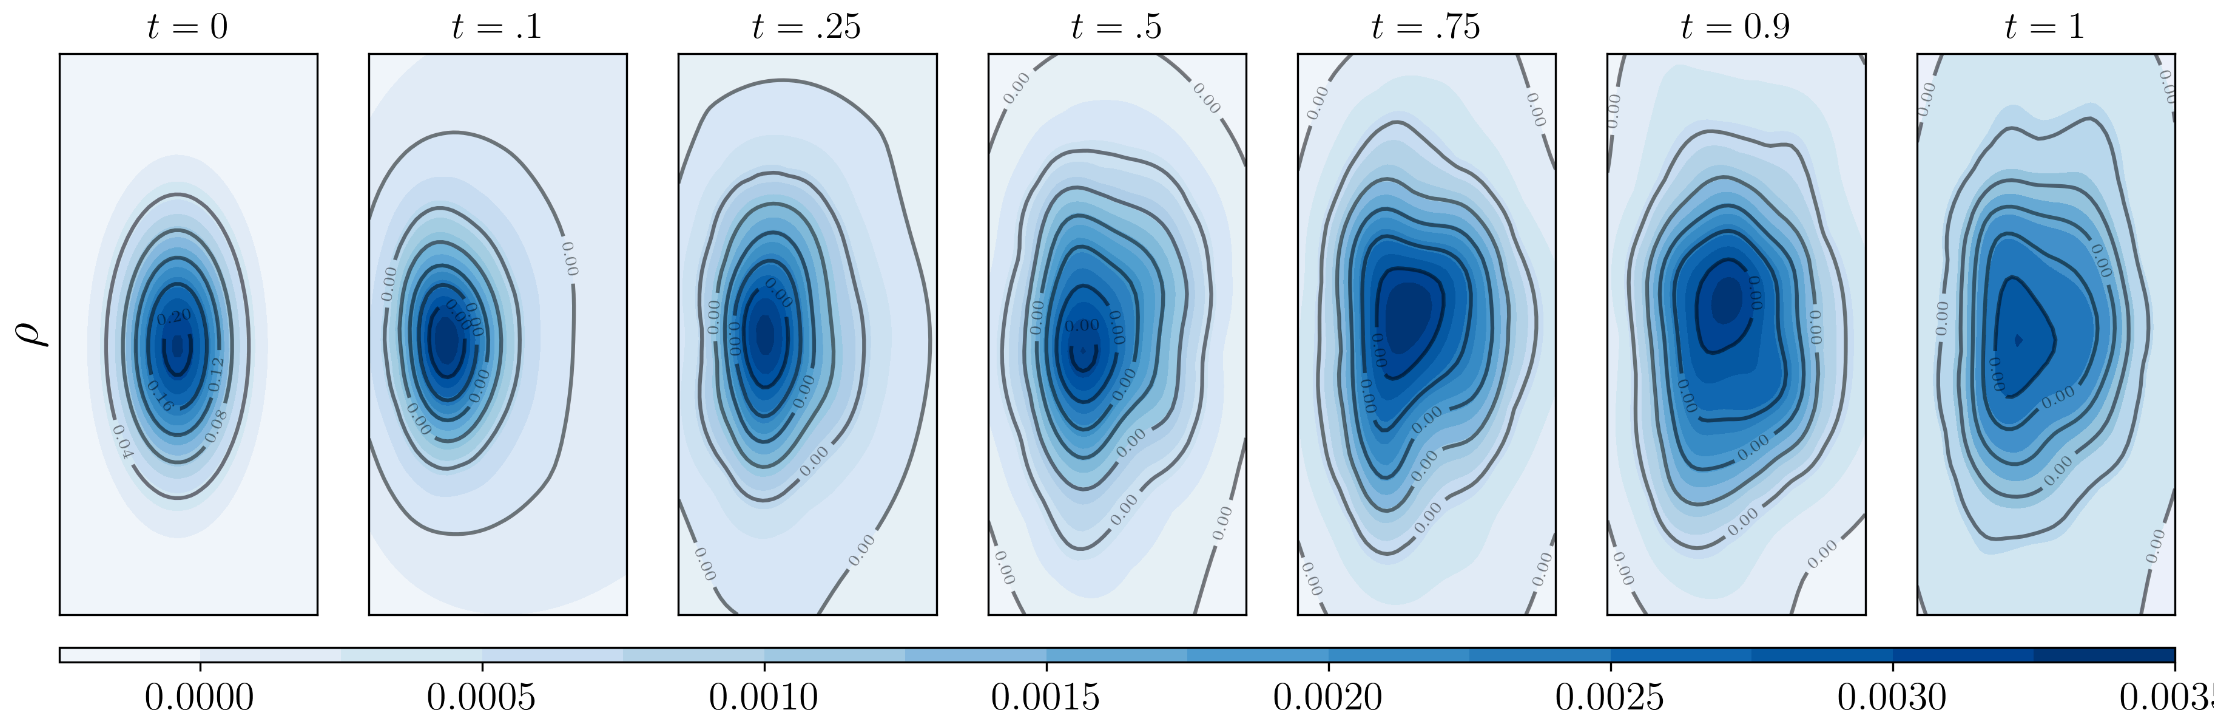

Optimal controlled joint PDF evolution:

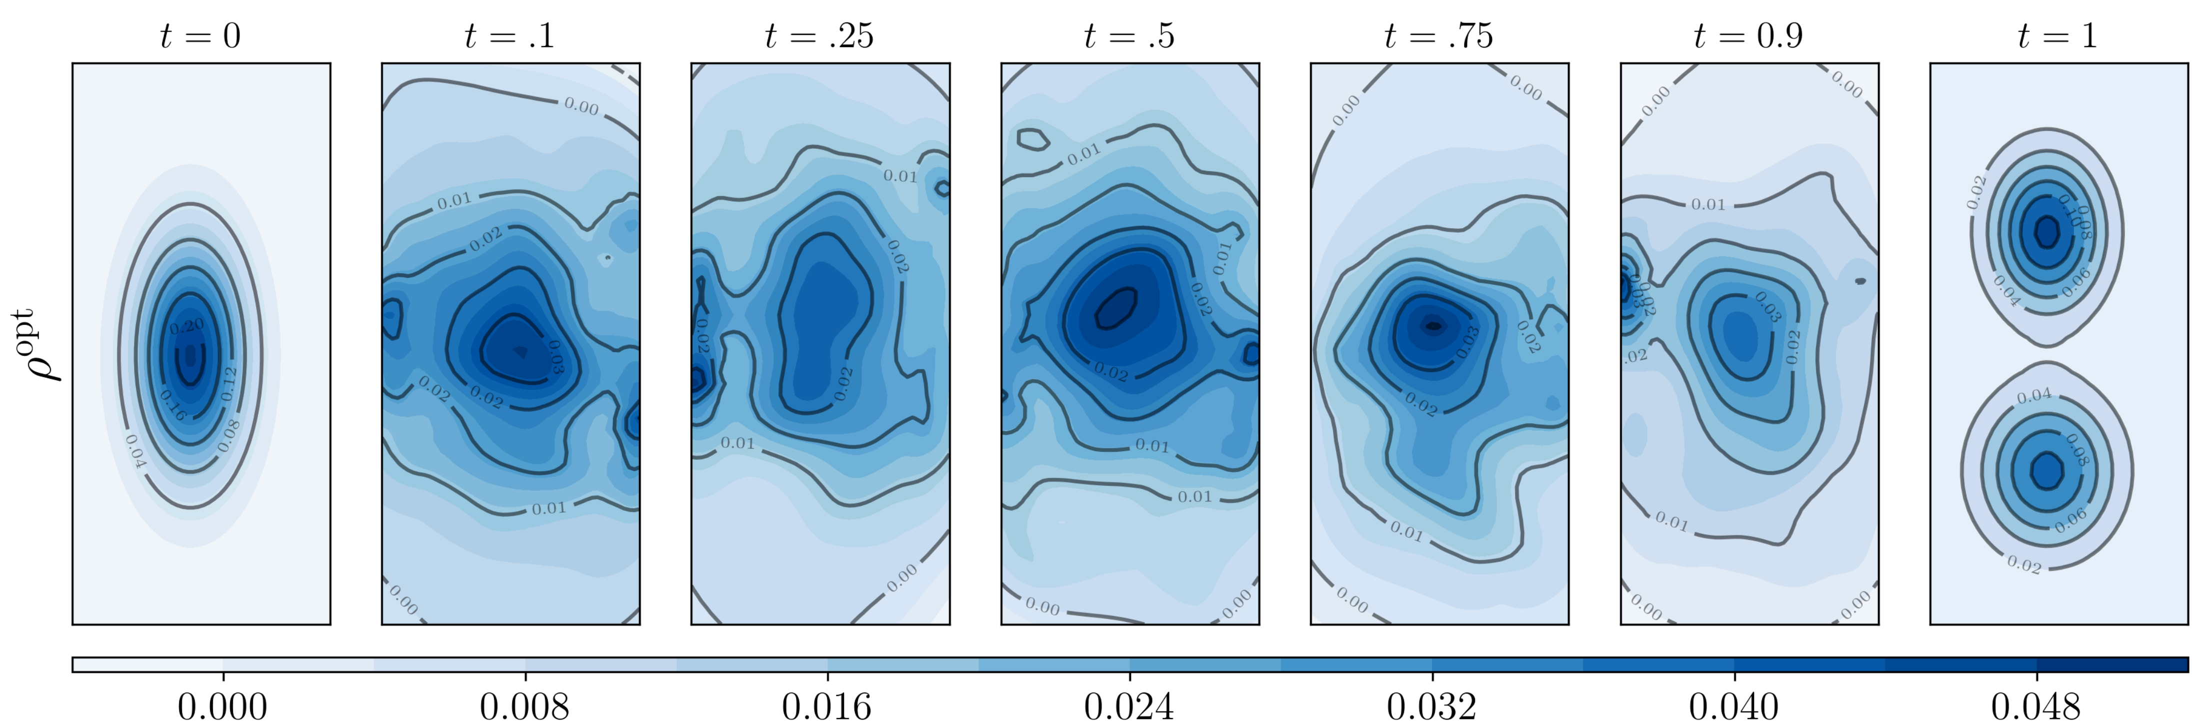

Feedback Density Control: Mixed Conservative-Dissipative Drift

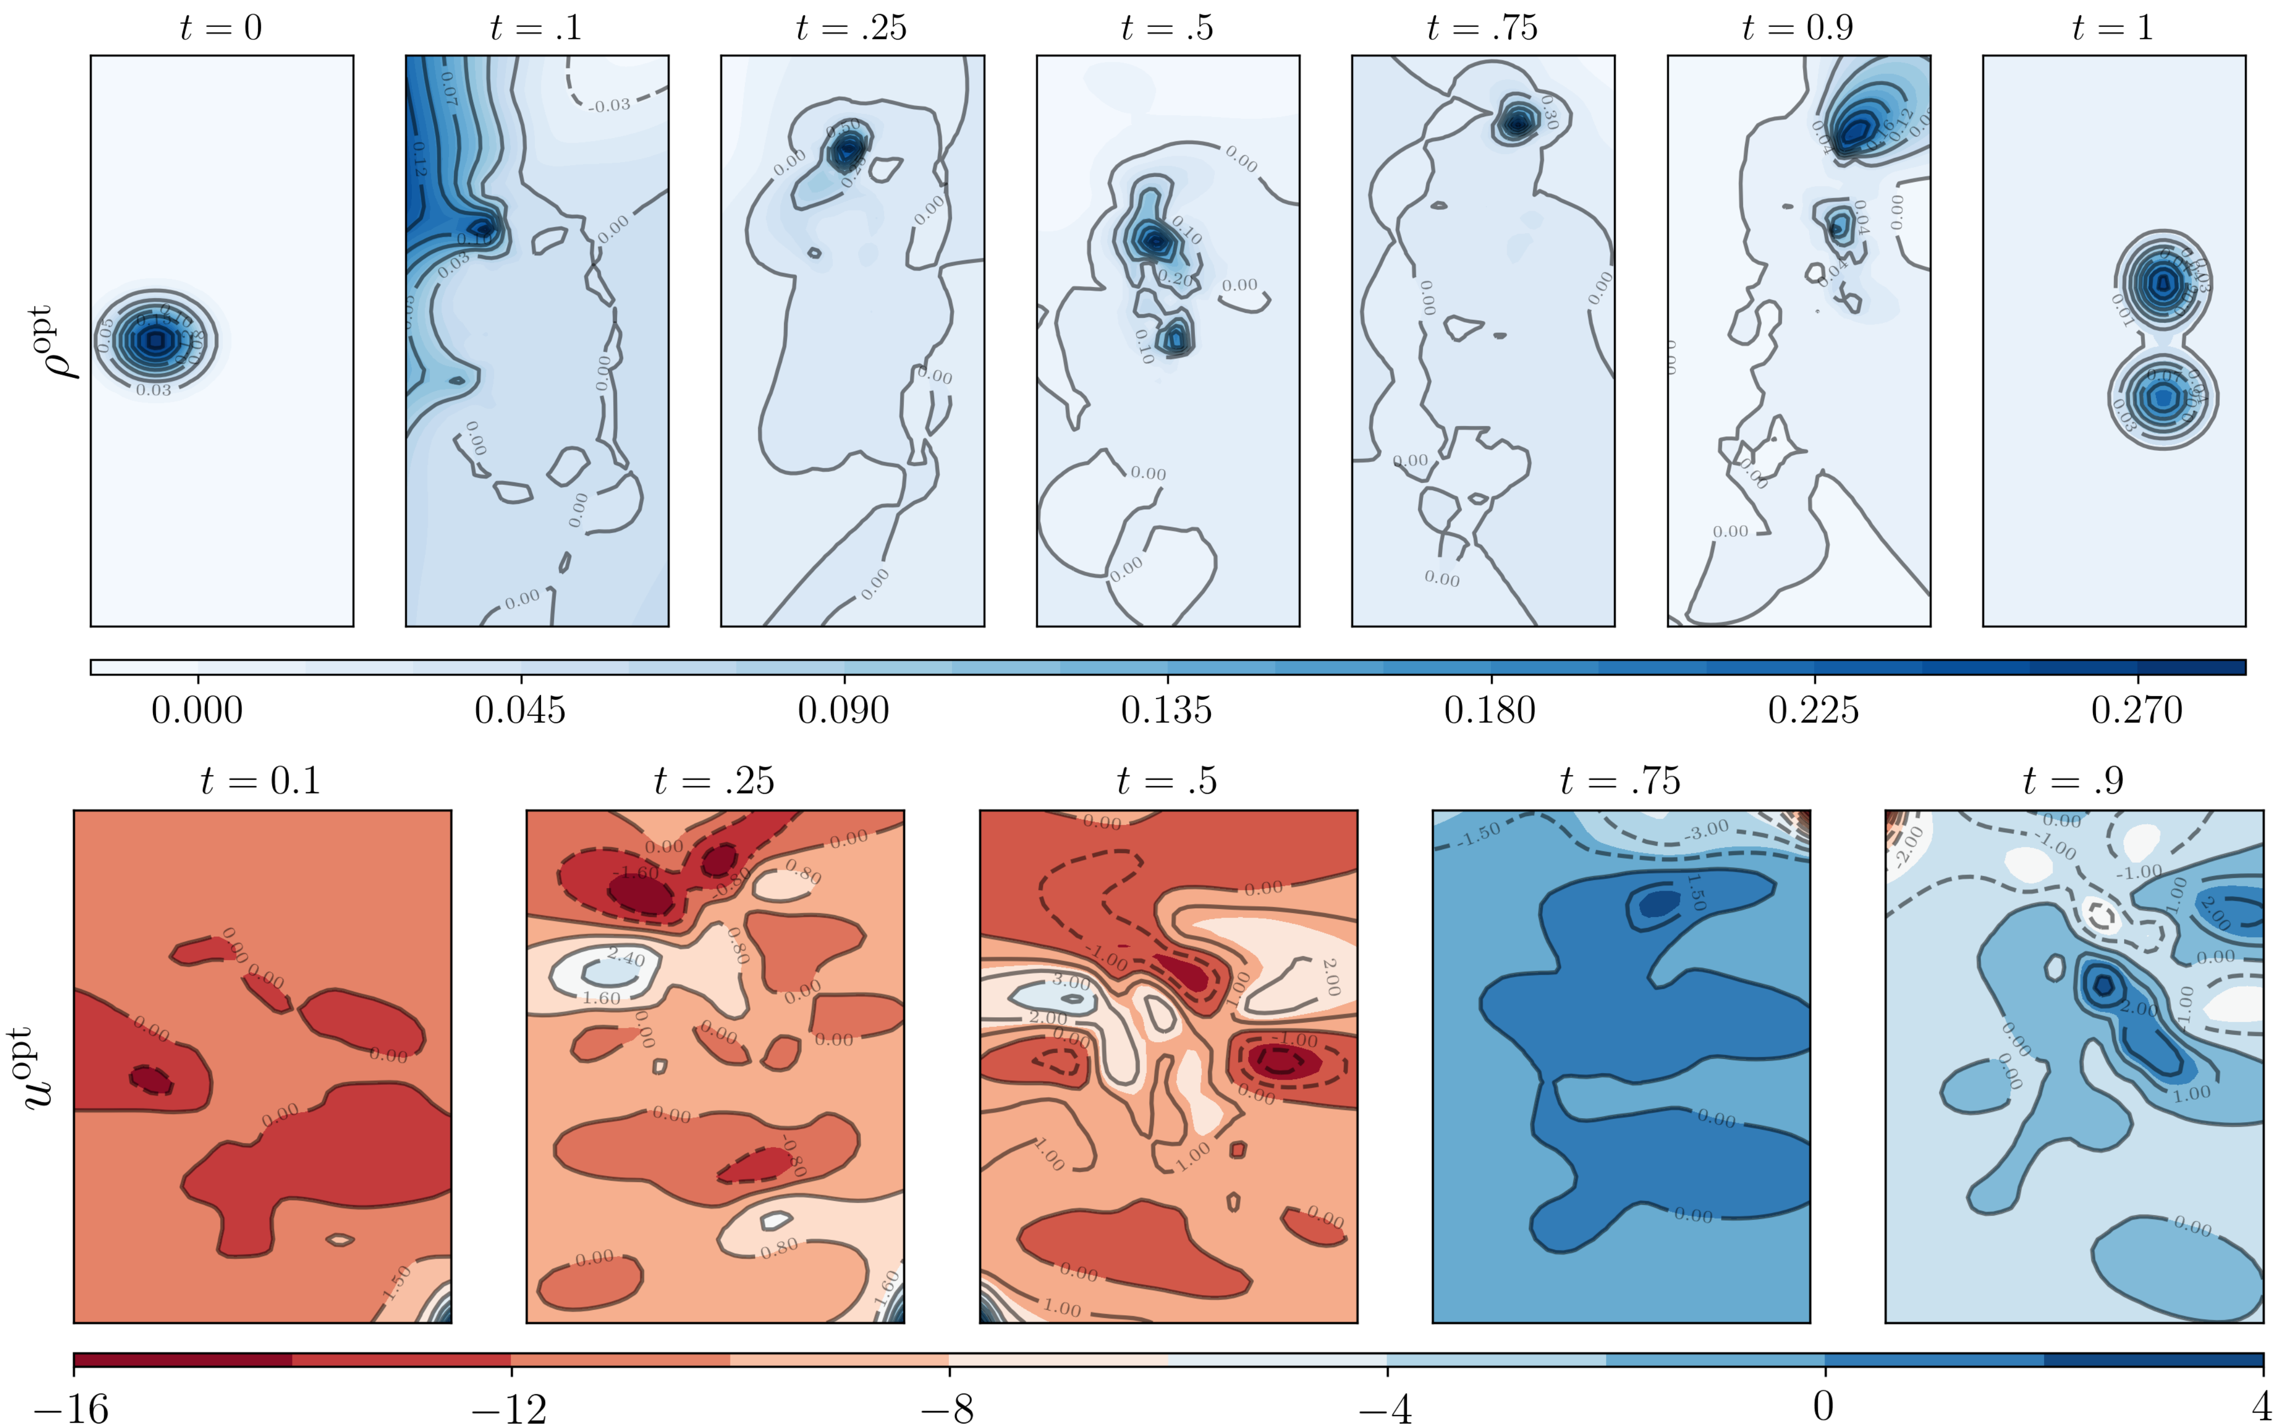

— K.F. Caluya and A.H., Wasserstein proximal algorithms for the Schrodinger bridge problem: density control with nonlinear drift, *IEEE TAC* 2021.

Density Prediction for Safe Automated Driving

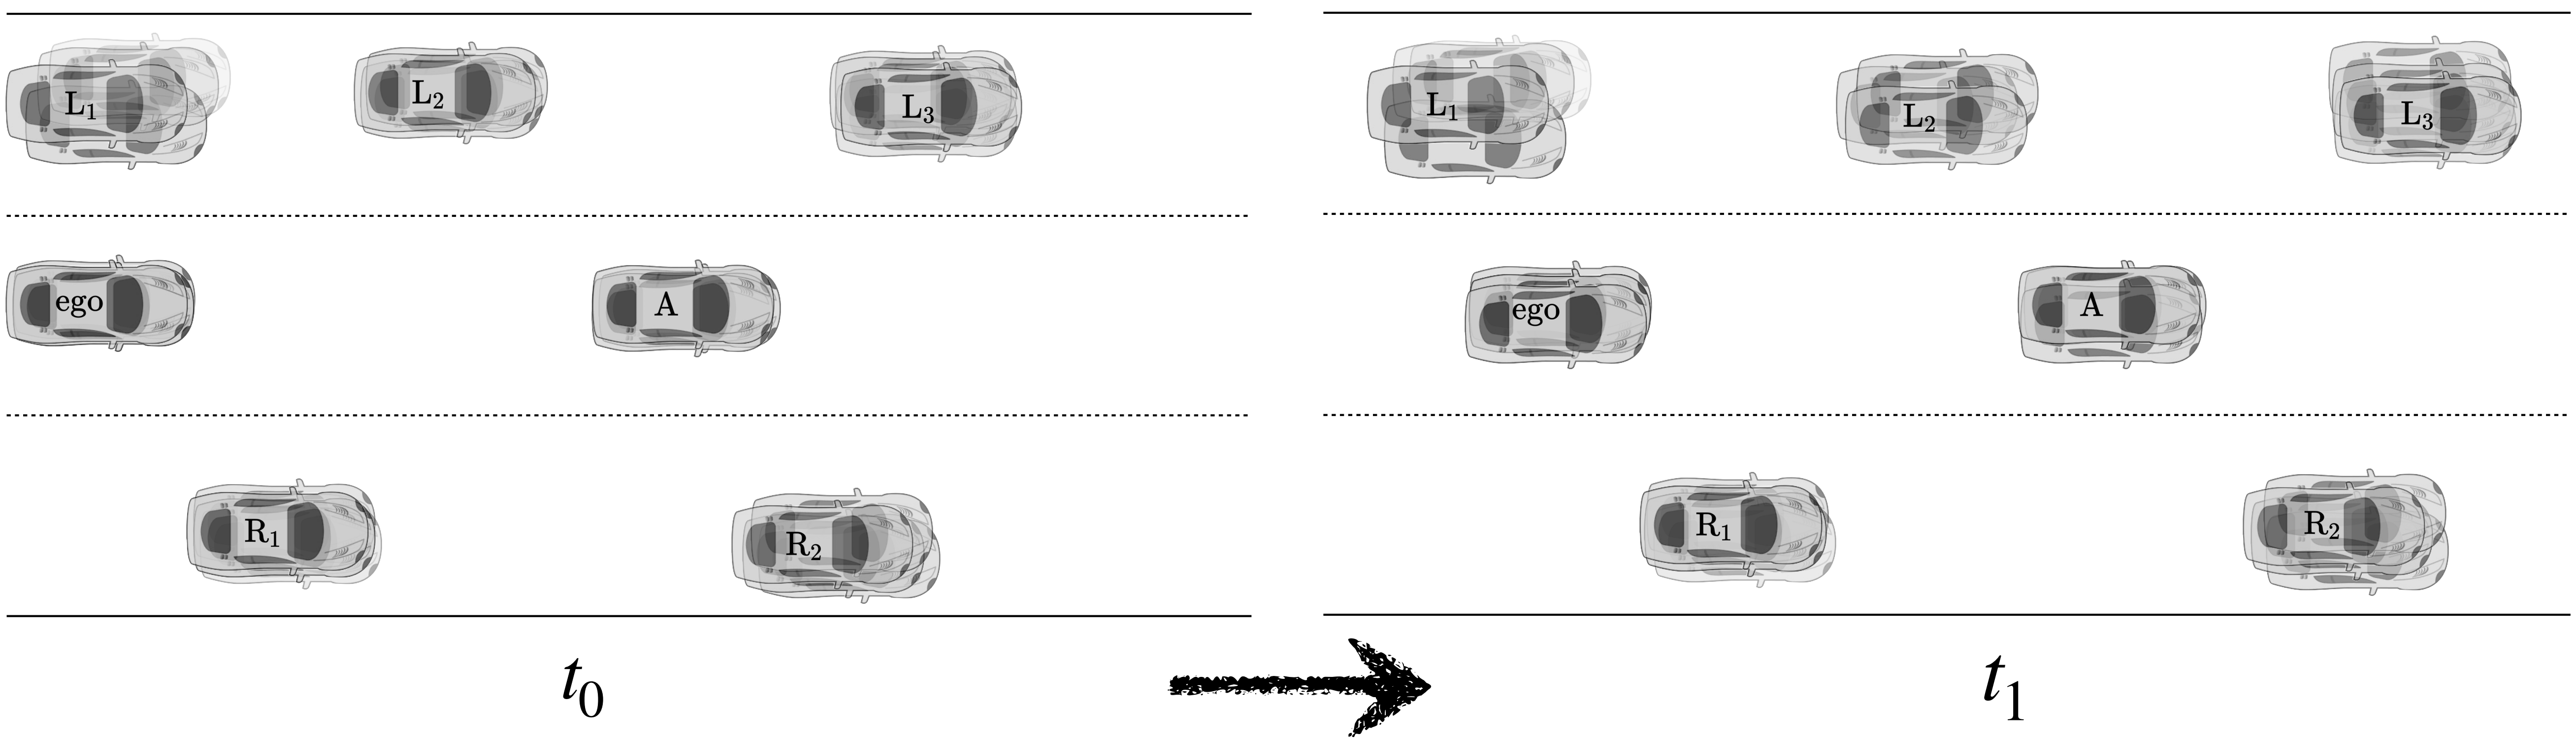

x marginals

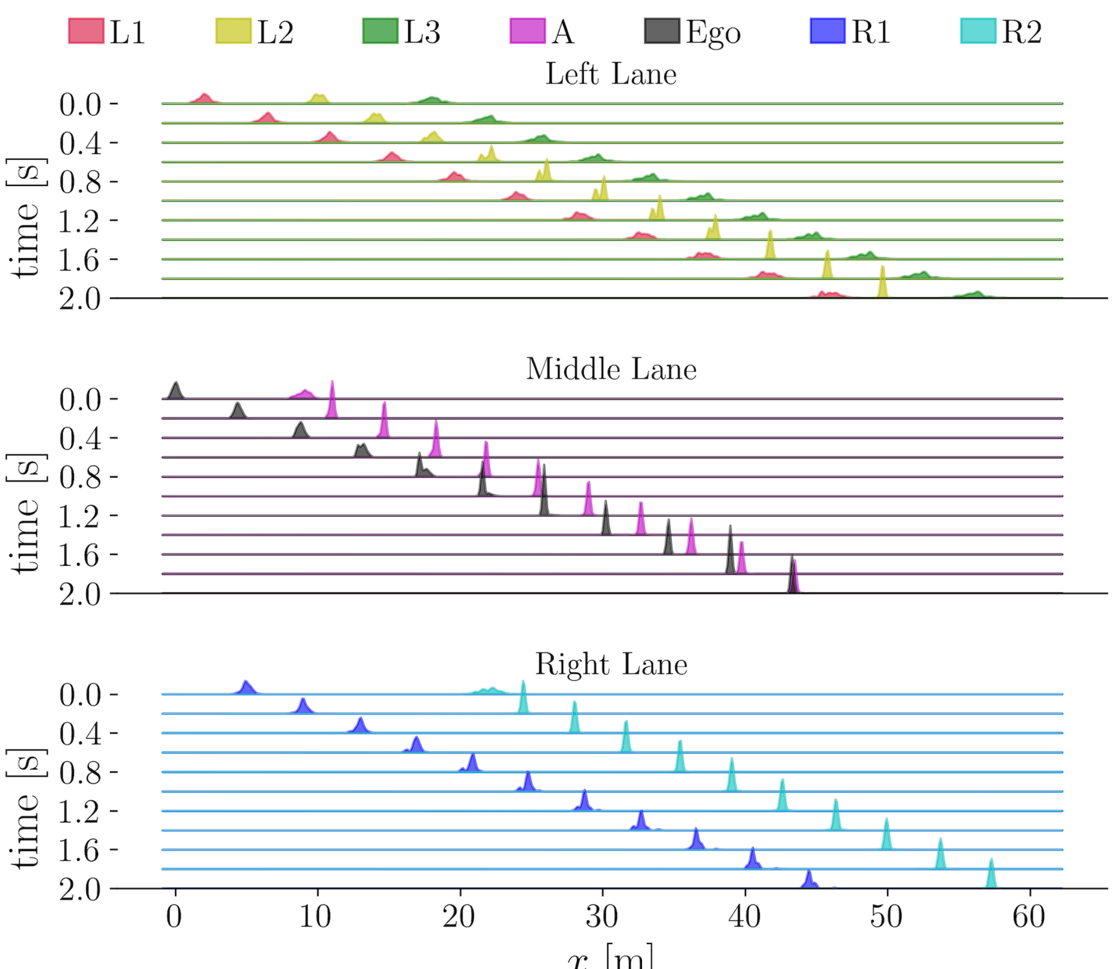

y marginals

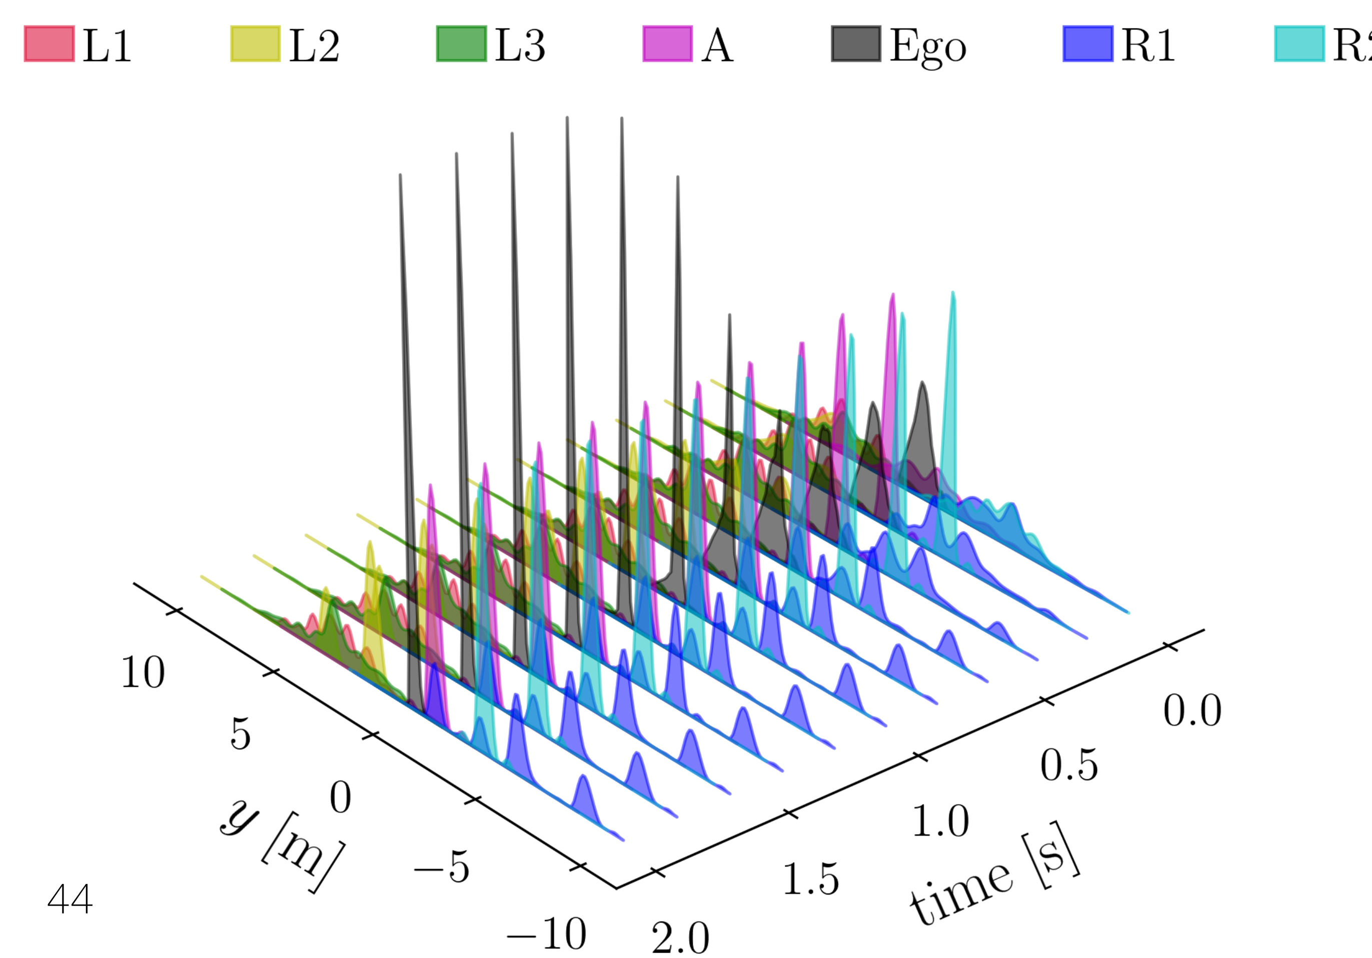

Density Control for Safe Automated Driving

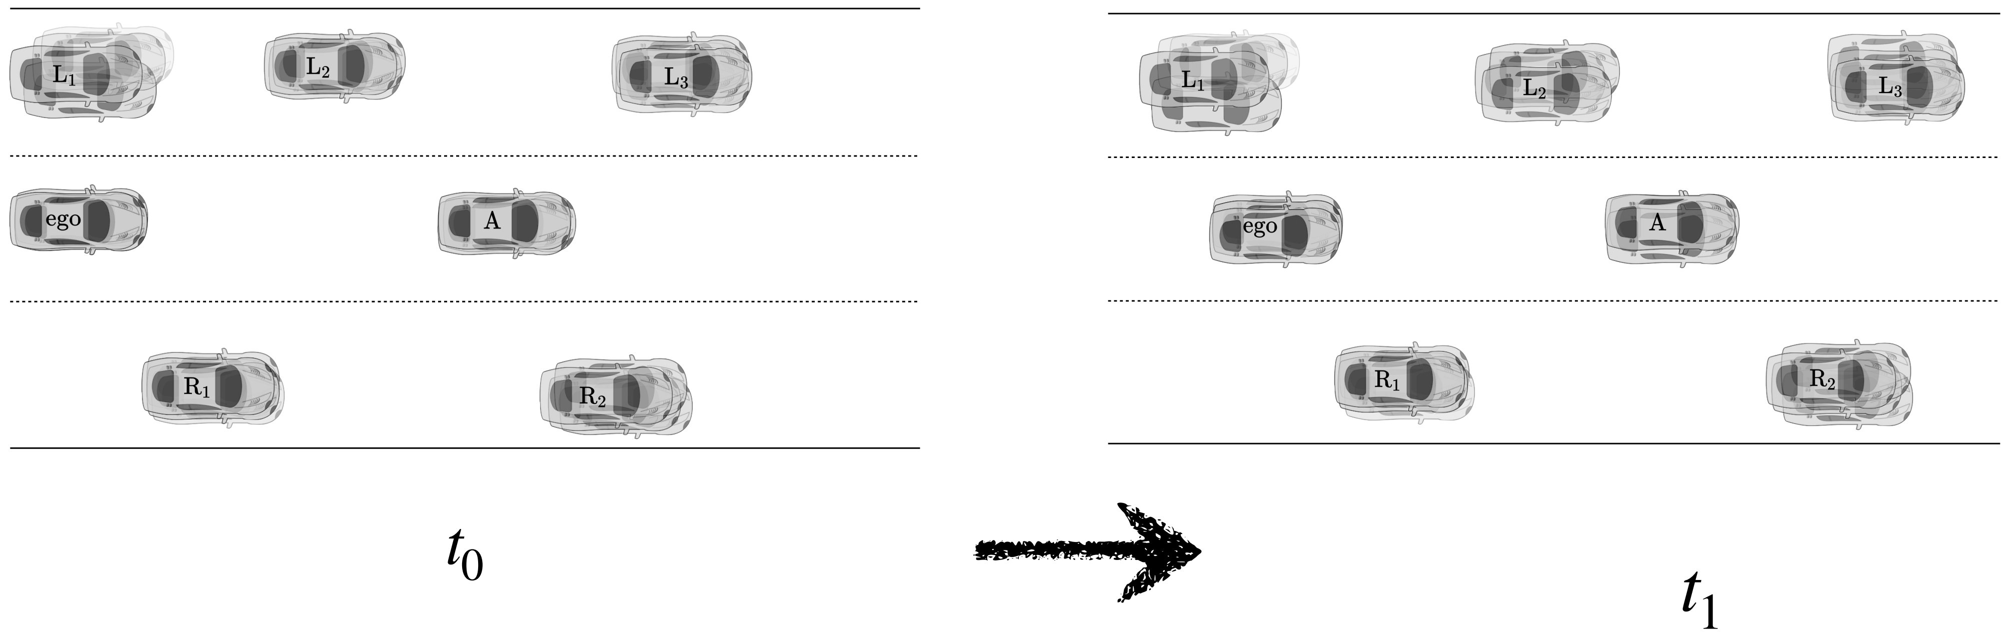

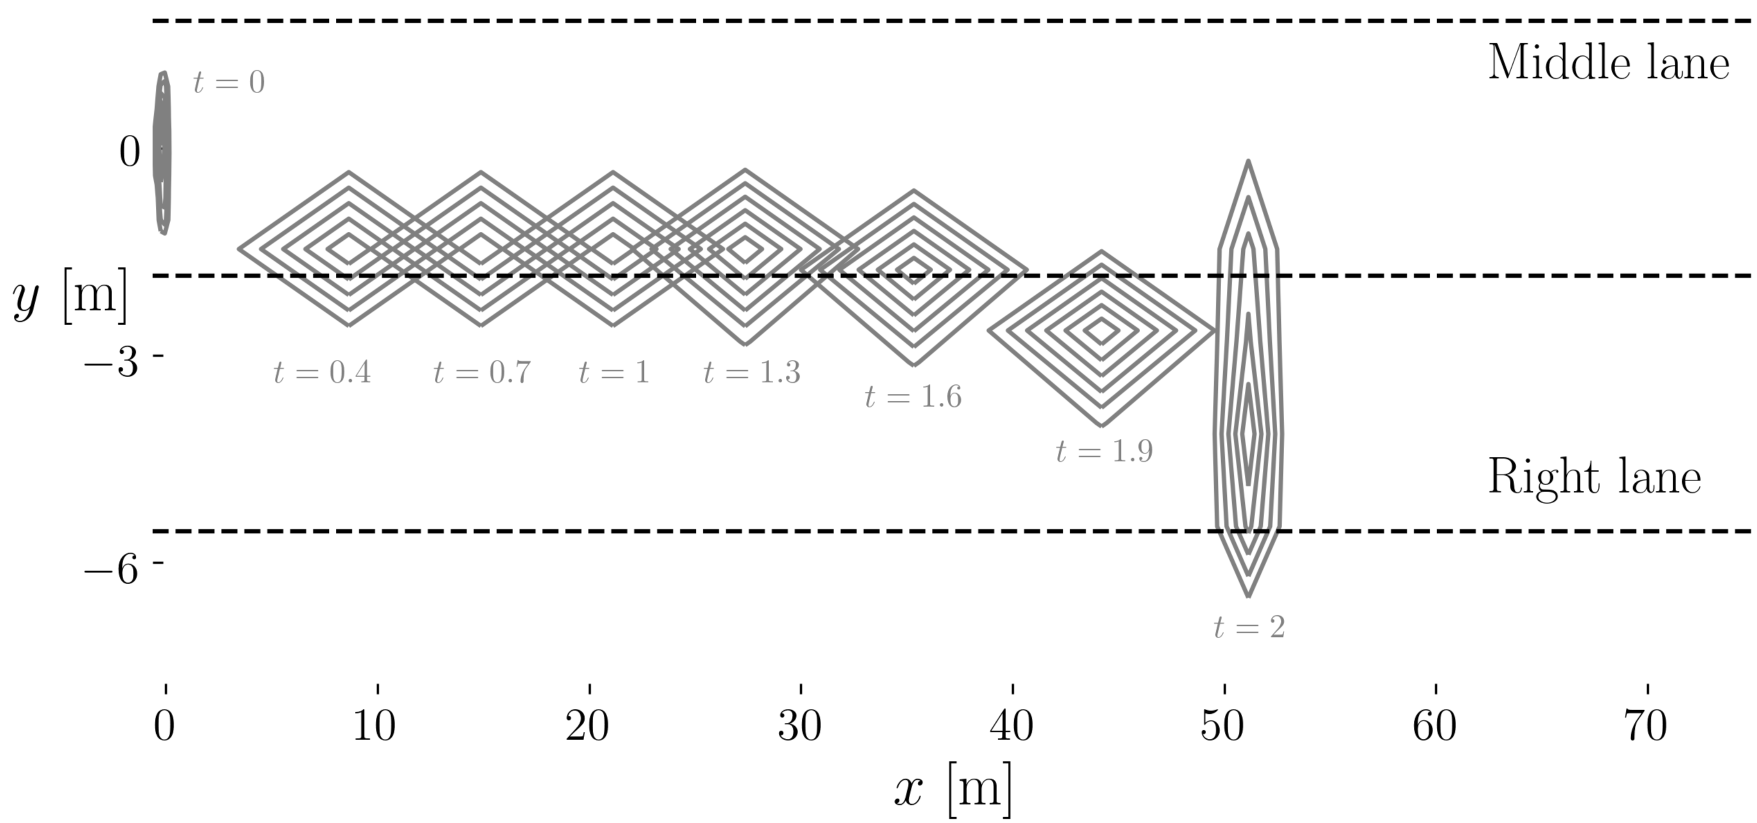

— S. Haddad, A.H., and B. Singh, Density-based stochastic reachability computation for occupancy prediction in automated driving, *IEEE TCST* 2022.

— S. Haddad, K.F. Caluya, A.H., and B. Singh, Prediction and optimal feedback steering of probability density functions for safe automated driving, *IEEE LCSS* 2021.

Learning a neural network as generalized gradient flow

Learning Neural Network from Data

$$(\text{feature vector, label}) = (\boldsymbol{x}_i, y_i) \in \mathbb{R}^d \times \mathbb{R}, \quad i = 1, \ldots, n$$

Consider shallow NN: 1 hidden layer with n_{H} neurons

NN parameter vector $\boldsymbol{\theta} := (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_{n_{\mathrm{H}}})^\top \in \mathbb{R}^{p n_{\mathrm{H}}}$

Approximating function:

$$\hat{f}(\boldsymbol{x}, \boldsymbol{\theta}) = \frac{1}{n_{\mathrm{H}}} \sum_{i=1}^{n_{\mathrm{H}}} \Phi(\boldsymbol{x}, \boldsymbol{\theta}_i), \quad \text{example: } \Phi(\boldsymbol{x}, \boldsymbol{\theta}_i) = a_i \sigma\left(\boldsymbol{w}_i^\top \boldsymbol{x} + b_i\right)$$

Population risk functional:

$$R(\hat{f}) = \mathbb{E}_{(\boldsymbol{x}, y)}\left[\left(y - \hat{f}(\boldsymbol{x}, \boldsymbol{\theta})\right)^2\right] \approx \frac{1}{n} \sum_{i=1}^{n} \left(y_i - \hat{f}(\boldsymbol{x}_i, \boldsymbol{\theta})\right)^2$$

Learning problem: $\underset{\boldsymbol{\theta} \in \mathbb{R}^{p n_{\mathrm{H}}}}{\text{minimize}} \; R(\hat{f})$

Mean Field Density Dynamics of SGD

Free energy functional: $F(\rho) := R(\hat{f}(\boldsymbol{x}, \rho))$

For quadratic loss:

$$F(\rho) = \underbrace{F_0}_{\text{independent of } \rho} + \underbrace{\int_{\mathbb{R}^p} V(\boldsymbol{\theta})\, \rho(\boldsymbol{\theta}) \mathrm{d}\boldsymbol{\theta}}_{\text{advection potential energy, linear in } \rho} + \underbrace{\int_{\mathbb{R}^p}\int_{\mathbb{R}^p} U(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}})\, \rho(\boldsymbol{\theta})\rho(\tilde{\boldsymbol{\theta}}) \mathrm{d}\boldsymbol{\theta} \mathrm{d}\tilde{\boldsymbol{\theta}}}_{\text{interaction potential energy, nonlinear in } \rho},$$

where

$$F_0 := \mathbb{E}_{(\boldsymbol{x},y)}\left[y^2\right], \qquad V(\boldsymbol{\theta}) := \mathbb{E}_{(\boldsymbol{x},y)}\left[-2y\Phi(\boldsymbol{x}, \boldsymbol{\theta})\right],$$

PDF dynamics for SGD:

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \nabla (\underbrace{V + U \circledast \rho}_{\frac{\delta F}{\delta \rho}})), \text{ where } (U \circledast \rho)(\boldsymbol{\theta}) := \int_{\mathbb{R}^p} U(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) \rho(\tilde{\boldsymbol{\theta}}) \mathrm{d}\tilde{\boldsymbol{\theta}}$$

This PDE is the gradient flow of functional F w.r.t. the Wasserstein metric W

Proximal Recursion for SGD Training of NN

$$\varrho_k(\tau,\boldsymbol{\theta}) = \underset{\varrho\in\mathcal{P}(\mathbb{R}^p)}{\arg\min}\ \frac{1}{2}\left(W\left(\varrho(\boldsymbol{\theta}),\varrho_{k-1}(\tau,\boldsymbol{\theta})\right)\right)^2 + \tau\, F(\varrho(\boldsymbol{\theta}))$$

$$= \text{prox}^{W}_{\tau F}\left(\varrho_{k-1}\right)$$

Case study: Wisconsin Breast Cancer (Diagnostic) Data Set

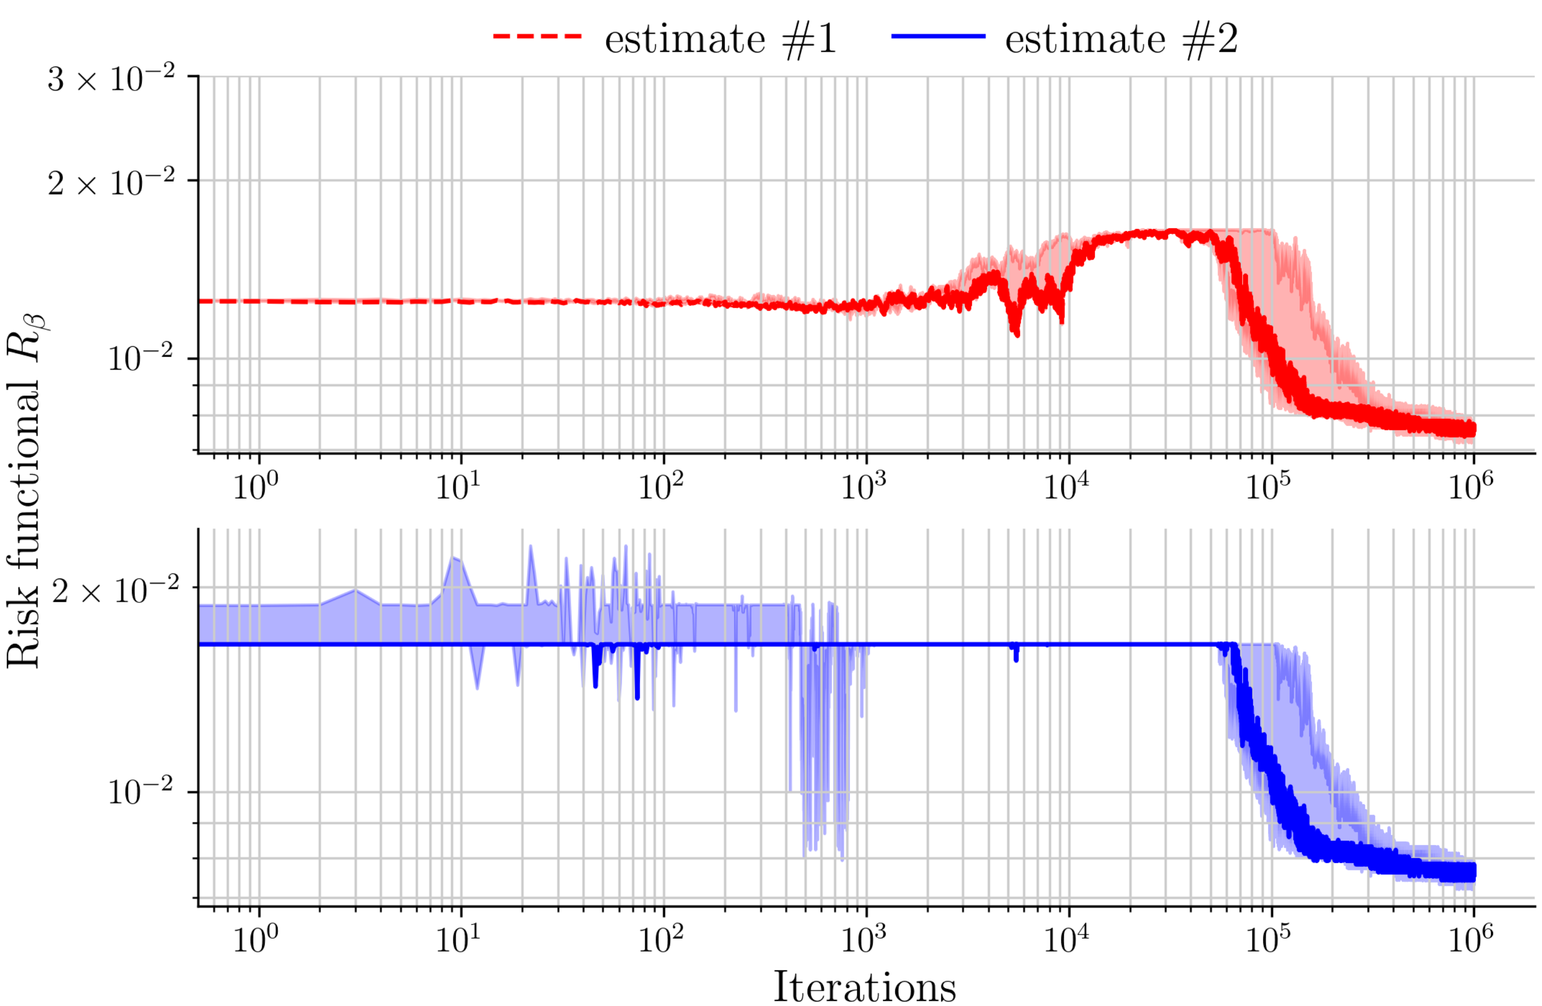

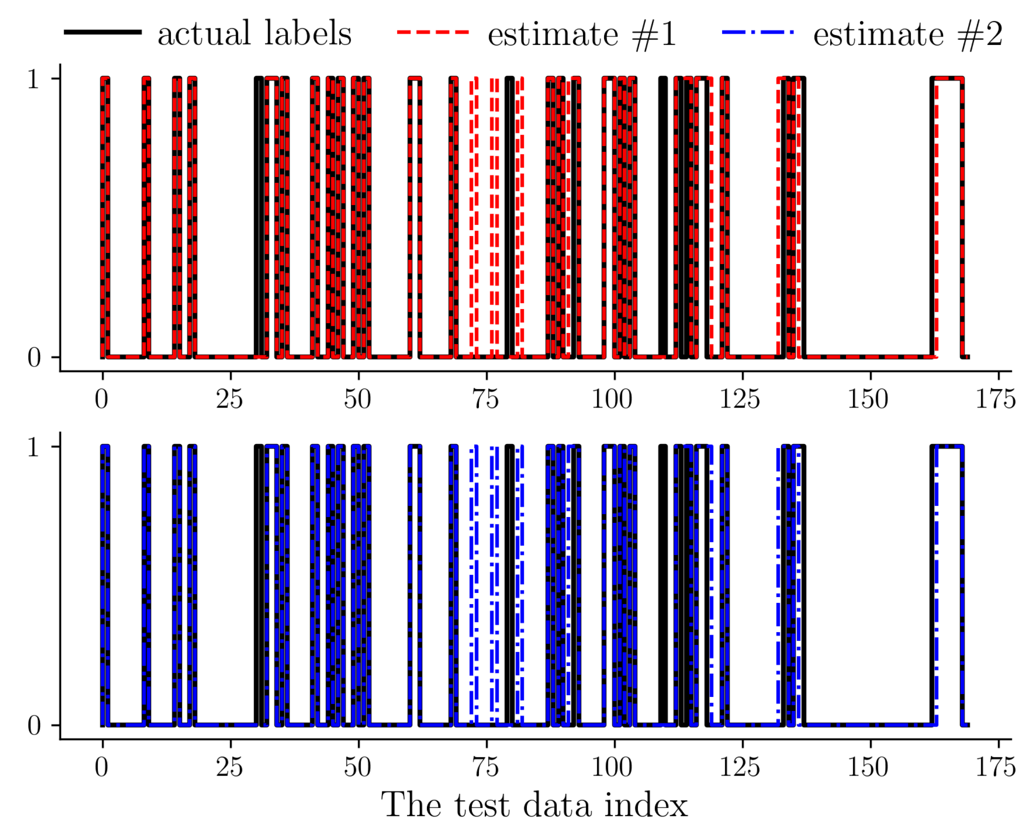

Classification accuracy for the WBDC dataset		
β	Estimate #1	Estimate #2
0.03	91.17%	92.35%
0.05	92.94%	92.94%
0.07	78.23%	92.94%

Mean Field Density Dynamics of Stoc. Heavy Ball

$$\partial_t \mu_t = -\nabla. \left[\mu_t . \begin{pmatrix} r \\ -\nabla F'([\mu_t]^\theta) - \gamma r \end{pmatrix}\right] + \gamma\beta^{-1}\Delta_r \mu_t$$

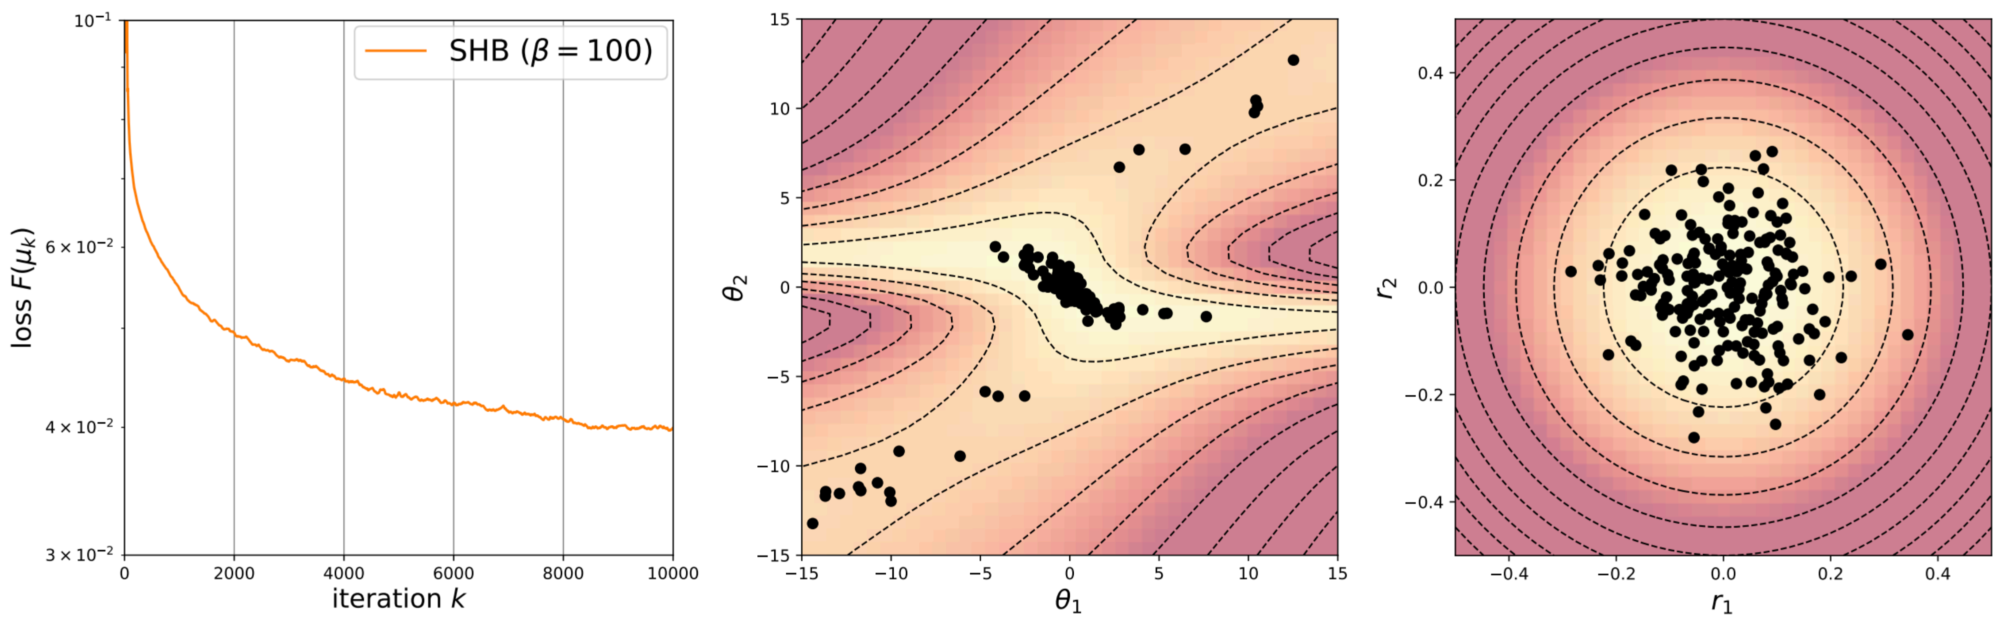

(a) Loss $F(\mu_k^n)$ (b) θ marginal (c) r marginal

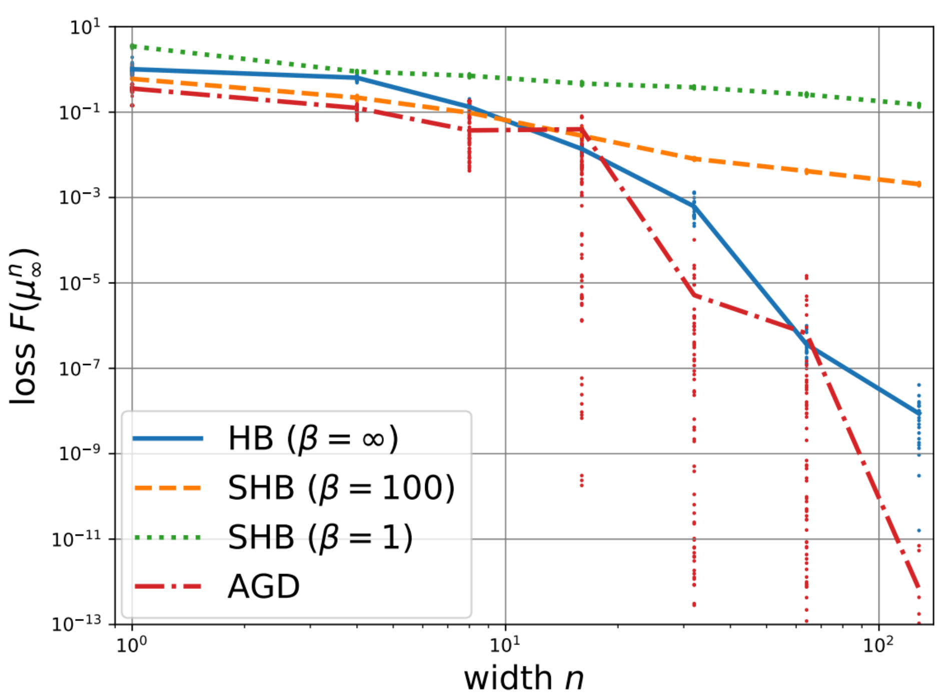

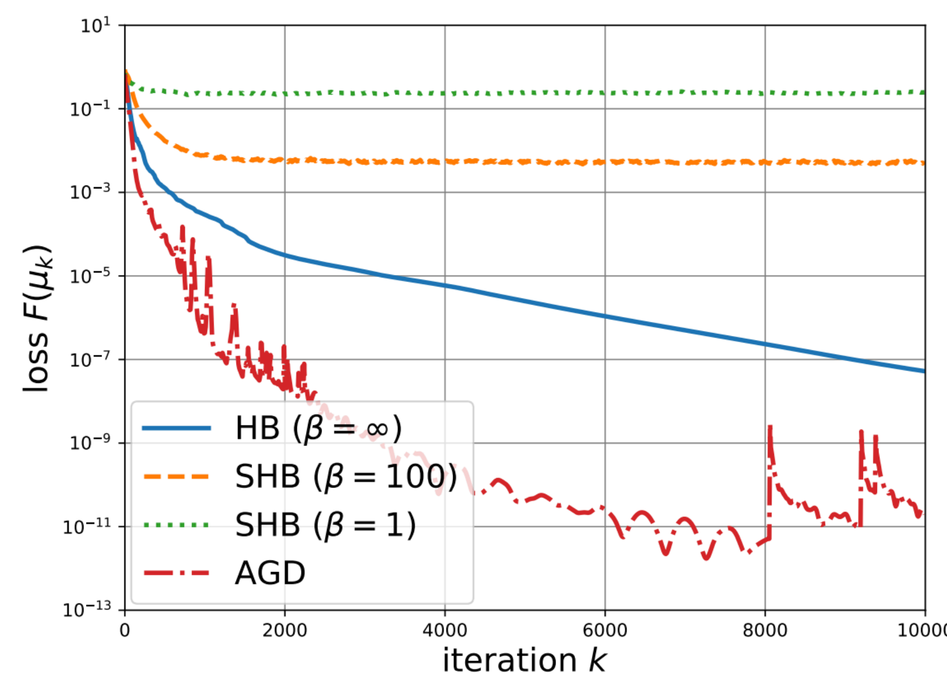

Summary and Applications

PDF prediction in Mars EDL

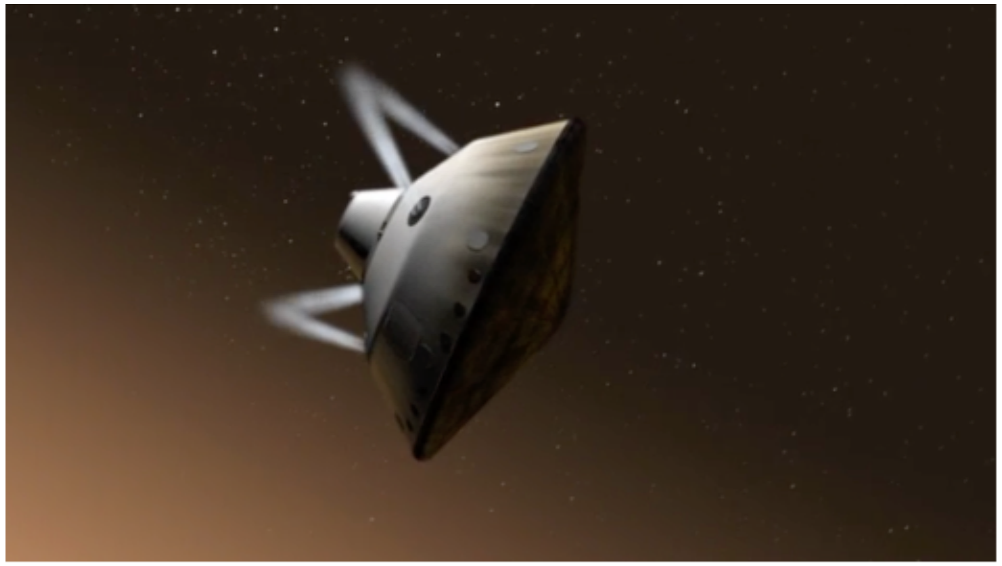

Predict heating rate uncertainty

Filtering in Mars EDL

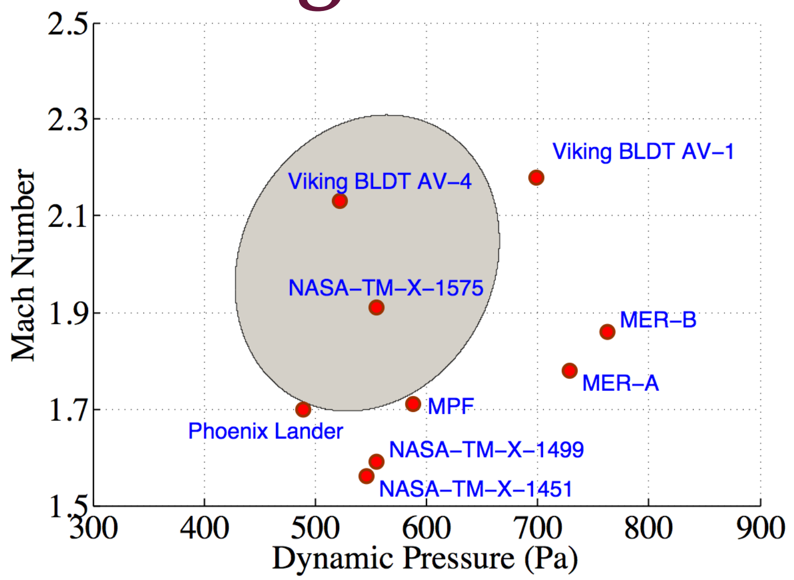

Estimate state to deploy parachute

PDF control in Mars EDL

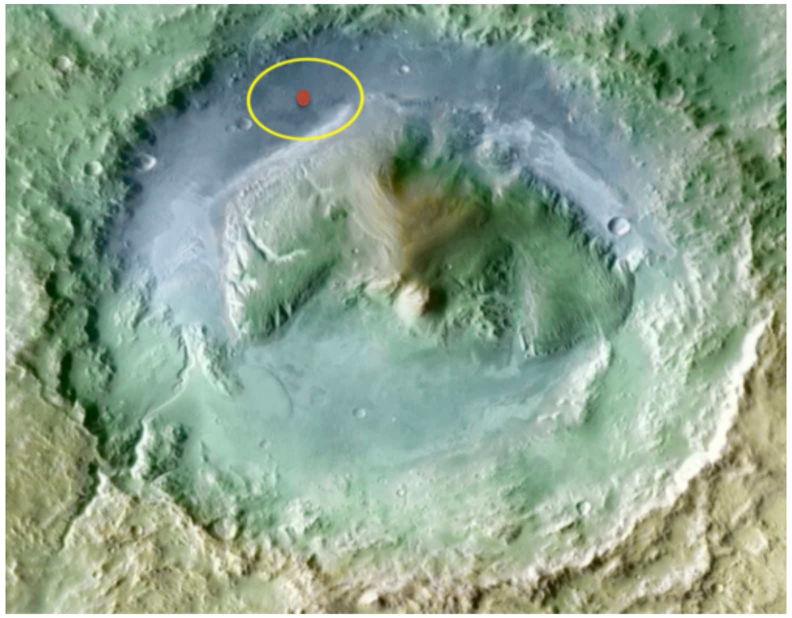

Gale Crater (4.49S, 137.42E)

Steer state PDF to achieve desired landing footprint accuracy

Learning in Mars EDL

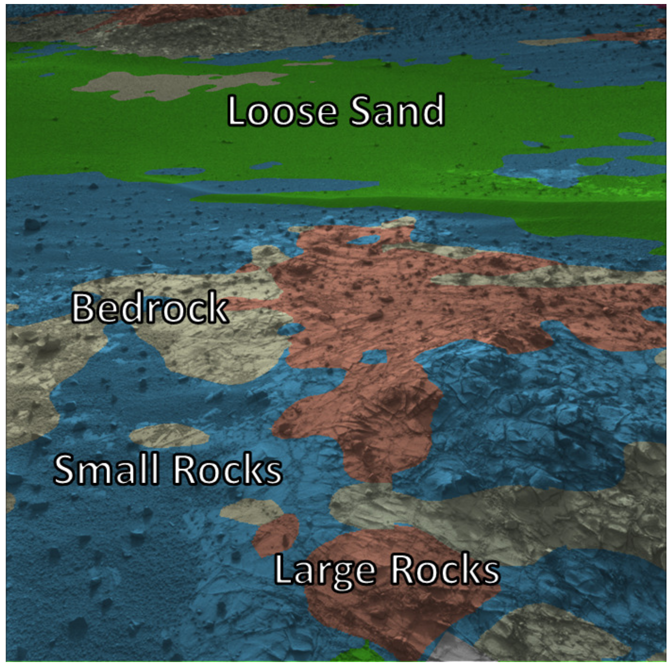

Learn surface feature from data

Thank You

Support:

Back up slides

Density Control with Hard State Constraints

SDE sample paths reflecting from a given boundary

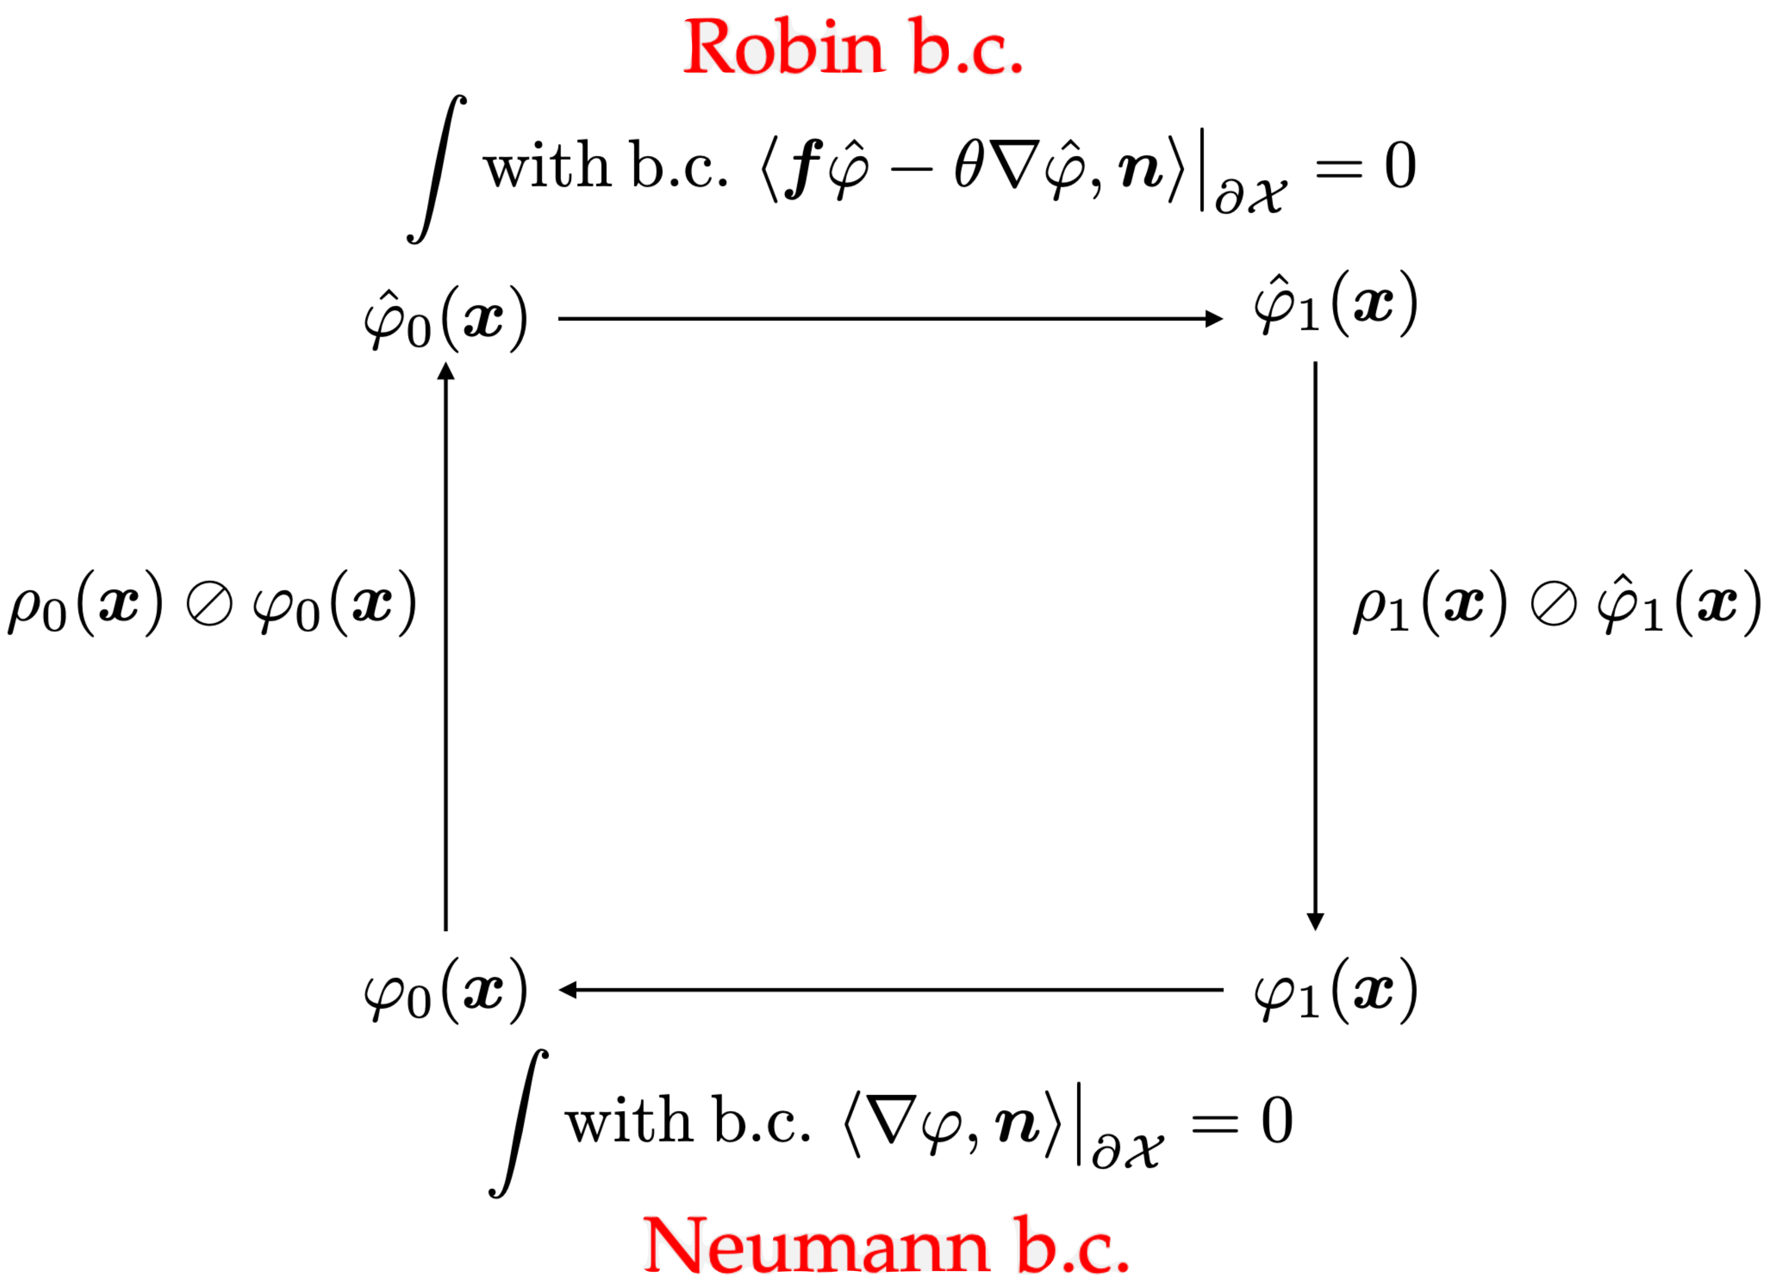

— K.F. Caluya and A.H., Reflecting Schrodinger bridge: density control with path constraints, *ACC* 2021.

Density Control with Hard State Constraints

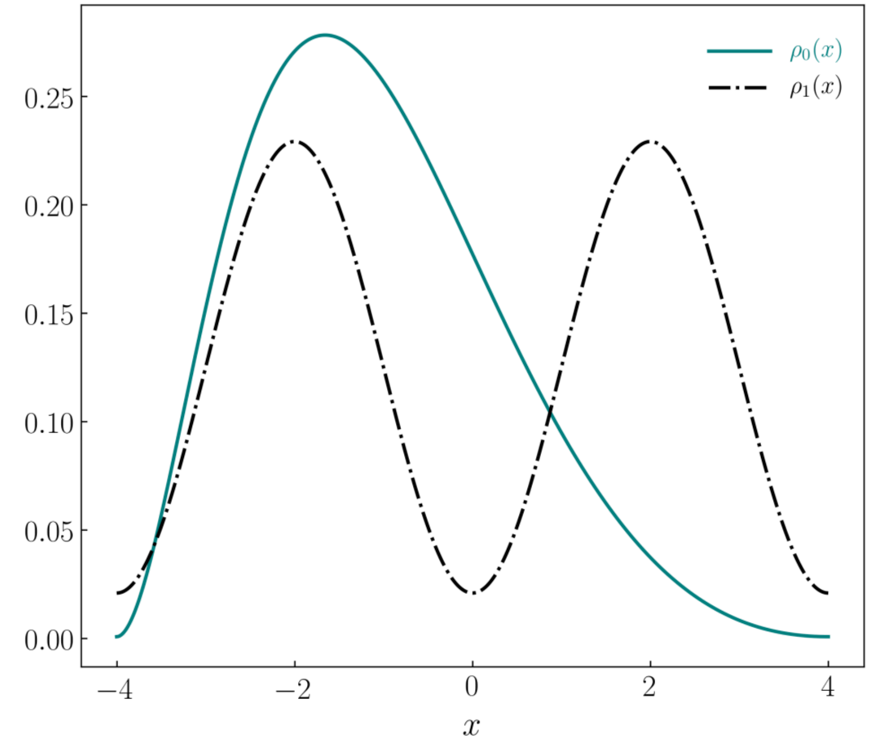

(a) Endpoint PDFs ρ_0, ρ_1

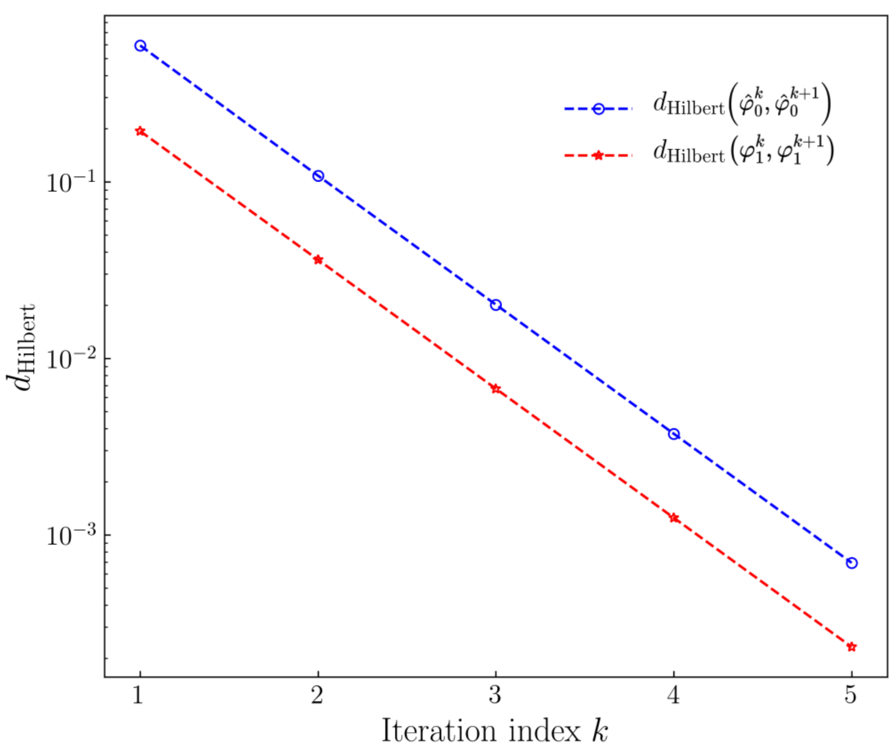

(b) Convergence of the fixed point recursion w.r.t. the Hilbert metric

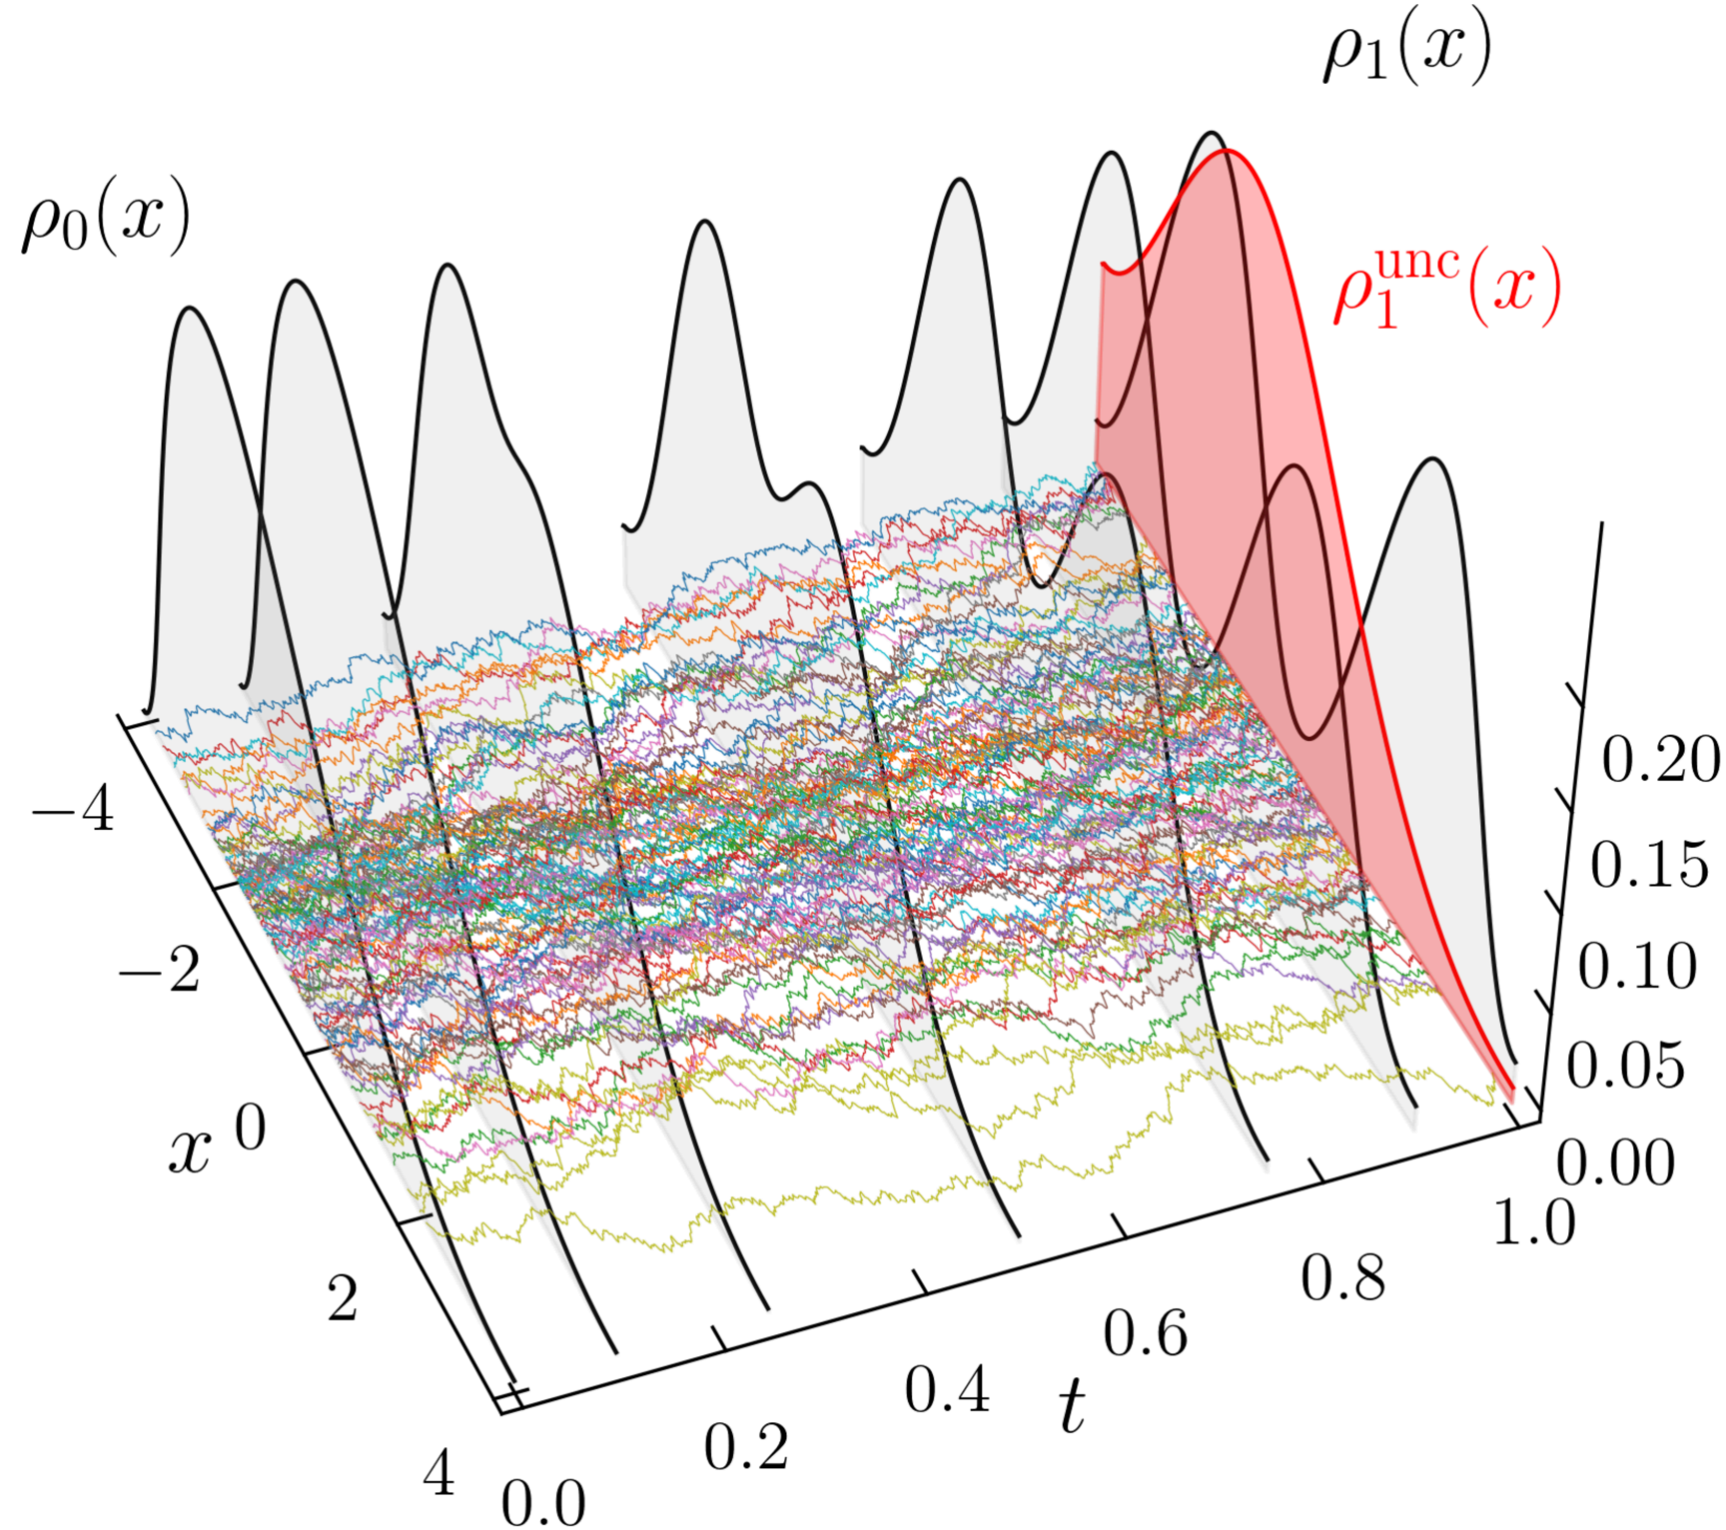

— K.F. Caluya and A.H., Reflecting Schrodinger bridge: density control with path constraints, *ACC* 2021.