Toward Feedback Control of Densities in Nonlinear Systems

Abhishek Halder

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

Joint work with Kenneth F. Caluya (UC Santa Cruz)

CDC 2019, Workshop: Uncertainty Synthesis, Nice

Overarching Theme

Systems-control theory for densities

What is density?

Probability Density Fn.

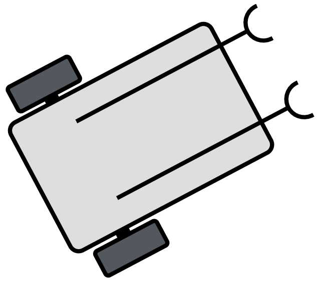

$$\boldsymbol{x}(t) \in \begin{pmatrix} x \\ y \\ \theta \end{pmatrix} \in \mathcal{X} \equiv \mathbb{R}^2 \times \mathbb{S}^1$$

Probability Density Fn.

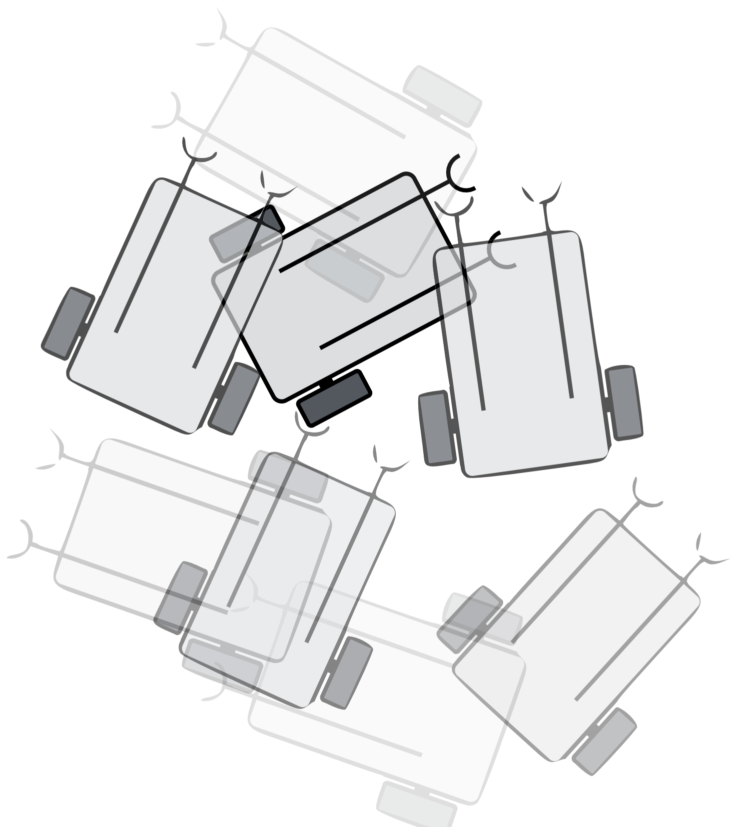

$$\boldsymbol{x}(t) \in \begin{pmatrix} x \\ y \\ \theta \end{pmatrix} \in \mathcal{X} \equiv \mathbb{R}^2 \times \mathbb{S}^1$$

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

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

Probability Density Fn.

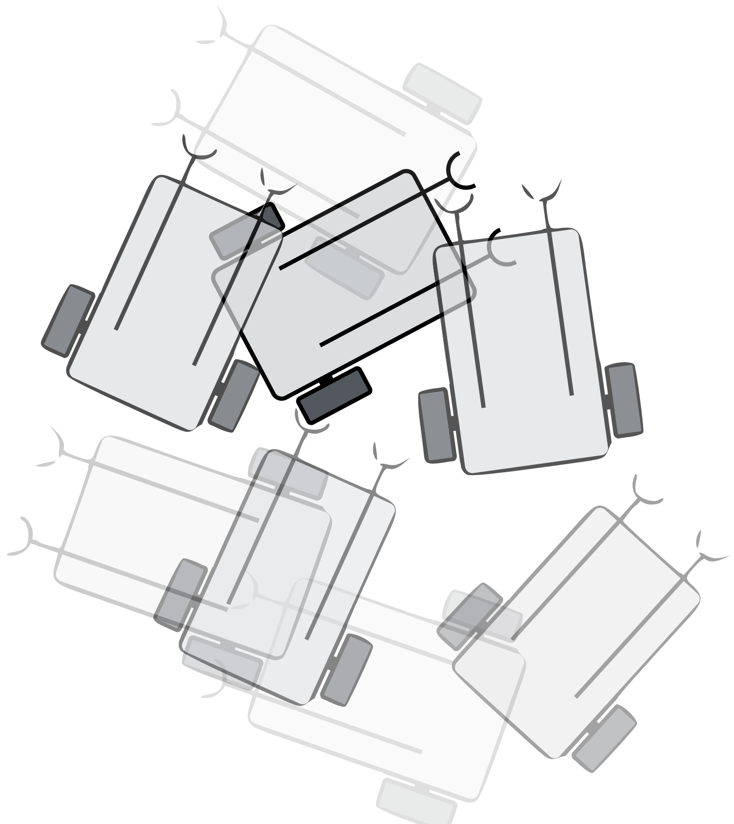

Population Density Fn.

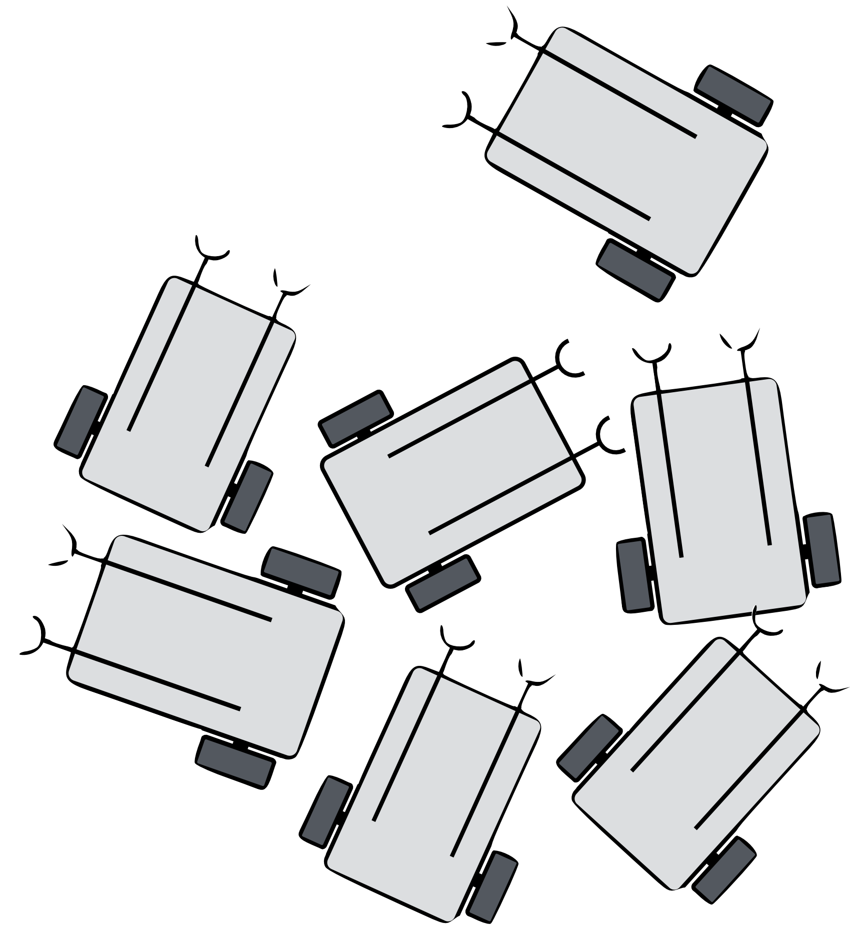

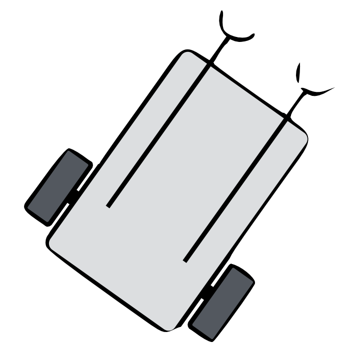

$$\boldsymbol{x}(t) \in \begin{pmatrix} x \\ y \\ \theta \end{pmatrix} \in \mathcal{X} \equiv \mathbb{R}^2 \times \mathbb{S}^1$$

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

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

Why care about densities?

Prediction Problem

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

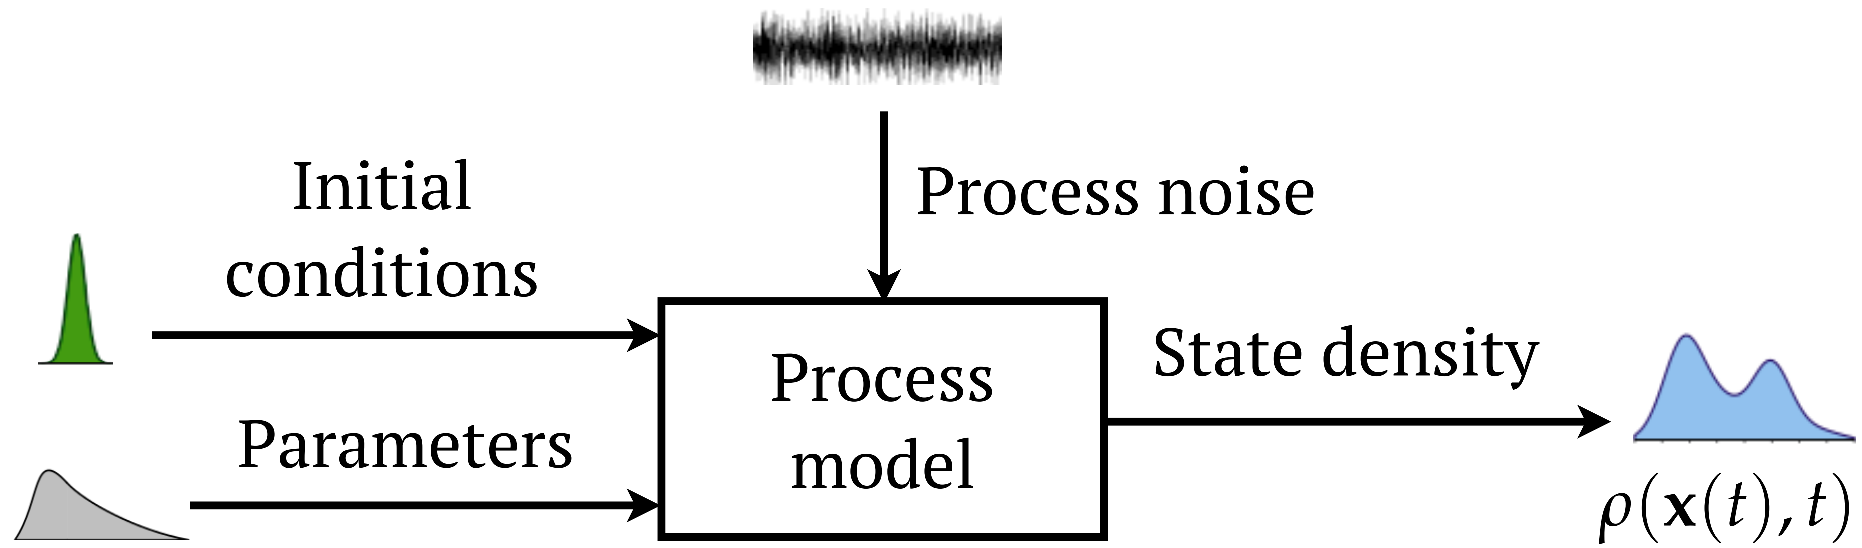

Trajectory flow:

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

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

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

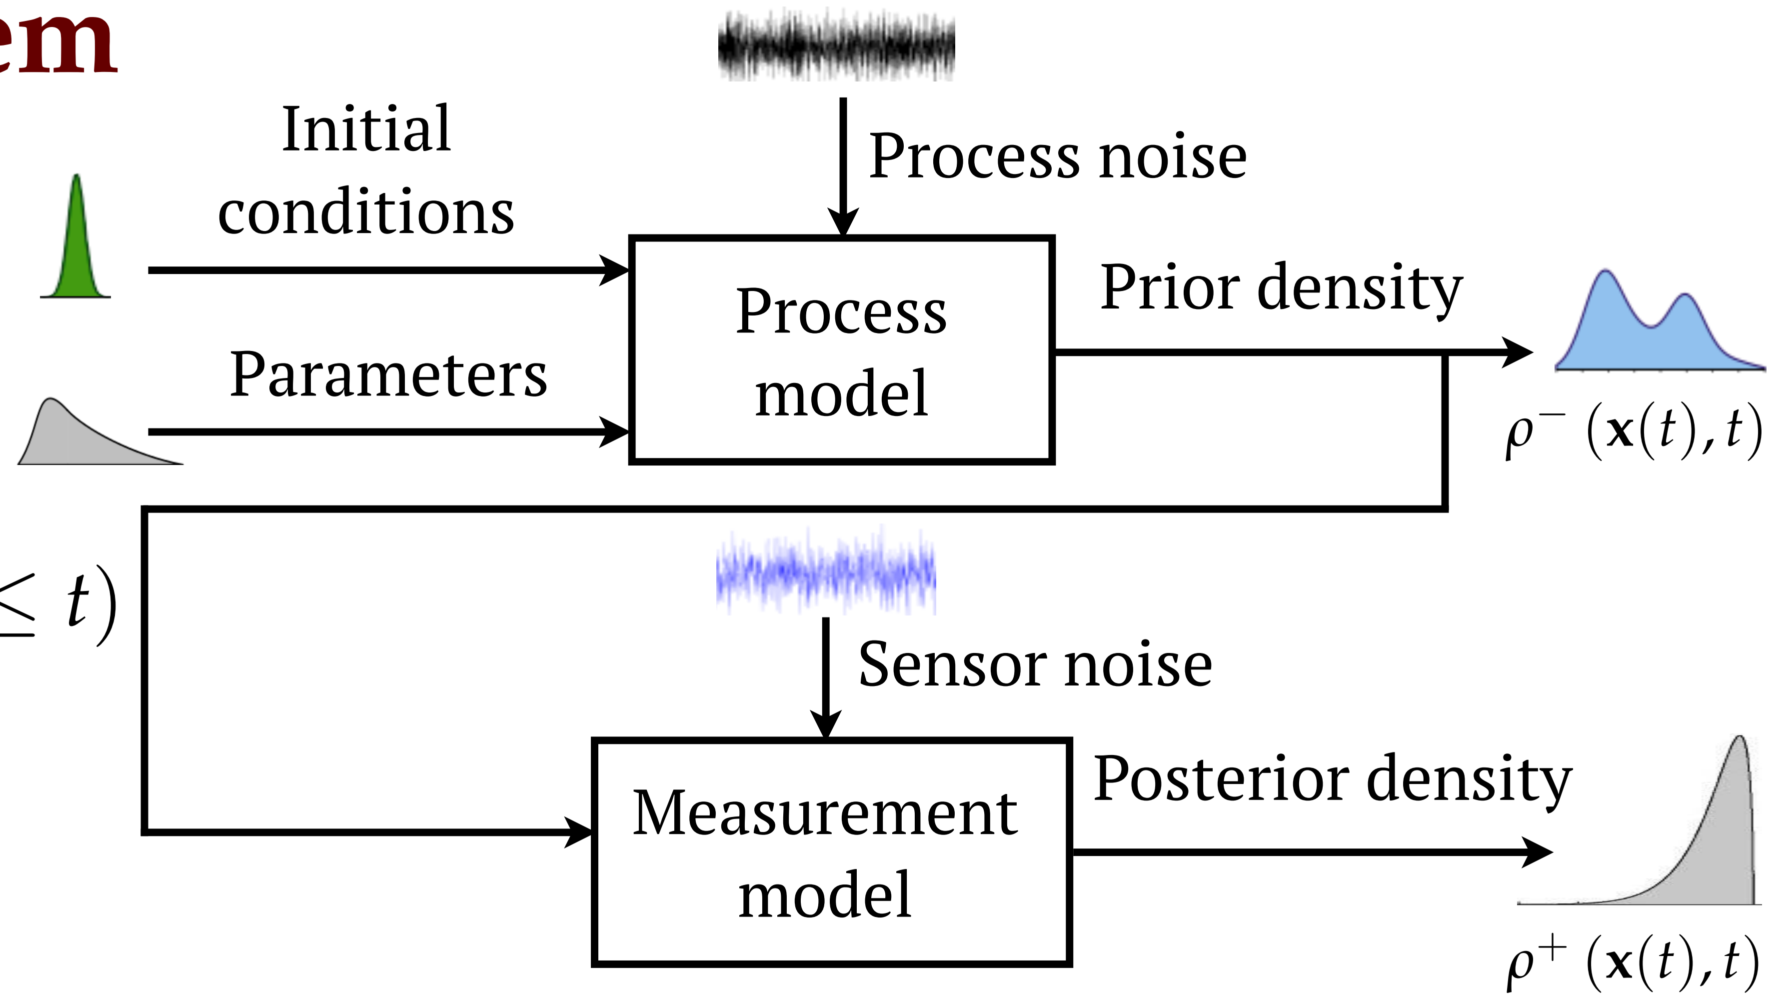

Trajectory flow:

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

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

Density flow:

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

Control Problem

Steer joint state PDF via feedback control over finite time horizon

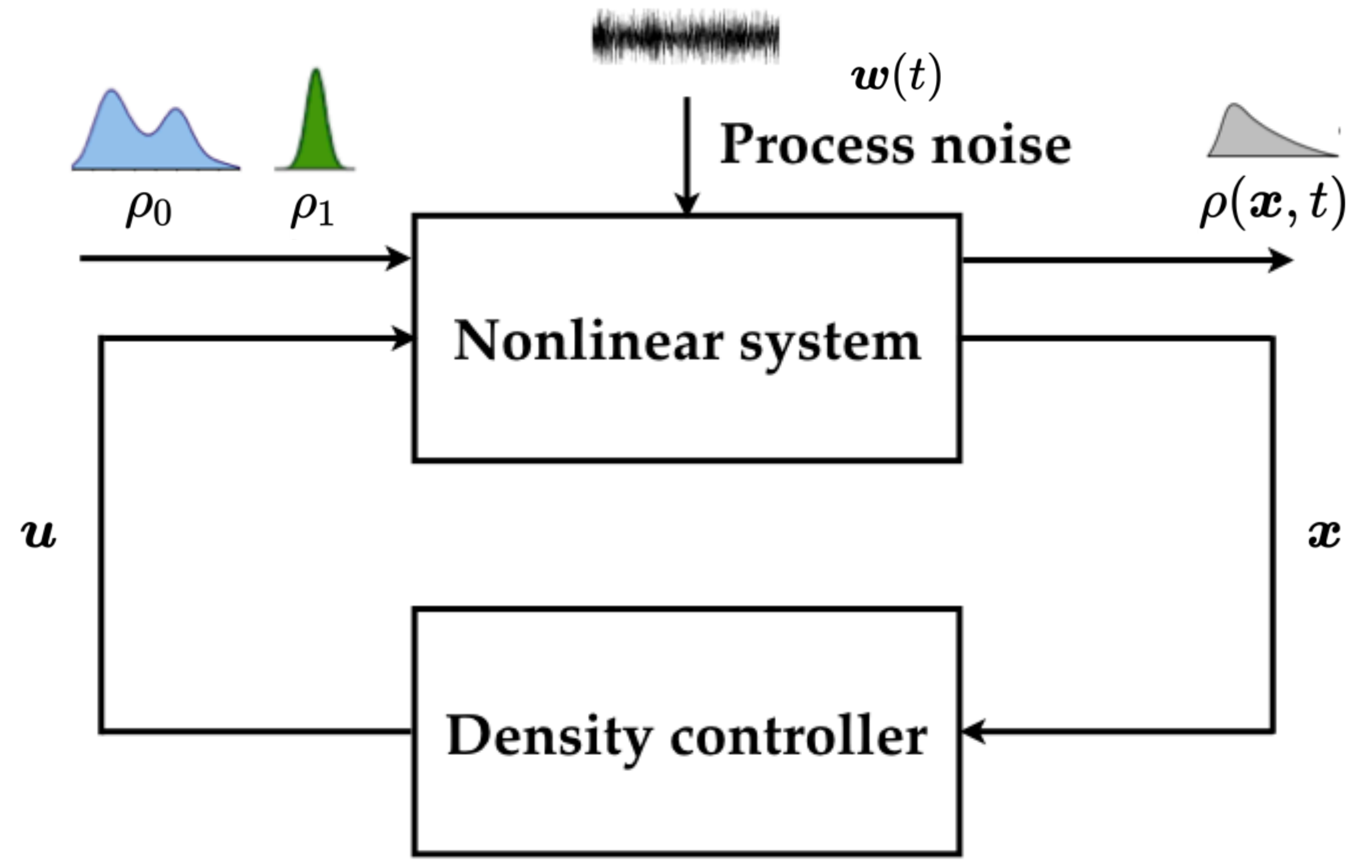

$$\begin{aligned}
&\underset{u\in\mathcal{U}}{\text{minimize}} \quad \mathbb{E}\left[\int_0^1 \|\boldsymbol{u}\|_2^2 \,\mathrm{d}t\right] \\
&\text{subject to} \\
&\mathrm{d}\boldsymbol{x} = \boldsymbol{f}\left(\boldsymbol{x},\boldsymbol{u},t\right)\,\mathrm{d}t + \boldsymbol{g}\left(\boldsymbol{x},t\right)\,\mathrm{d}\boldsymbol{w}, \\
&\boldsymbol{x}\left(t=0\right) \sim \rho_0, \quad \boldsymbol{x}\left(t=1\right) \sim \rho_1
\end{aligned}$$

PDFs in Mars Entry-Descent-Landing

Prediction Problem

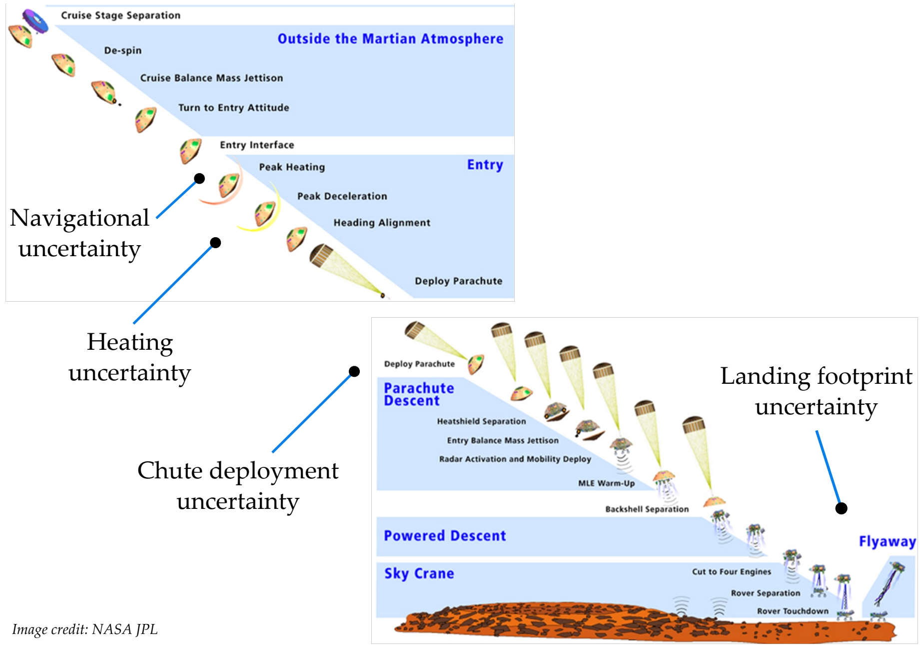

Image credit: NASA JPL

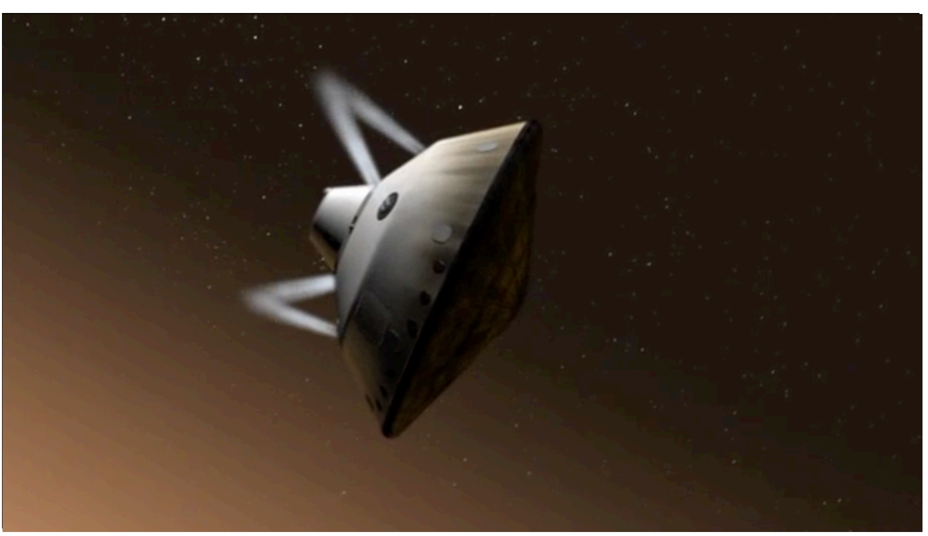

Predict heating rate uncertainty

Filtering Problem

Control Problem

PDFs in Mars Entry-Descent-Landing

Prediction Problem

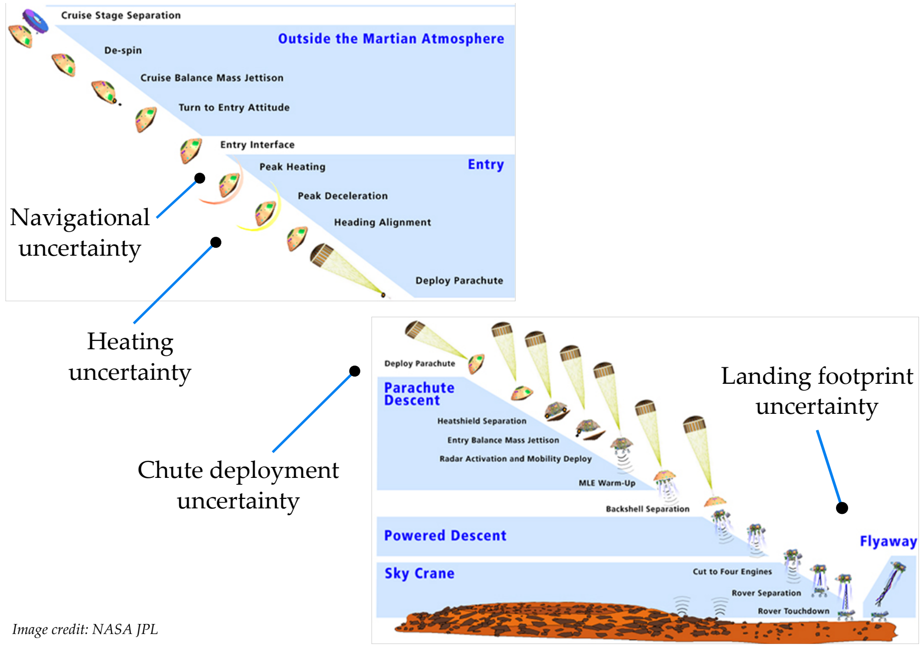

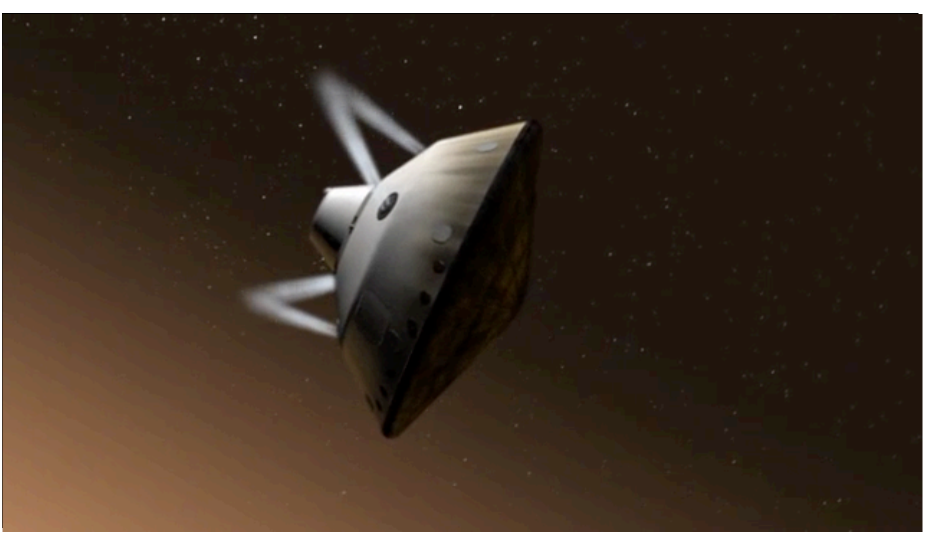

Predict heating rate uncertainty

Filtering Problem

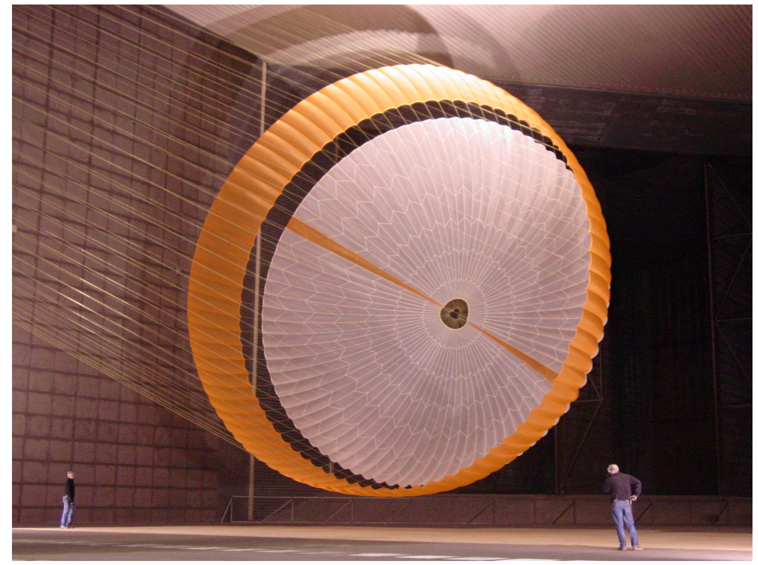

Supersonic parachute

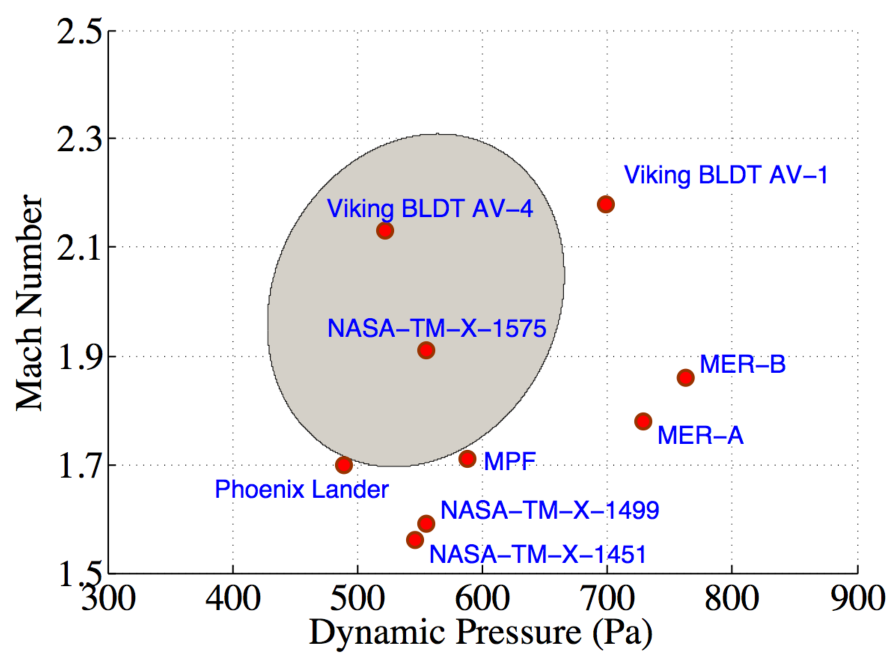

Estimate state to deploy parachute

Control Problem

PDFs in Mars Entry-Descent-Landing

Prediction Problem

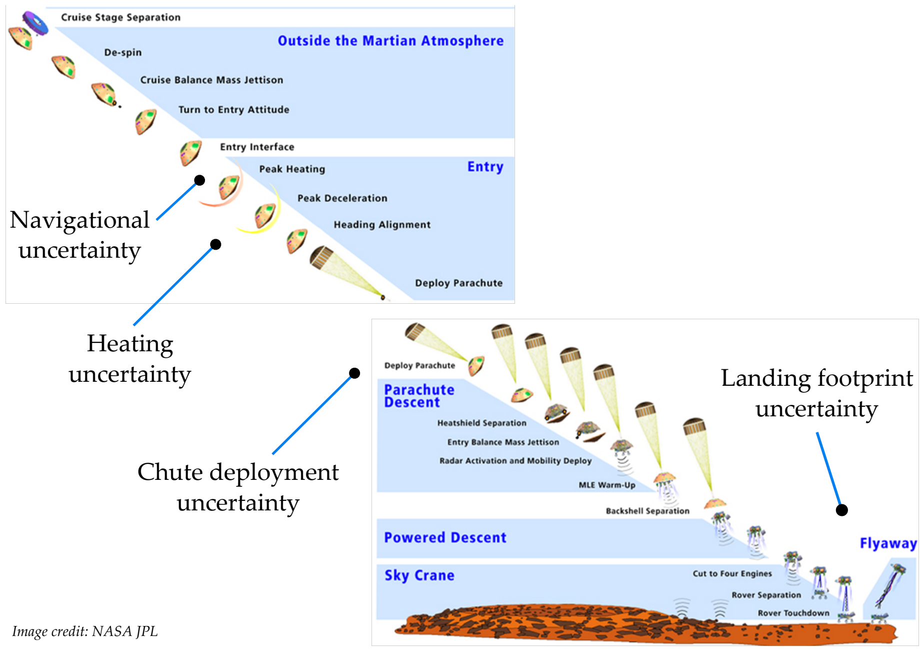

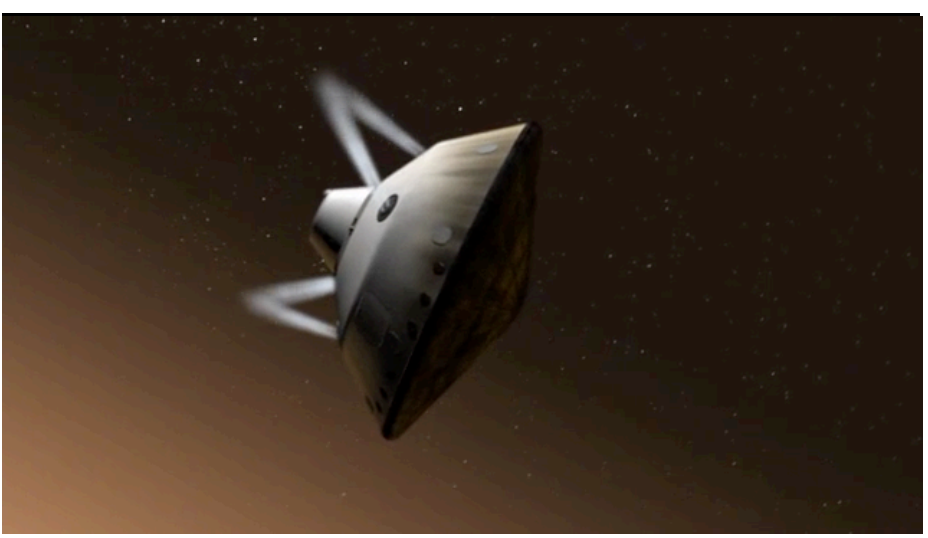

Filtering Problem

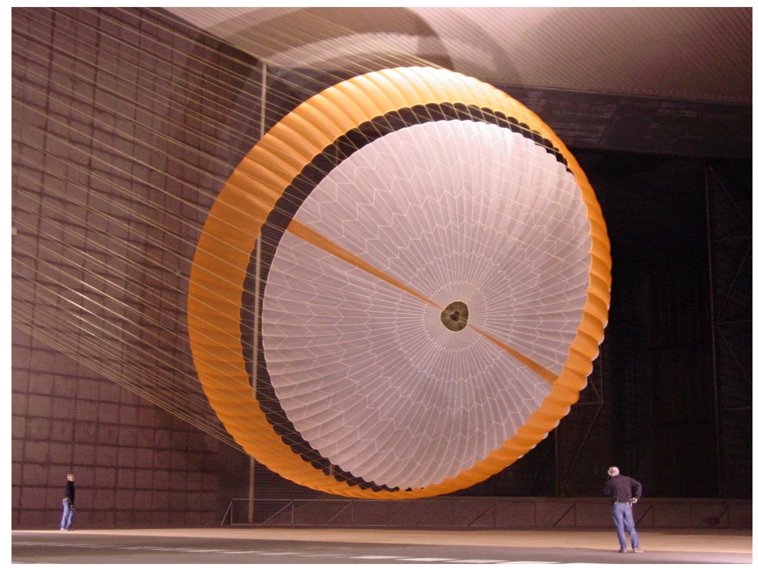

Supersonic parachute

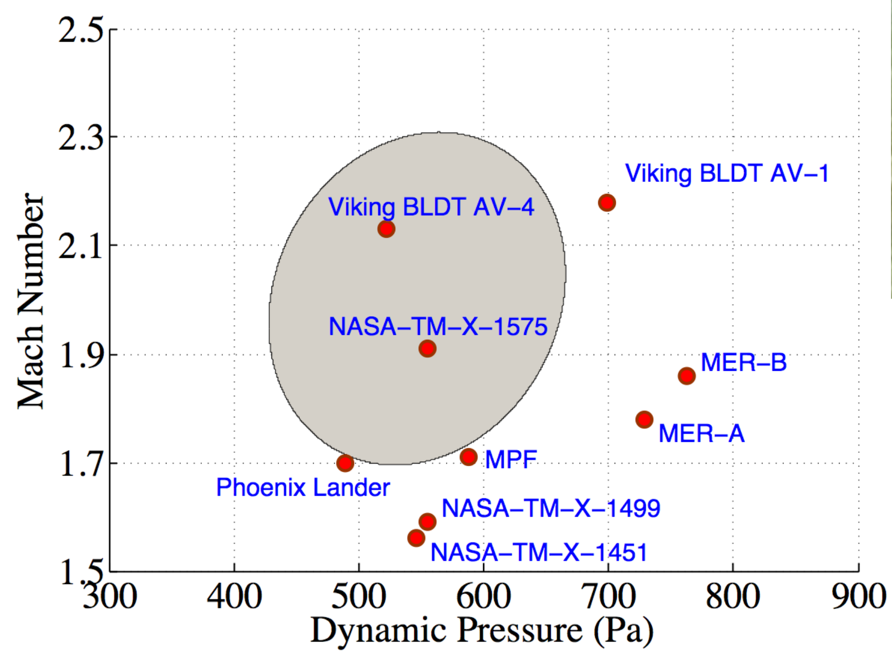

Control Problem

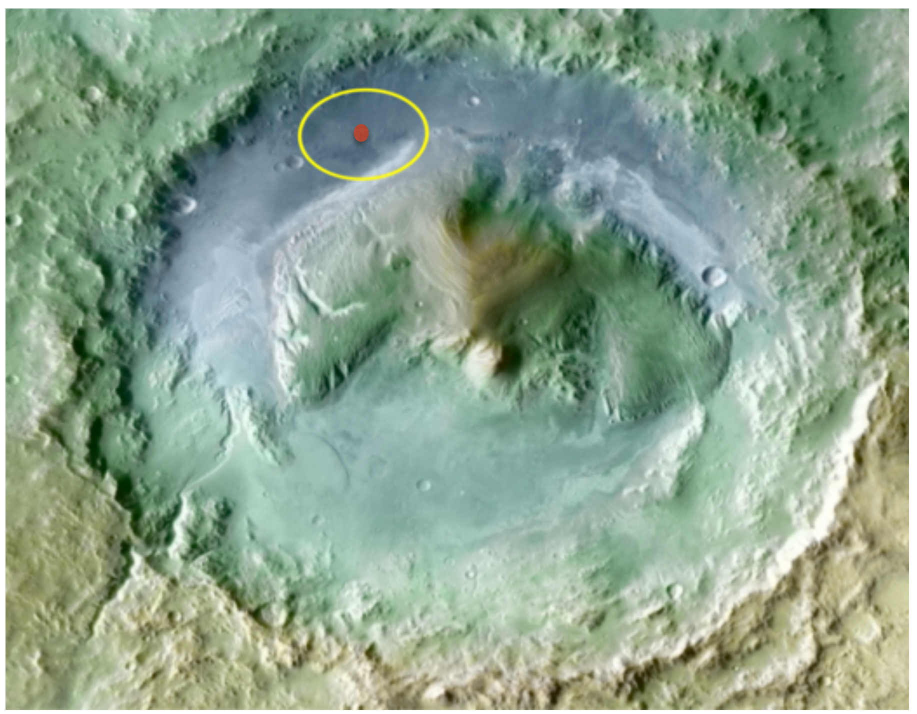

Gale Crater (4.49S, 137.42E)

Predict heating rate uncertainty

Estimate state to deploy parachute

Steer state PDF to achieve desired landing footprint accuracy

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:

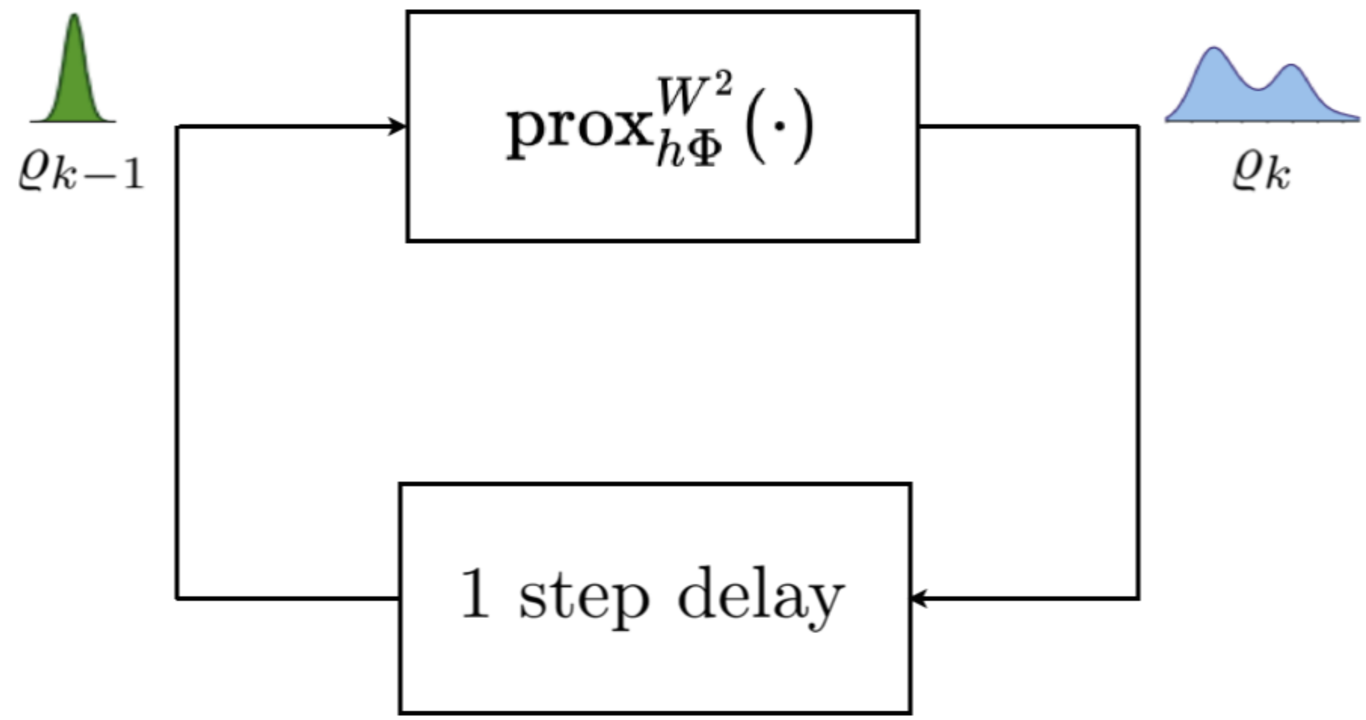

Proximal operator: $\varrho_k = \mathrm{prox}_{h\Phi}^{W^2}(\varrho_{k-1}) := \underset{\varrho\in\mathcal{P}_2(\mathcal{X})}{\arg\inf}\left\{\frac{1}{2}W^2(\varrho,\varrho_{k-1}) + h\Phi(\varrho)\right\}$

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

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

Geometric Meaning of Gradient Flow

Gradient Flow in $\mathcal{X}$

$$\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}t} = -\nabla\varphi(\boldsymbol{x}), \quad \boldsymbol{x}(0) = \boldsymbol{x}_0$$

Recursion:

$$\begin{aligned}\boldsymbol{x}_k &= \boldsymbol{x}_{k-1} - h\nabla\varphi(\boldsymbol{x}_k)\\ &= \underset{\boldsymbol{x}\in\mathcal{X}}{\arg\min}\left\{\frac{1}{2}\|\boldsymbol{x}-\boldsymbol{x}_{k-1}\|_2^2 + h\varphi(\boldsymbol{x})\right\}\\ &=: \operatorname{prox}_{h\varphi}^{\|\cdot\|_2}(\boldsymbol{x}_{k-1})\end{aligned}$$

Convergence:

$$\boldsymbol{x}_k \to \boldsymbol{x}(t=kh) \quad \text{as} \quad h\downarrow 0$$

φ as Lyapunov function:

$$\frac{\mathrm{d}}{\mathrm{d}t}\varphi = -\|\nabla\varphi\|_2^2 \;\leq\; 0$$

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

$$\frac{\partial\rho}{\partial t} = -\nabla^{W}\Phi(\rho), \quad \rho(\boldsymbol{x},0) = \rho_0$$

Recursion:

$$\begin{aligned}\rho_k &= \rho(\cdot, t=kh)\\ &= \underset{\rho\in\mathcal{P}_2(\mathcal{X})}{\arg\min}\left\{\frac{1}{2}W^2(\rho,\rho_{k-1}) + h\Phi(\rho)\right\}\\ &=: \operatorname{prox}_{h\Phi}^{W^2}(\rho_{k-1})\end{aligned}$$

Convergence:

$$\rho_k \to \rho(\cdot, t=kh) \quad \text{as} \quad h\downarrow 0$$

Φ as Lyapunov functional:

$$\frac{\mathrm{d}}{\mathrm{d}t}\Phi = -\mathbb{E}_{\rho}\left[\left\|\nabla\frac{\delta\Phi}{\delta\rho}\right\|_2^2\right] \;\leq\; 0$$

Geometric Meaning of Gradient Flow

Gradient Flow in $\mathcal{X}$

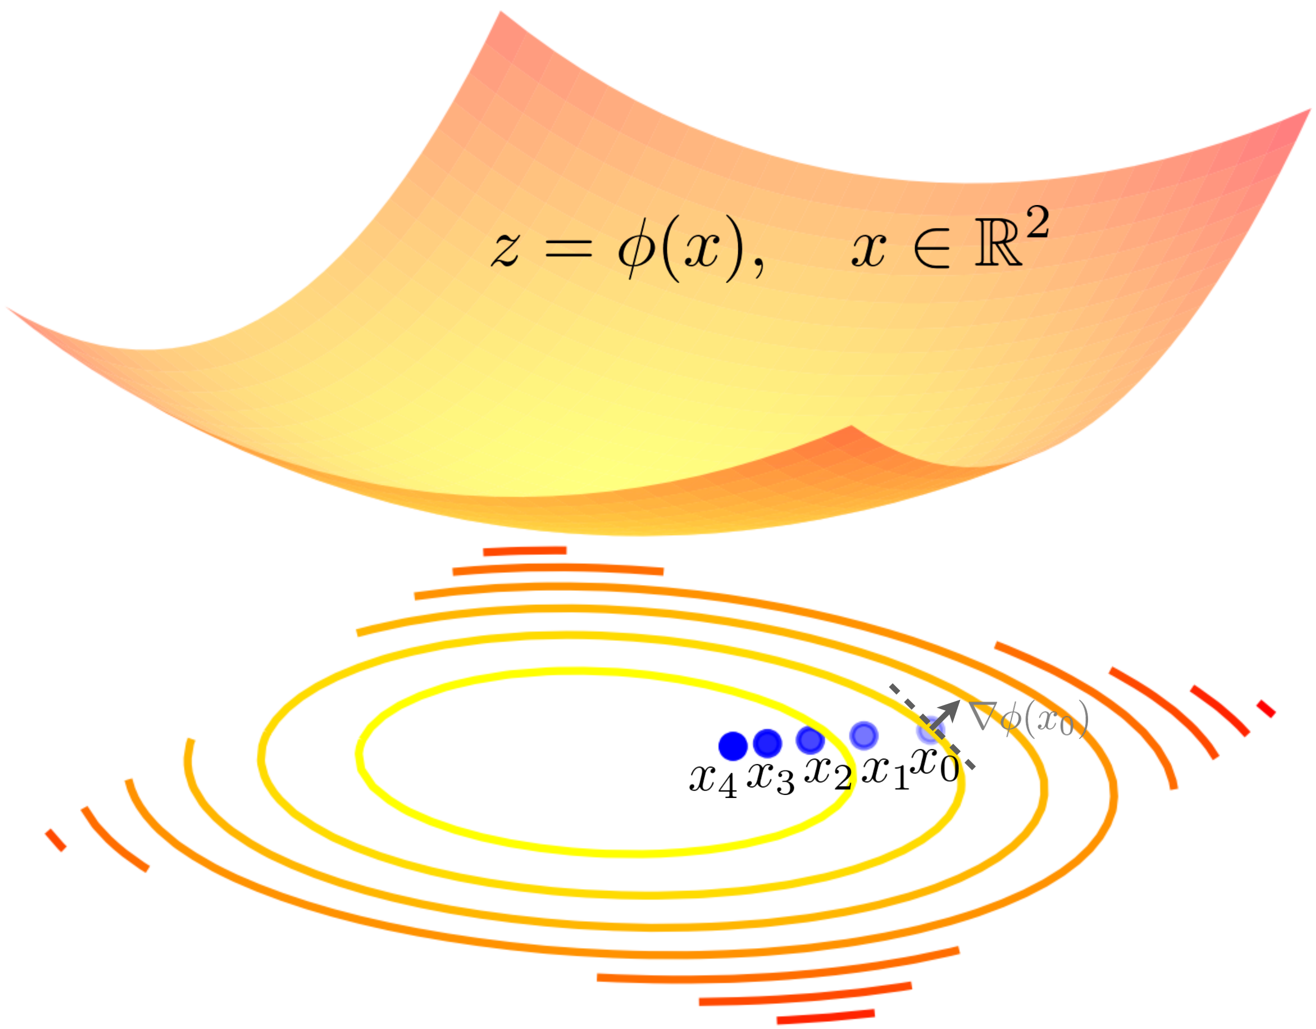

Gradient Flow in $\mathcal{P}_2(\mathcal{X})$

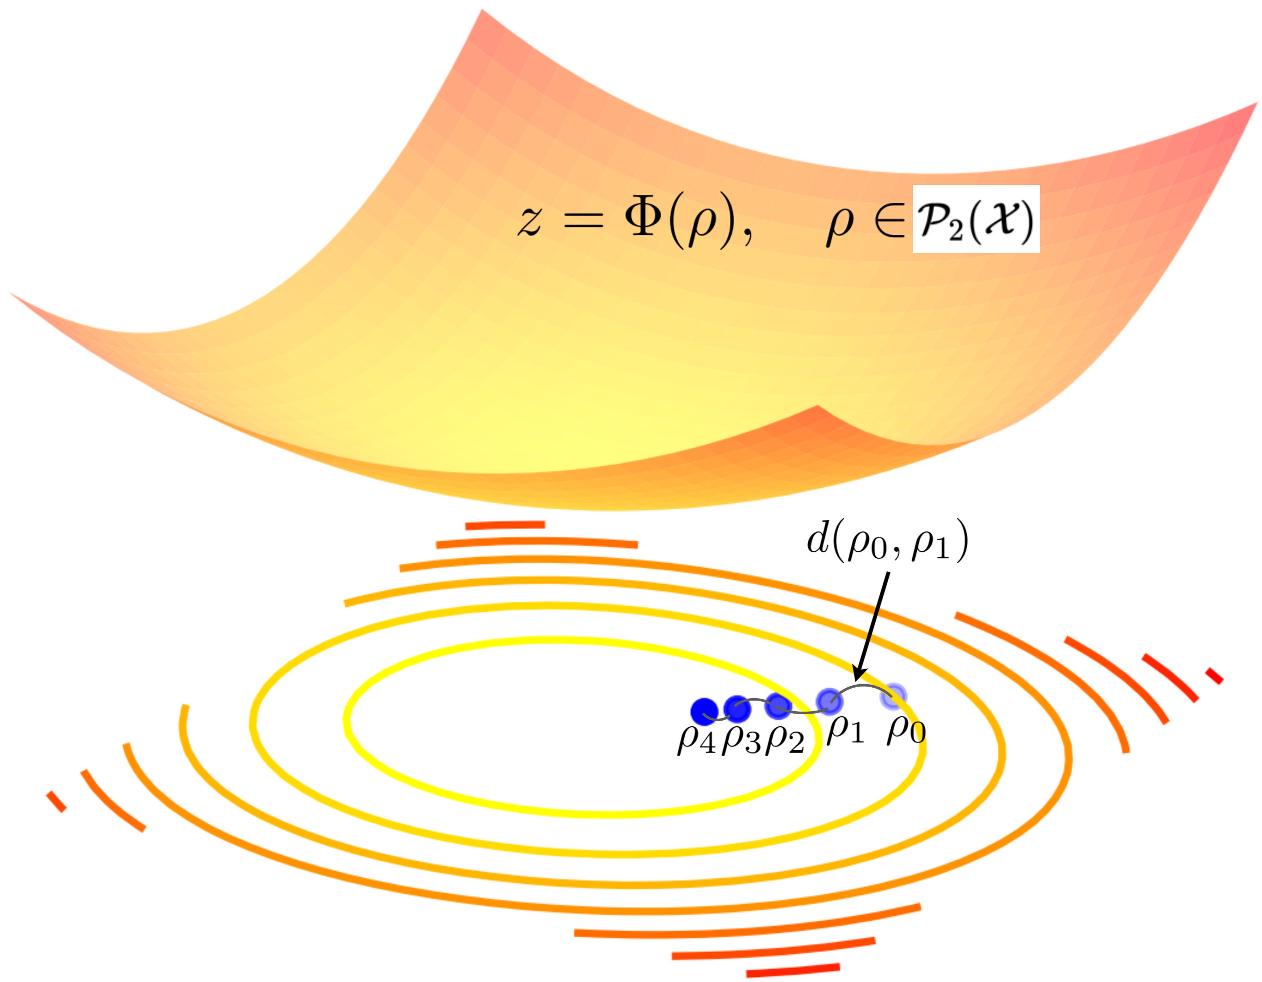

Algorithm: Gradient Ascent on the Dual Space

Uncertainty propagation via point clouds

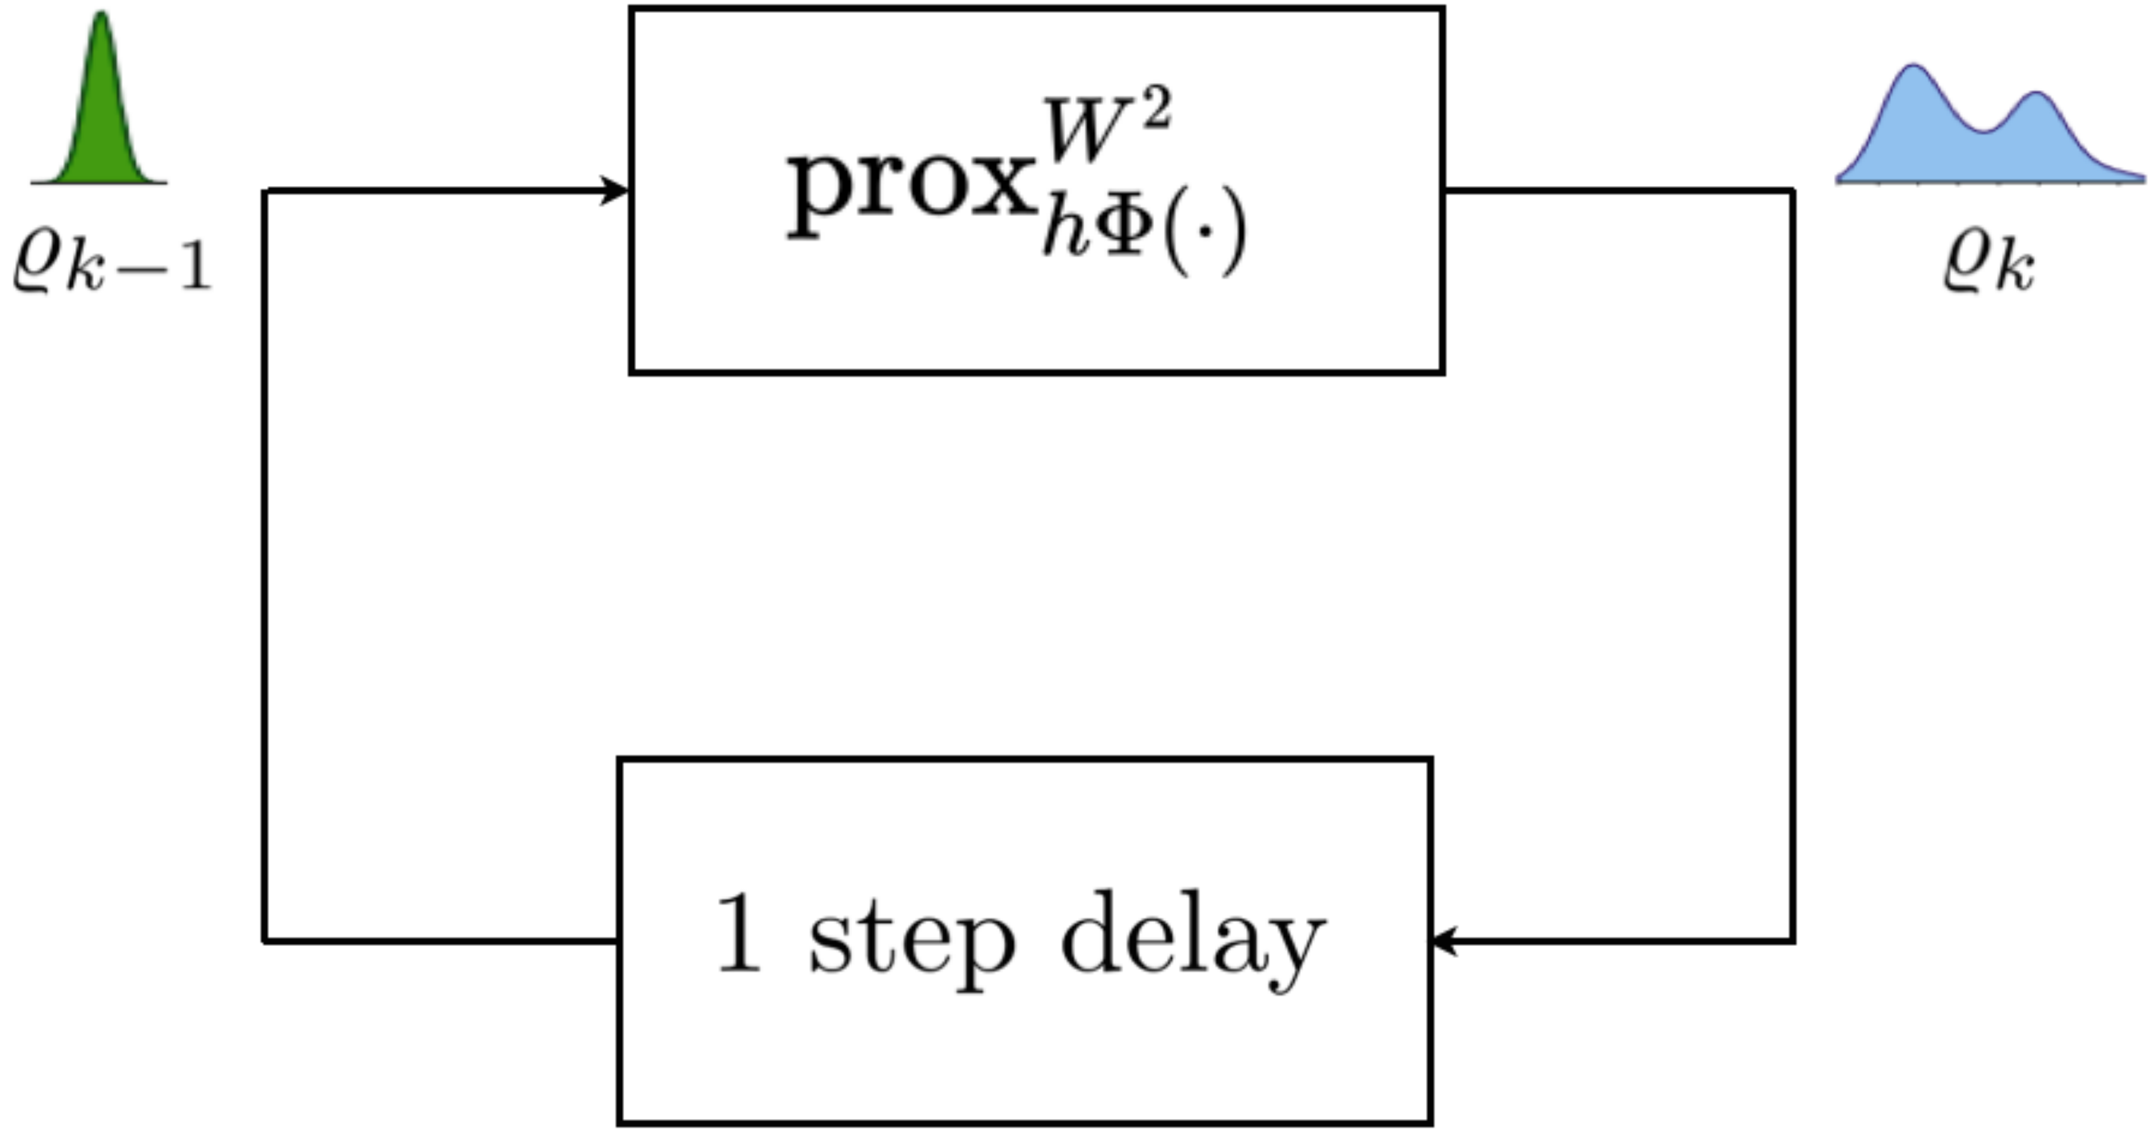

No spatial discretization or function approximation

Algorithm: Gradient Ascent on the Dual Space

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

$\Updownarrow$ **Proximal Recursion**

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

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**

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

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**

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

$\Downarrow$ **Entropic Regularization**

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

$\Updownarrow$ **Dualization**

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

Recursion on the Cone

$$\boldsymbol{y} = e^{\frac{\boldsymbol{\lambda}_0^*}{\epsilon}} h \Big\downarrow \qquad \Big\downarrow \boldsymbol{z} = e^{\frac{\boldsymbol{\lambda}_1^*}{\epsilon}} h$$

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

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

$$\boxed{\begin{aligned} \boldsymbol{y} \odot \boldsymbol{\Gamma}_k \boldsymbol{z} &= \varrho_{k-1} \\ \boldsymbol{z} \odot \boldsymbol{\Gamma}_k^\top \boldsymbol{y} &= \boldsymbol{\xi}_{k-1} \odot \boldsymbol{z}^{-\beta\epsilon/2h} \end{aligned}} \longrightarrow \varrho_k = \boldsymbol{z} \odot \boldsymbol{\Gamma}_k^\top \boldsymbol{y}$$

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 $\varrho_k = \boldsymbol{z}^* \odot (\boldsymbol{\Gamma}_k{}^\top \boldsymbol{y}^*)$

Algorithmic Setup

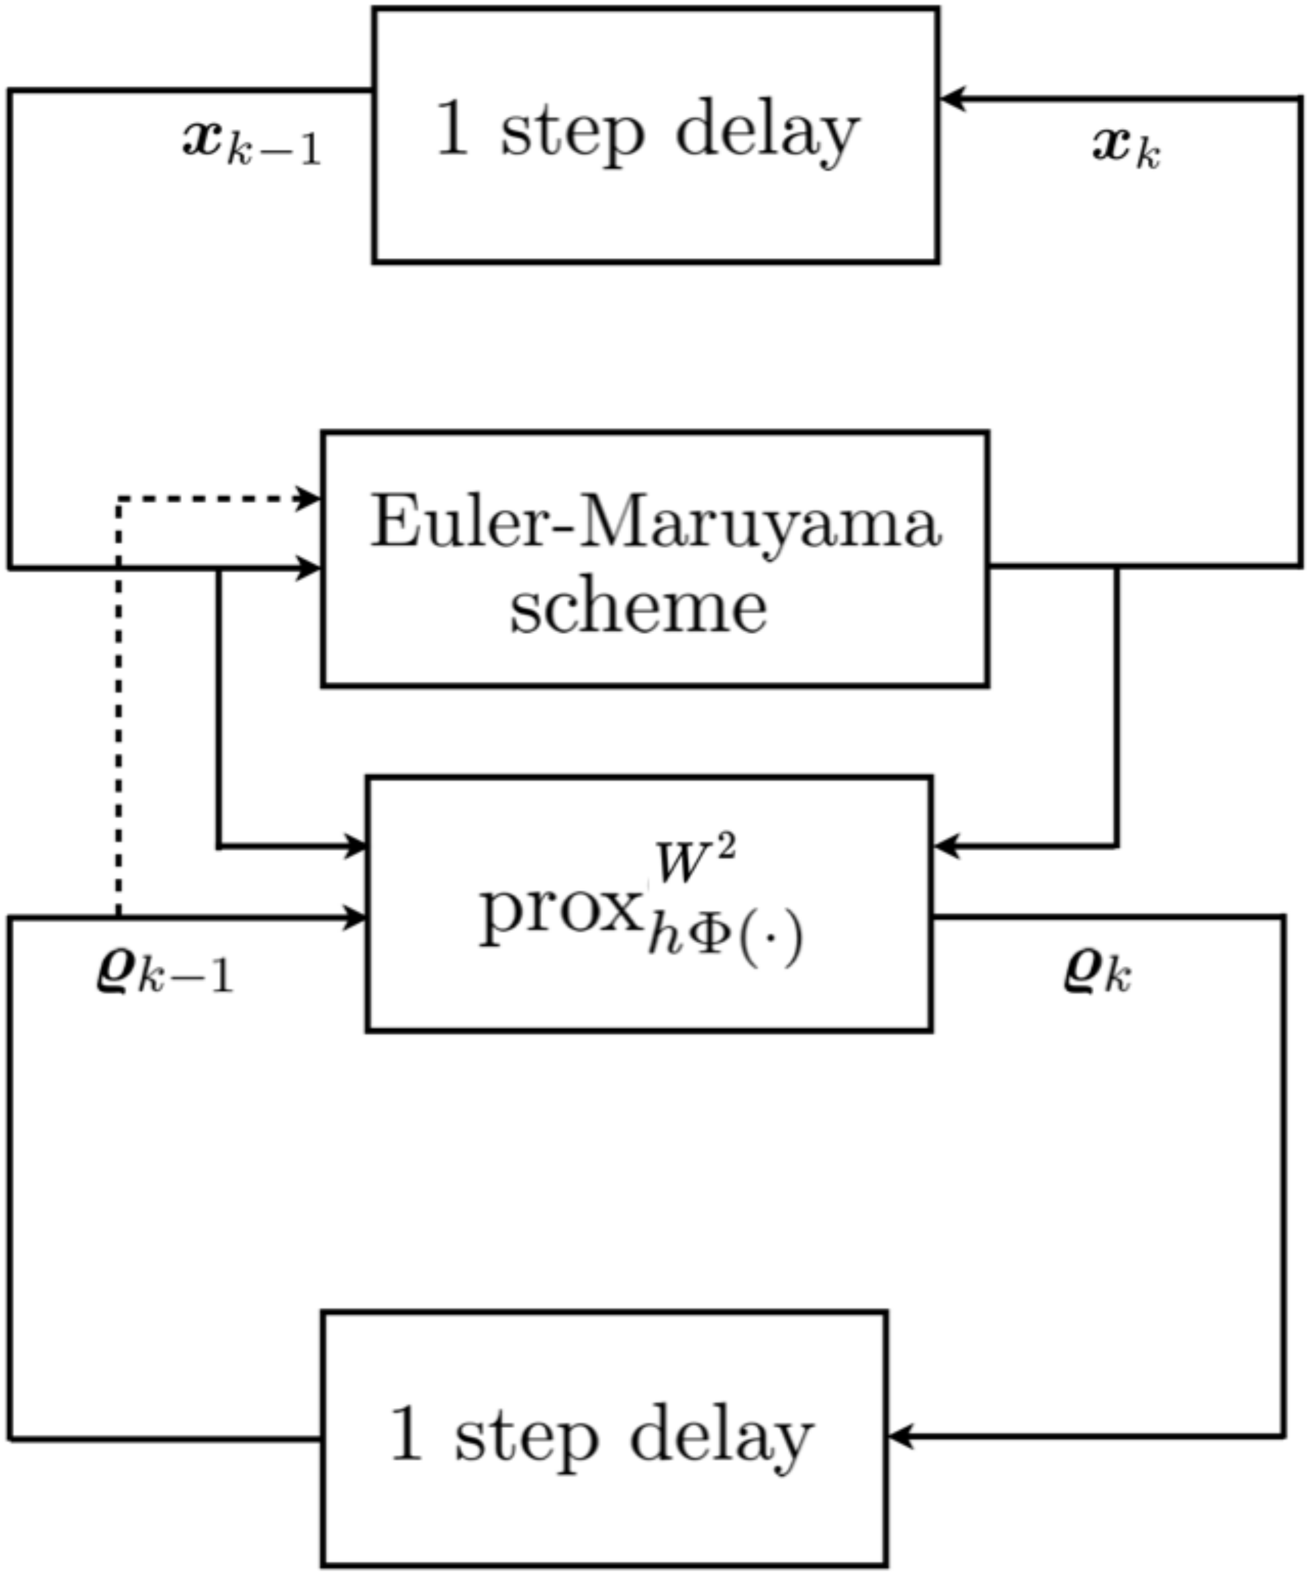

Theorem: Block co-ordinate iteration of $(\mathbf{y}, \mathbf{z})$ recursion is contractive on $\mathbb{R}^n_{>0} \times \mathbb{R}^n_{>0}$.

Proximal Prediction: 2D Linear Gaussian

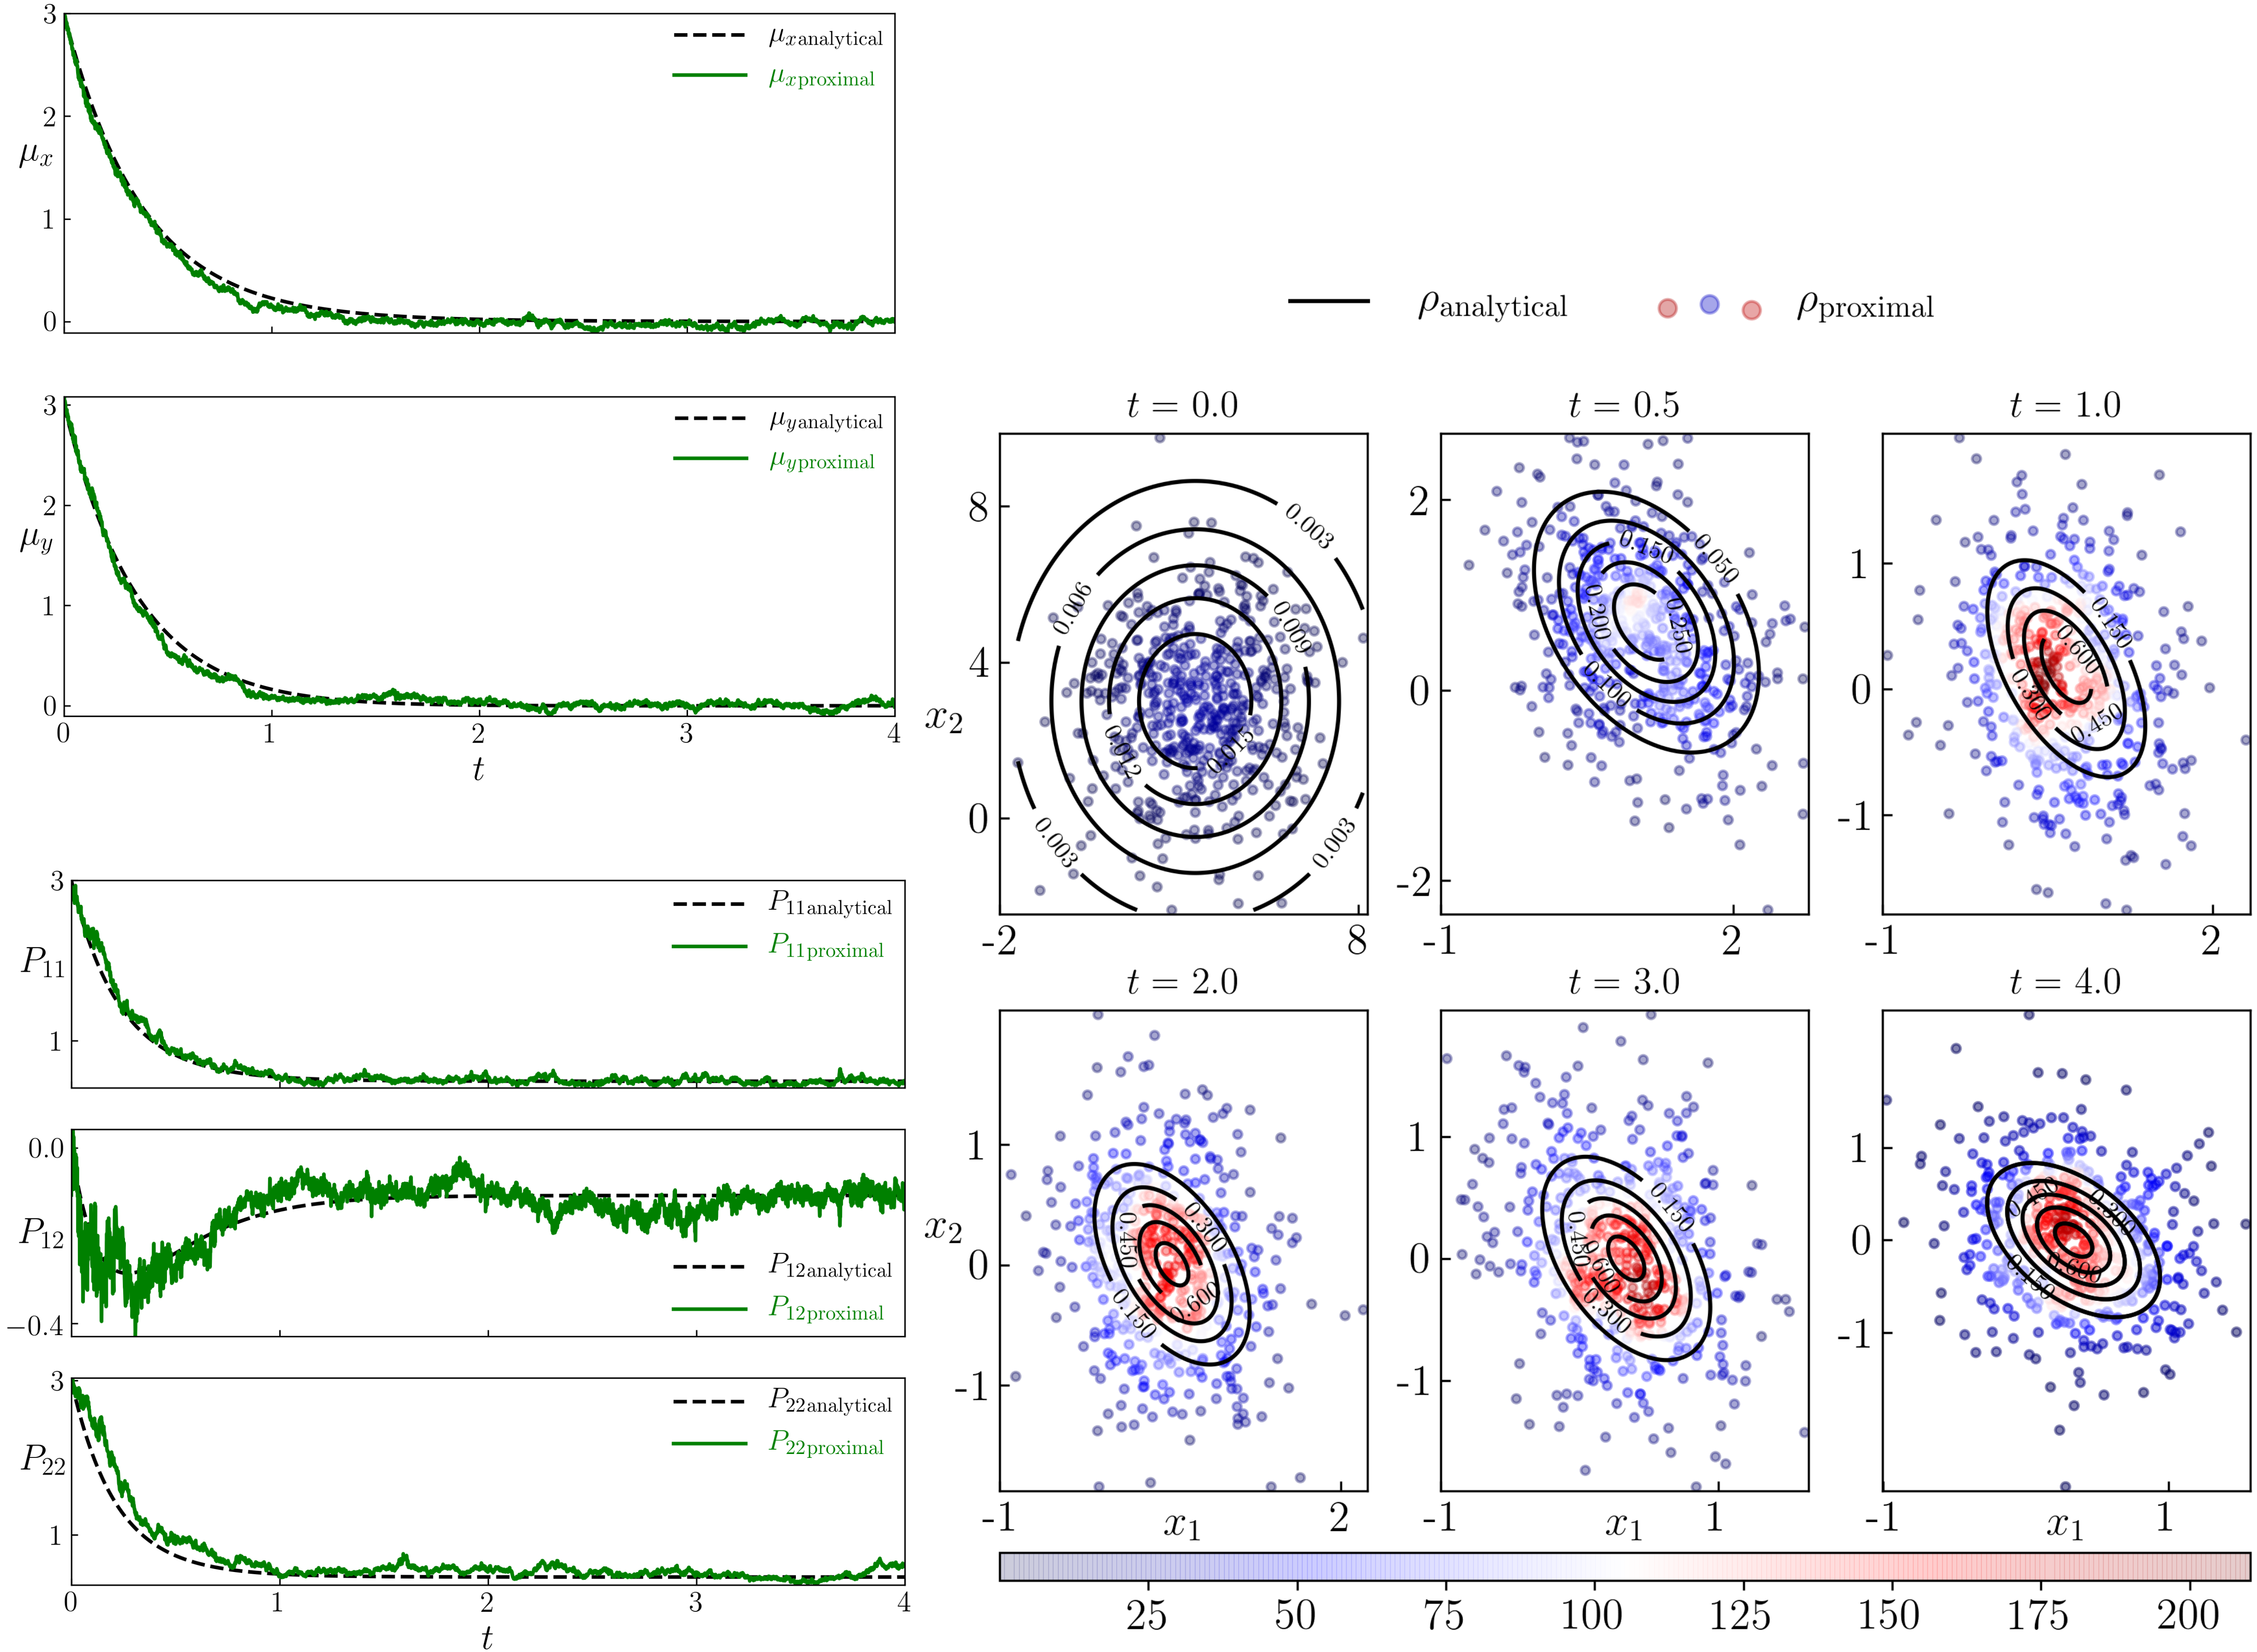

Proximal Prediction: Nonlinear Non-Gaussian

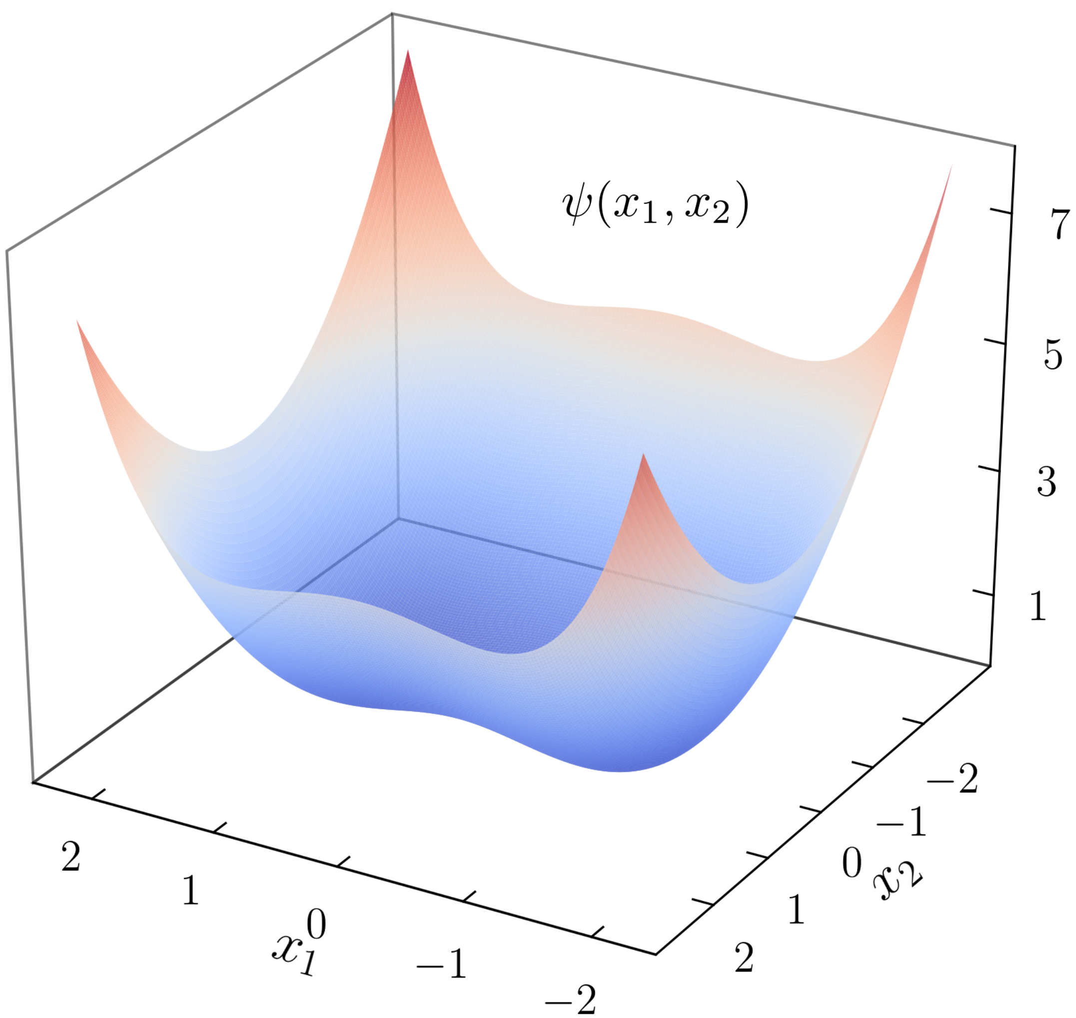

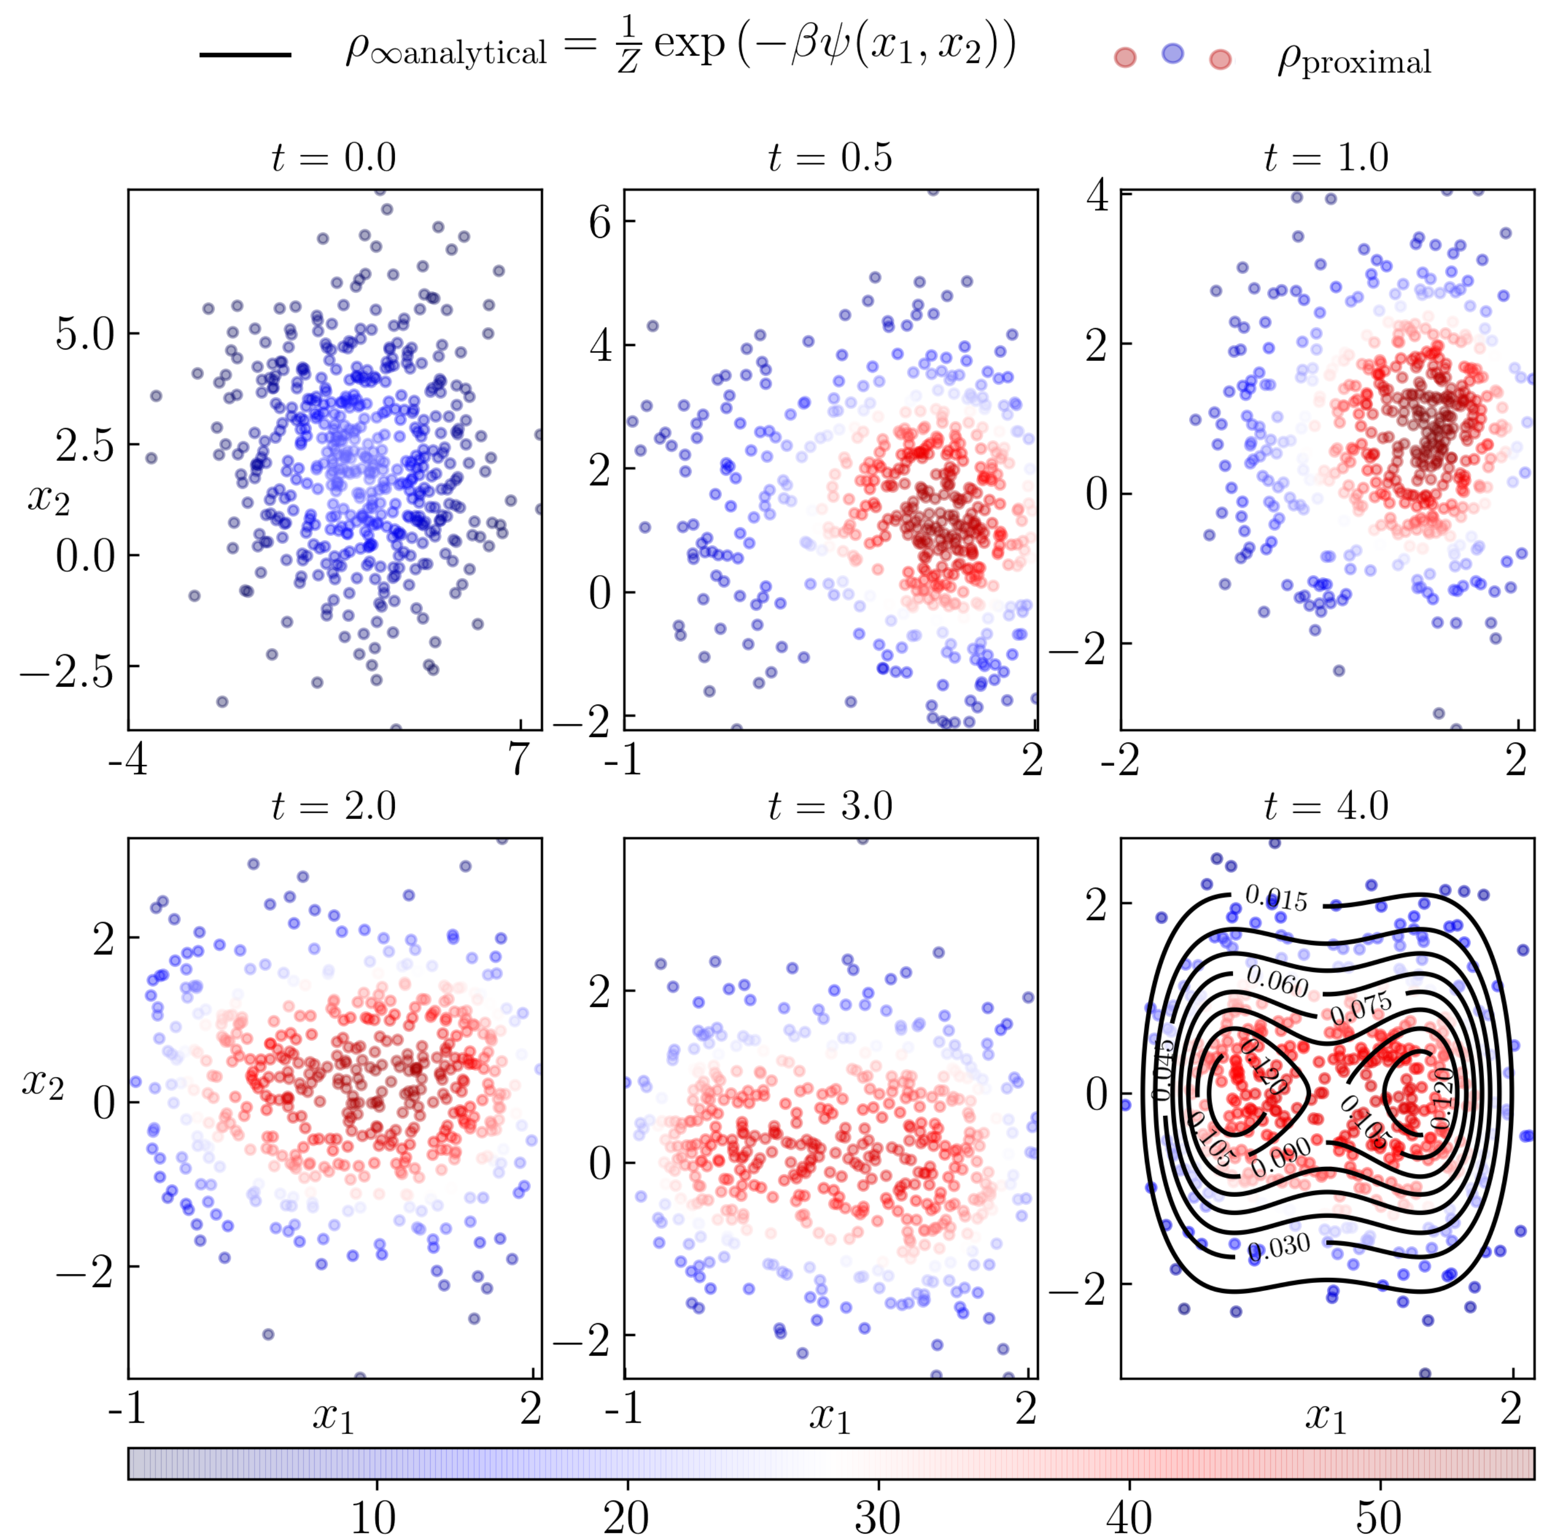

Computational Time: Nonlinear Non-Gaussian

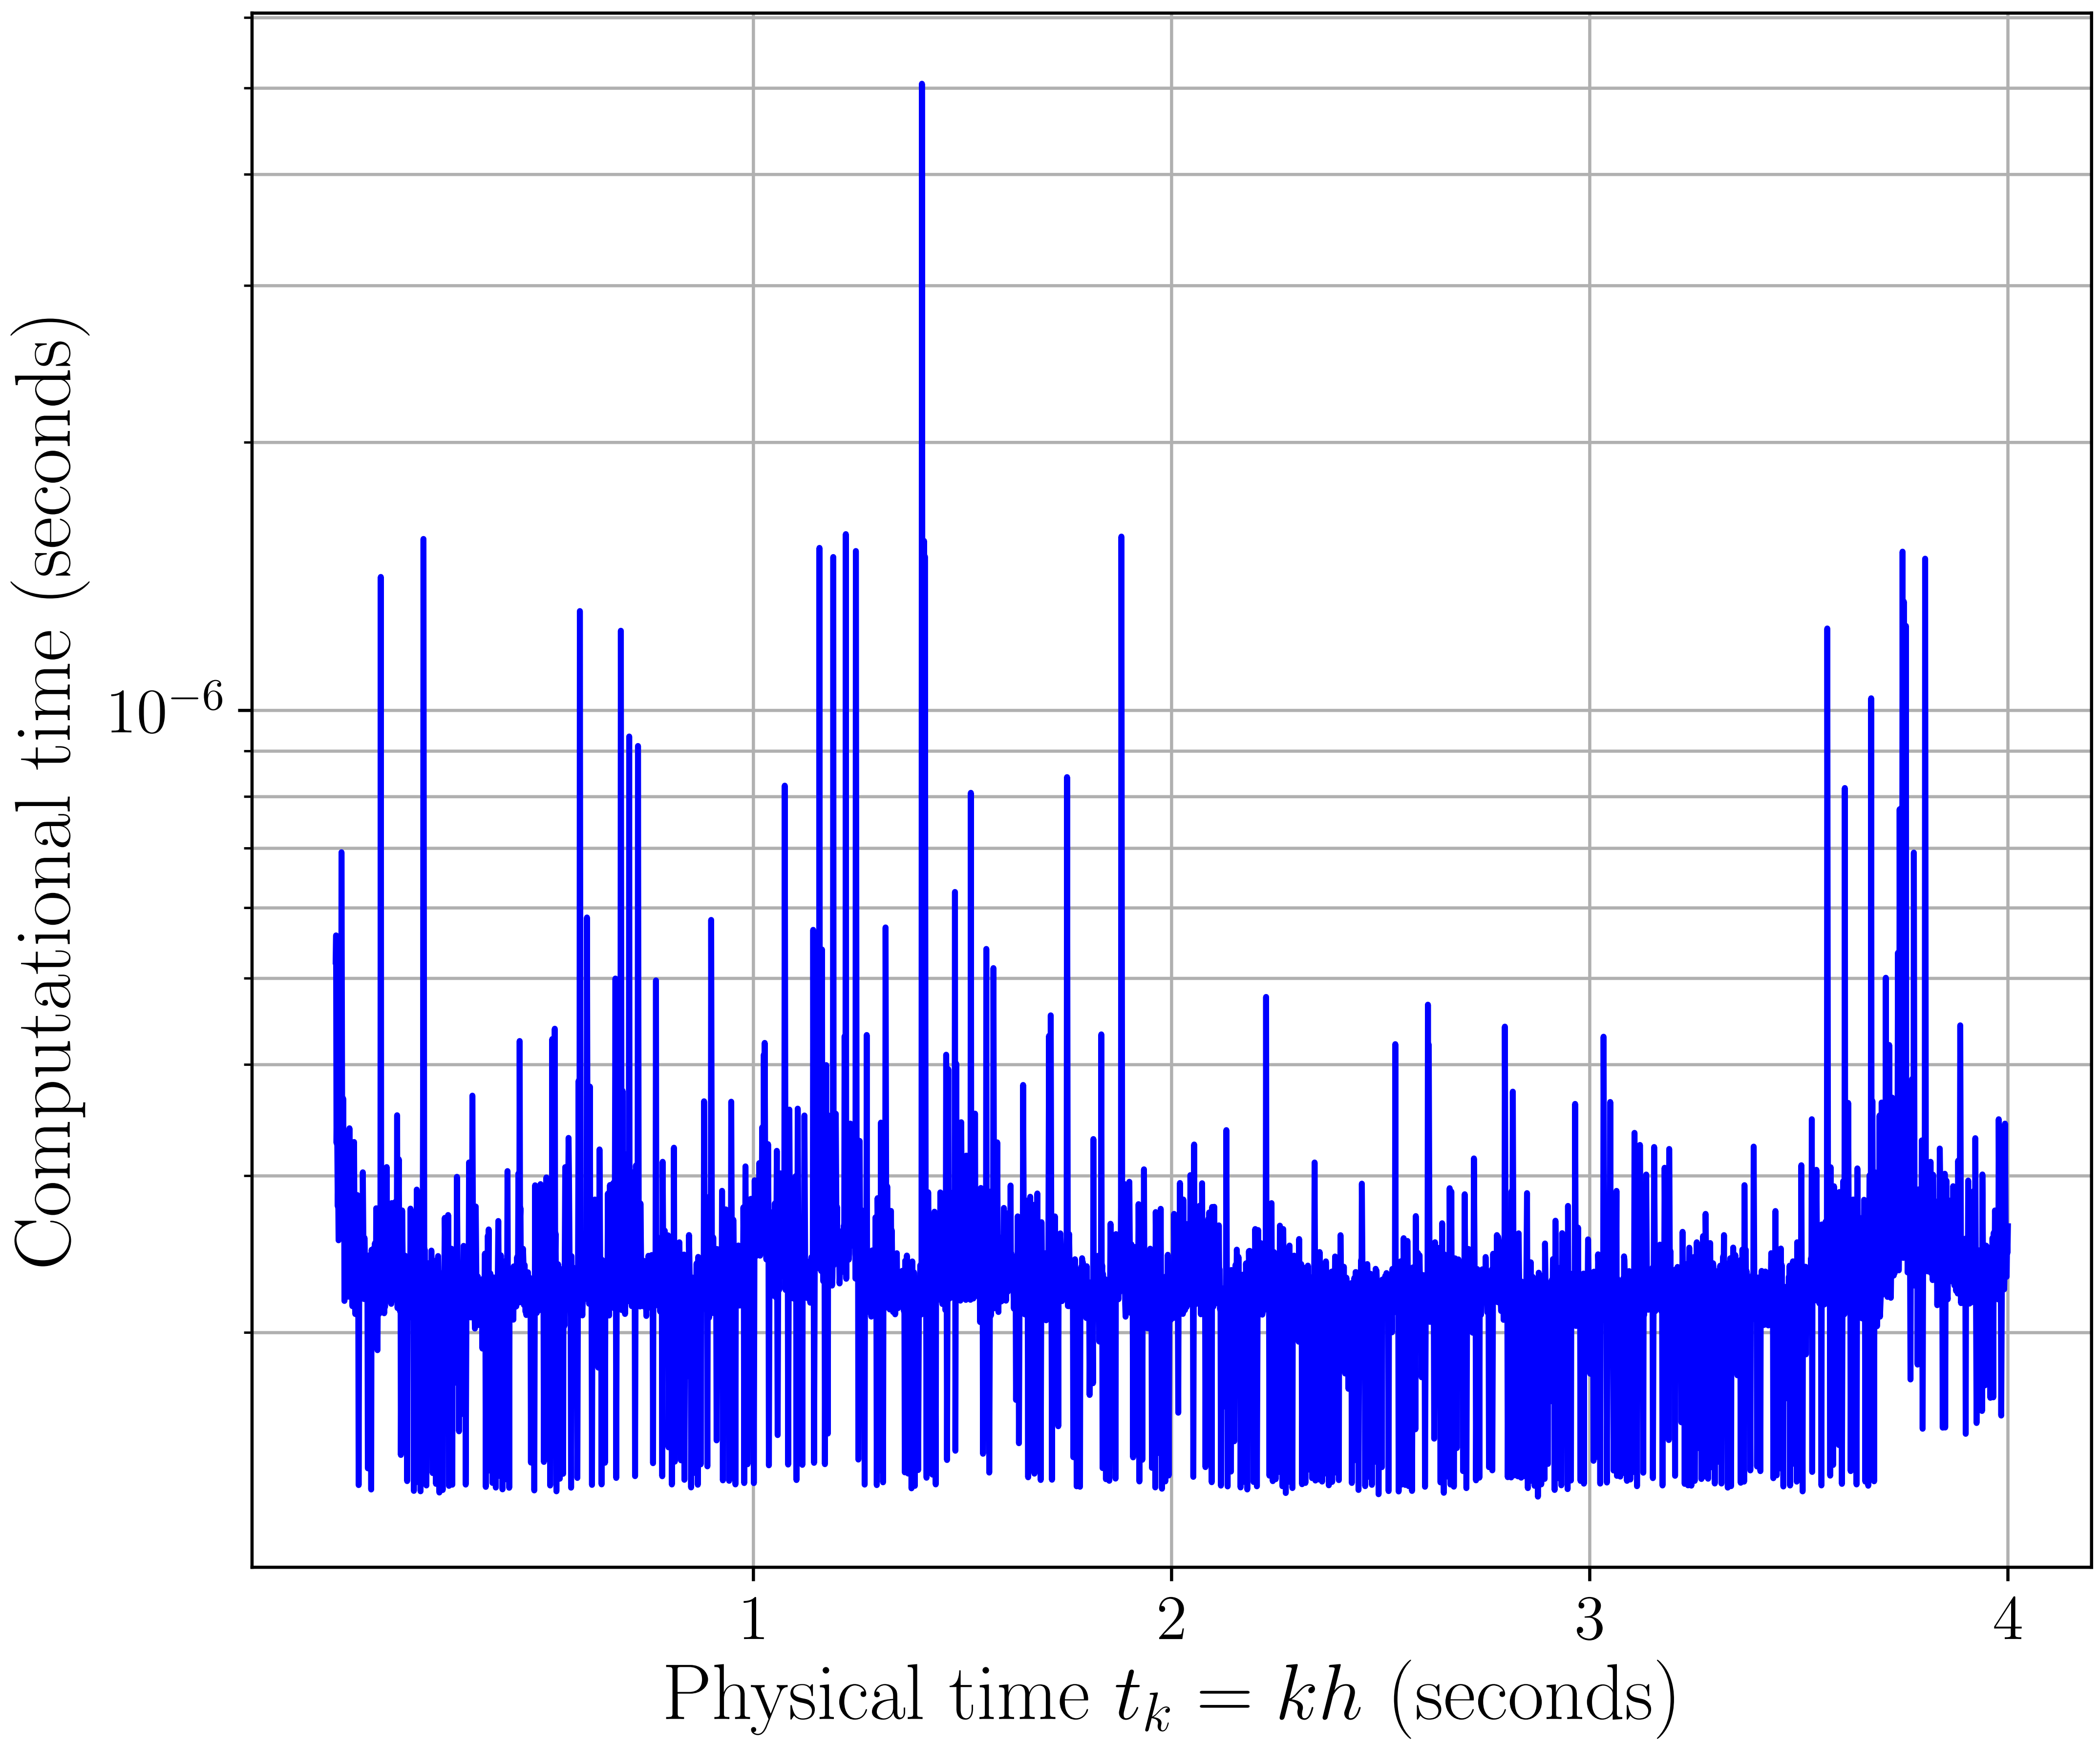

Proximal Prediction: Satellite in Geocentric Orbit

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

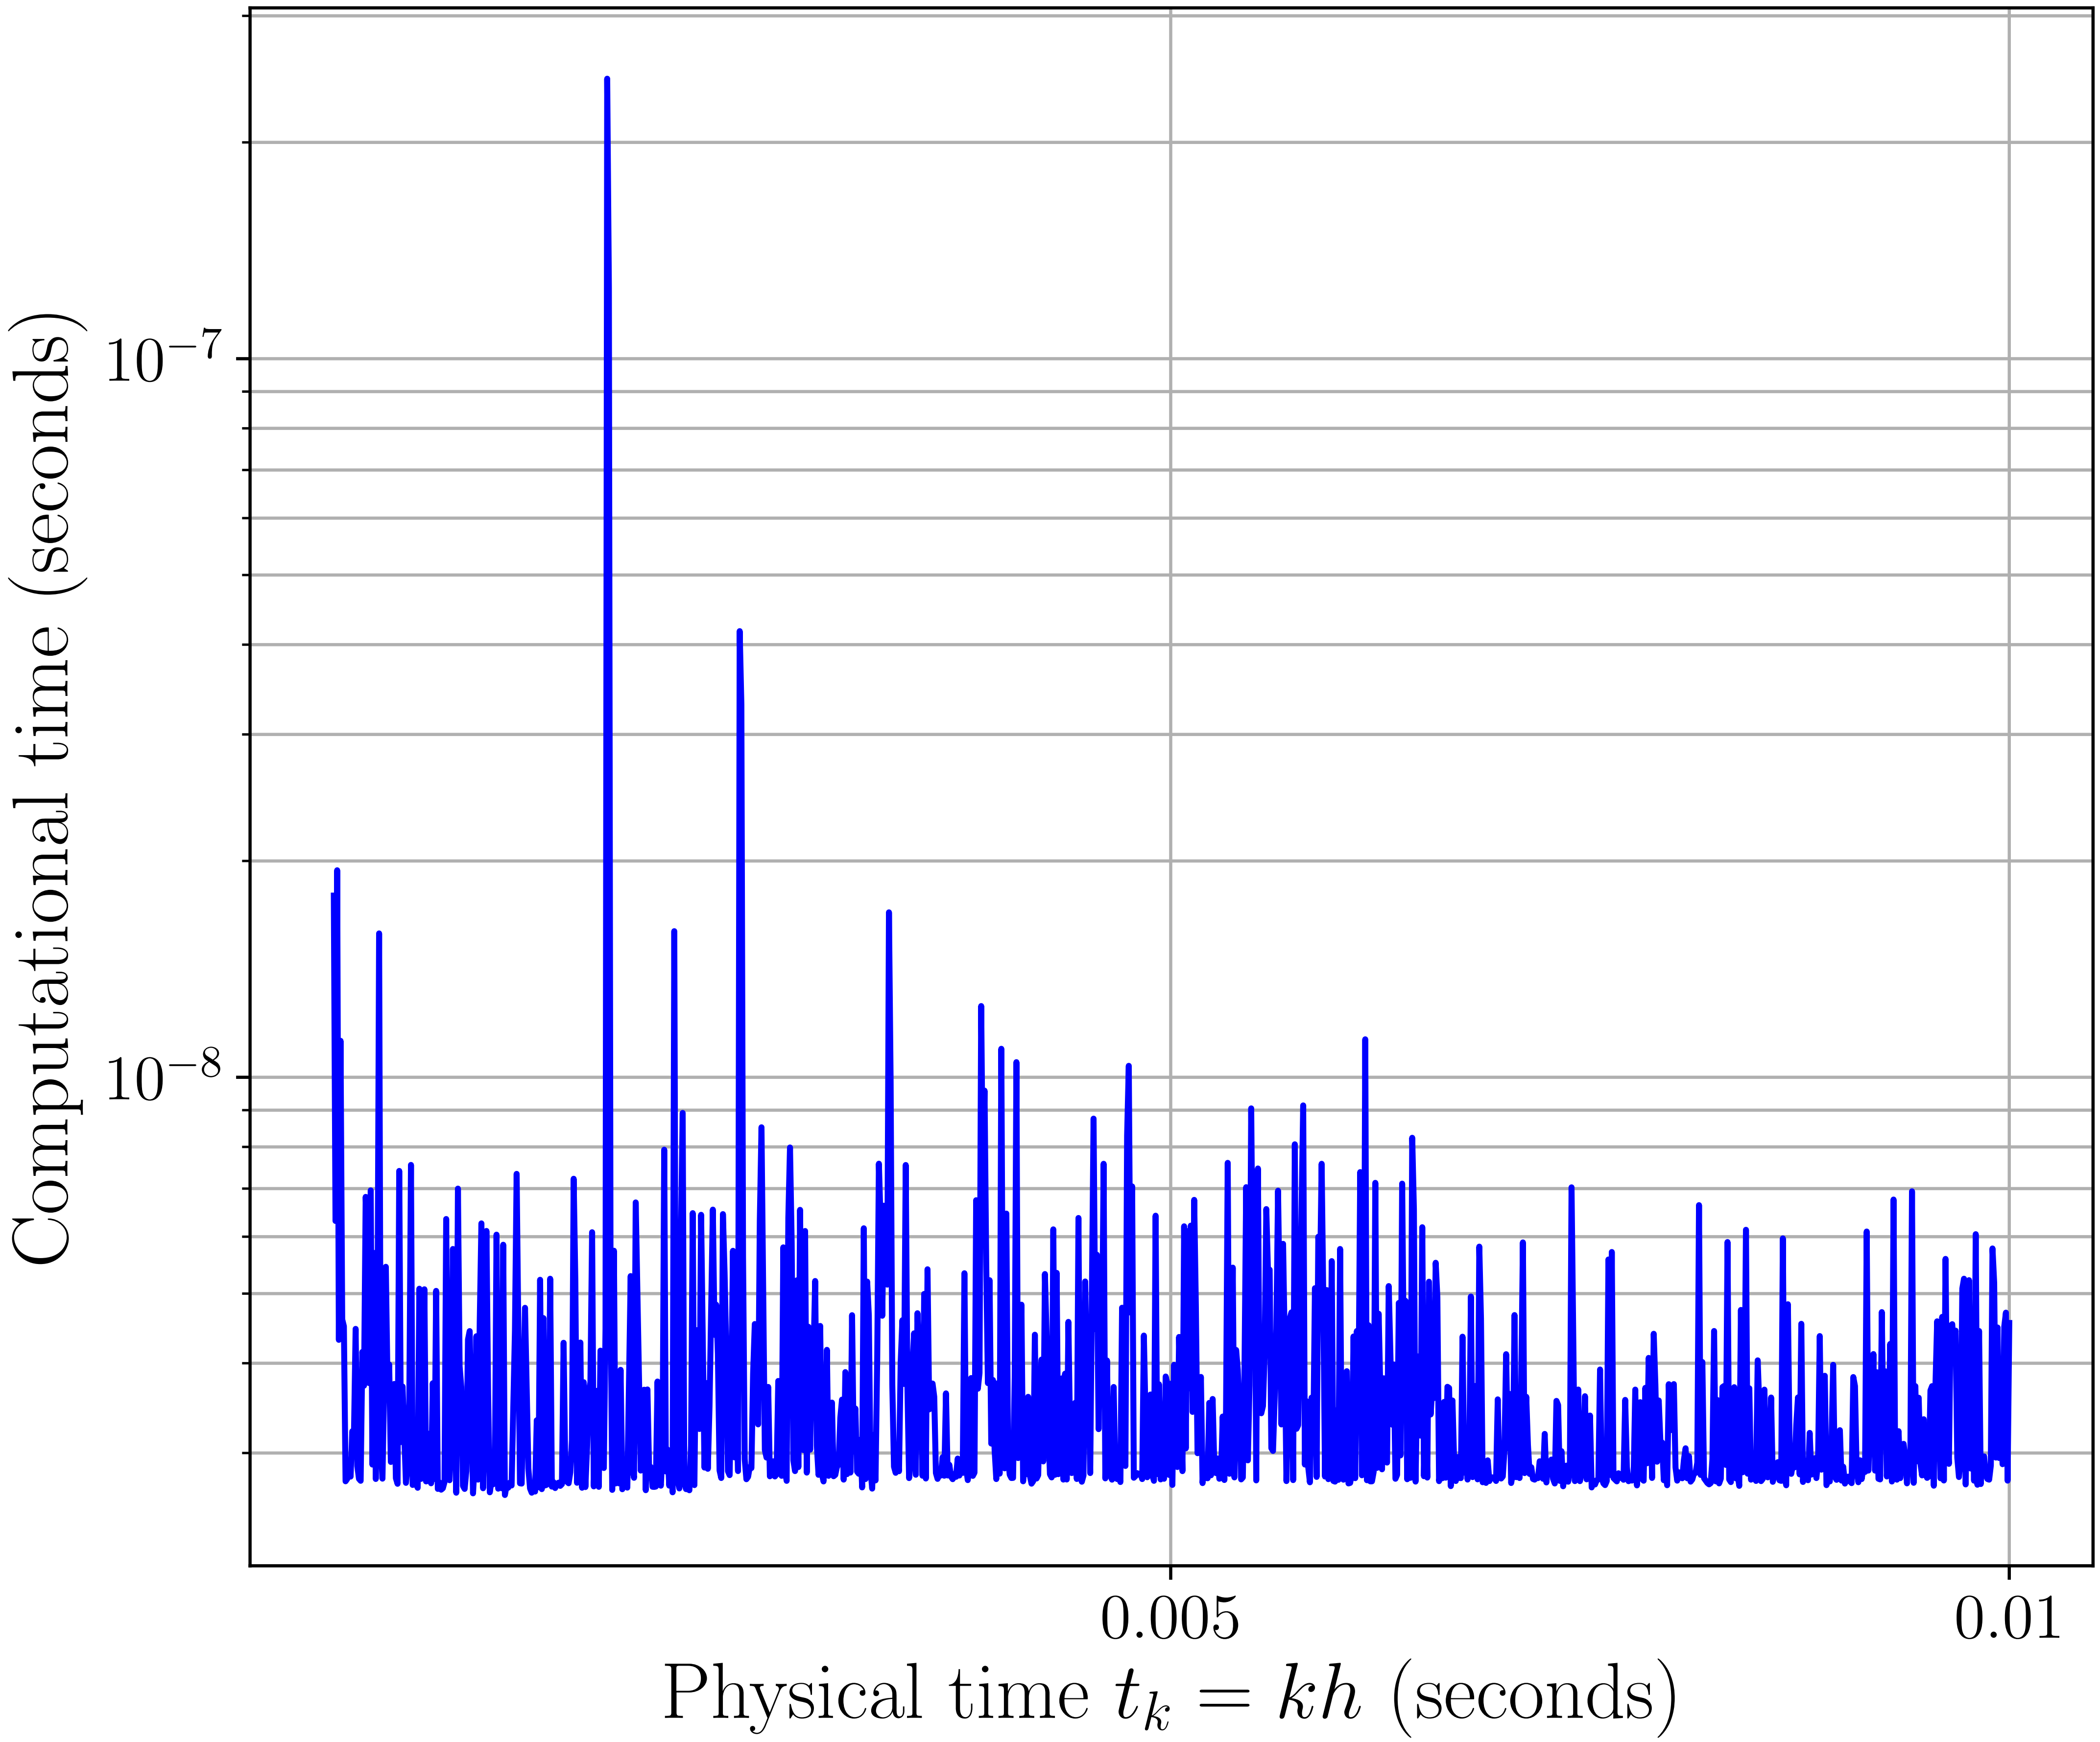

Details on Proximal Prediction

Extensions: mean-field models for nonlocal interaction, state-dependent diffusions

Publications:

— K.F. Caluya, and A.H., Proximal Recursion for Solving the Fokker-Planck Equation, *ACC 2019*.

— K.F. Caluya, and A.H., Gradient Flow Algorithms for Density Propagation in Stochastic Systems, *IEEE Trans. Automatic Control* 2020, doi: 10.1109/TAC.2019.2951348.

Git repo: github.com/kcaluya/UncertaintyPropagation

Solving density control as Wasserstein gradient flow

Finite Horizon Feedback Density Control

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

subject to

$$\mathrm{d}\boldsymbol{x} = \left\{\boldsymbol{f}\left(\boldsymbol{x},t\right) + \boldsymbol{B}(t)\boldsymbol{u}\left(\boldsymbol{x},t\right)\right\}\mathrm{d}t + \sqrt{2\epsilon}\boldsymbol{B}(t)\,\mathrm{d}\boldsymbol{w},$$

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

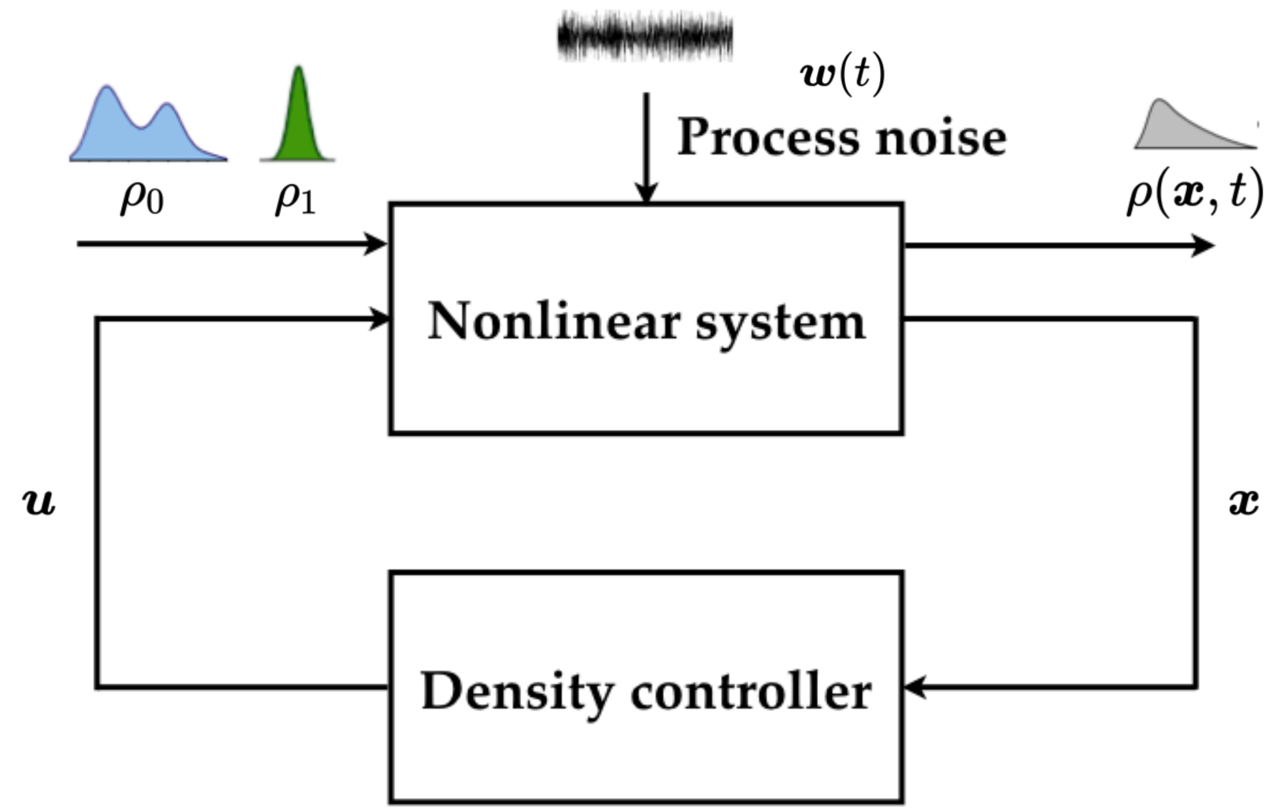

Finite Horizon Feedback Density Control

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

subject to

$$\mathrm{d}\boldsymbol{x} = \left\{\boldsymbol{f}(\boldsymbol{x},t) + \boldsymbol{B}(t)\boldsymbol{u}(\boldsymbol{x},t)\right\}\mathrm{d}t + \sqrt{2\epsilon}\boldsymbol{B}(t)\,\mathrm{d}\boldsymbol{w},$$

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

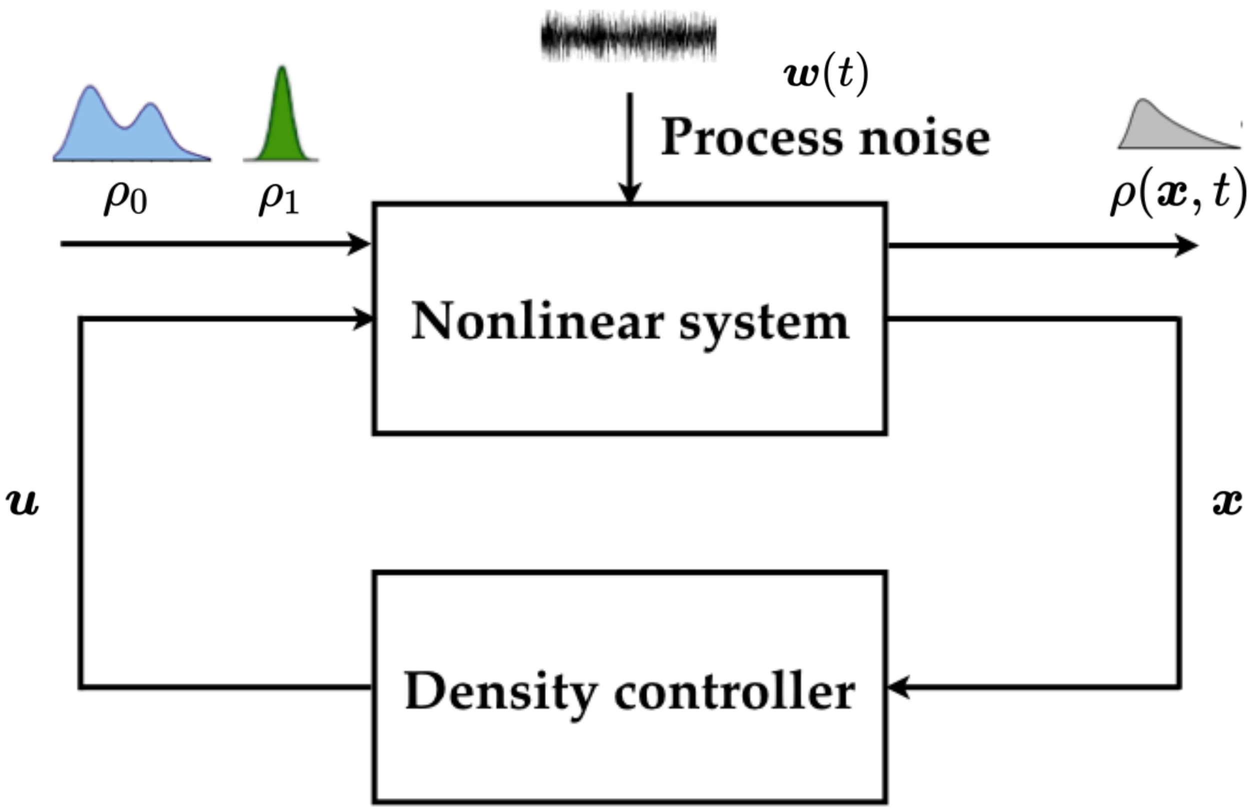

Necessary conditions for optimality: coupled nonlinear PDEs (FPK + HJB)

$$\frac{\partial \rho^{\mathrm{opt}}}{\partial t} + \nabla \cdot \left(\rho^{\mathrm{opt}}\left(\boldsymbol{f} + \boldsymbol{B}(t)^\top \nabla\psi\right)\right) = \epsilon \boldsymbol{1}^\top \left(\boldsymbol{D}(t) \odot \mathrm{Hess}\left(\rho^{\mathrm{opt}}\right)\right)\boldsymbol{1},$$

$$\frac{\partial \psi}{\partial t} + \frac{1}{2}\|\boldsymbol{B}(t)^\top \nabla\psi\|_2^2 + \langle \nabla\psi, \boldsymbol{f}\rangle = -\epsilon\langle \boldsymbol{D}(t), \mathrm{Hess}(\psi)\rangle$$

Boundary conditions:

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

Optimal control:

$$\boldsymbol{u}^{\mathrm{opt}}(\boldsymbol{x},t) = \boldsymbol{B}(t)^\top \nabla\psi$$

Feedback Synthesis via the Schrödinger System

Schrödinger's (until recently) forgotten papers:

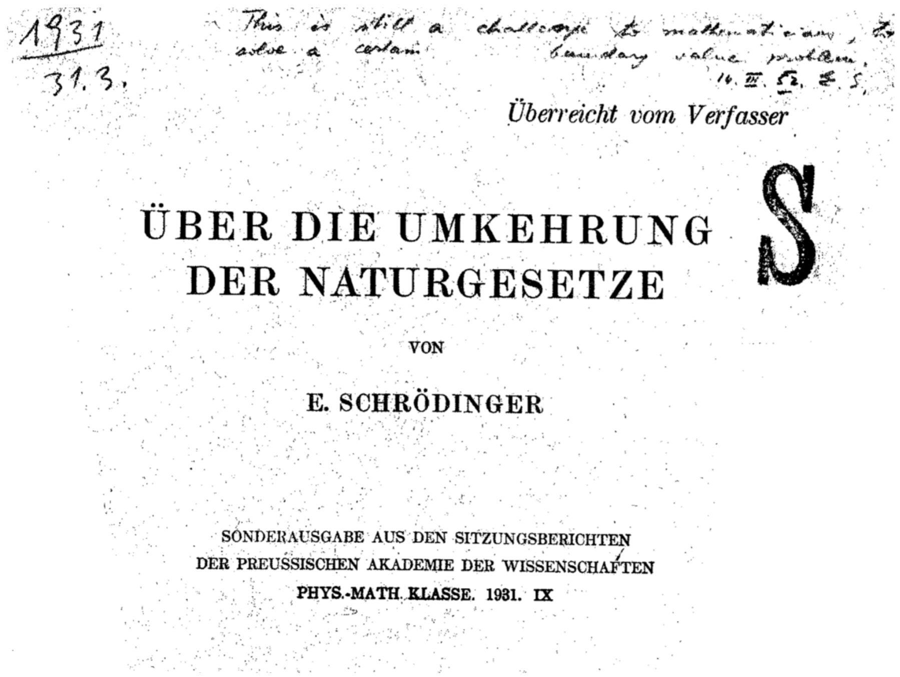
1931
31.3.

This is still a challenge to mathematicians, to solve a certain boundary value problem. 14.III.53. 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

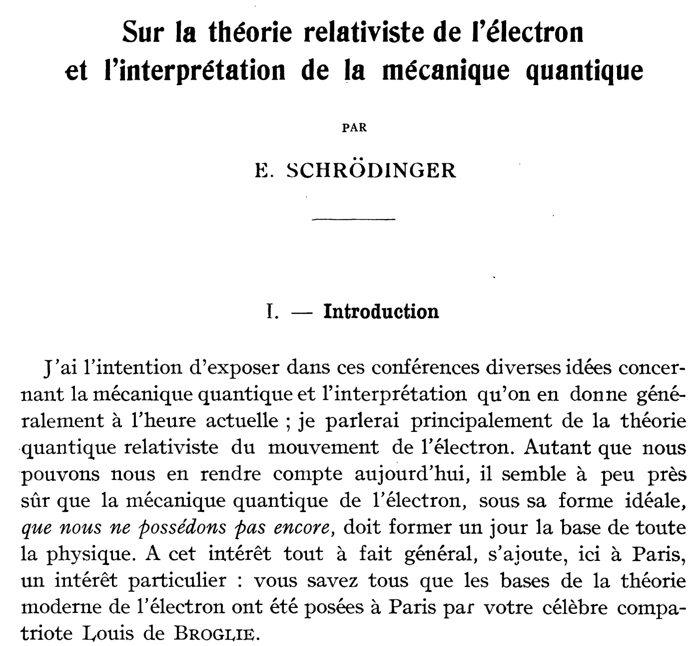
Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique

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 transform: $\left(\rho^{\text{opt}}, \psi\right) \mapsto \left(\varphi, \hat{\varphi}\right)$

$$\varphi\left(x,t\right) = \exp\left(\frac{\psi\left(x,t\right)}{2\epsilon}\right),$$

$$\hat{\varphi}\left(x,t\right) = \rho^{\text{opt}}\left(x,t\right)\exp\left(-\frac{\psi\left(x,t\right)}{2\epsilon}\right),$$

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

Optimal control: $u^{\text{opt}}(x,t) = 2\epsilon B(t)^{\top}\nabla\log\varphi\left(x,t\right)$

Feedback Synthesis via the Schrödinger System

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

$$\underbrace{\frac{\partial \hat{\varphi}}{\partial t} = -\nabla \cdot (\hat{\varphi}\boldsymbol{f}) + \epsilon \mathbf{1}^{\top}\left(\boldsymbol{D}(t) \odot \operatorname{Hess}(\hat{\varphi})\right)\mathbf{1}}_{\text{forward Kolmogorov PDE}}, \quad \varphi_0 \hat{\varphi}_0 = \rho_0,$$

$$\underbrace{\frac{\partial \varphi}{\partial t} = -\langle \nabla \varphi, \boldsymbol{f} \rangle - \epsilon \langle \boldsymbol{D}(t), \operatorname{Hess}(\varphi) \rangle}_{\text{backward Kolmogorov PDE}}, \qquad \varphi_1 \hat{\varphi}_1 = \rho_1.$$

Wasserstein proximal algorithm $\longrightarrow$ fixed point recursion over $(\hat{\varphi}_0, \varphi_1)$

(Contractive in Hilbert metric)

— Y. Chen, T.T. Georgiou, and M. Pavon, Entropic and displacement interpolation: a computational approach using the Hilbert metric, *SIAM J. Applied Mathematics*, 2016.

Feedback Density Control: Zero Prior Dynamics

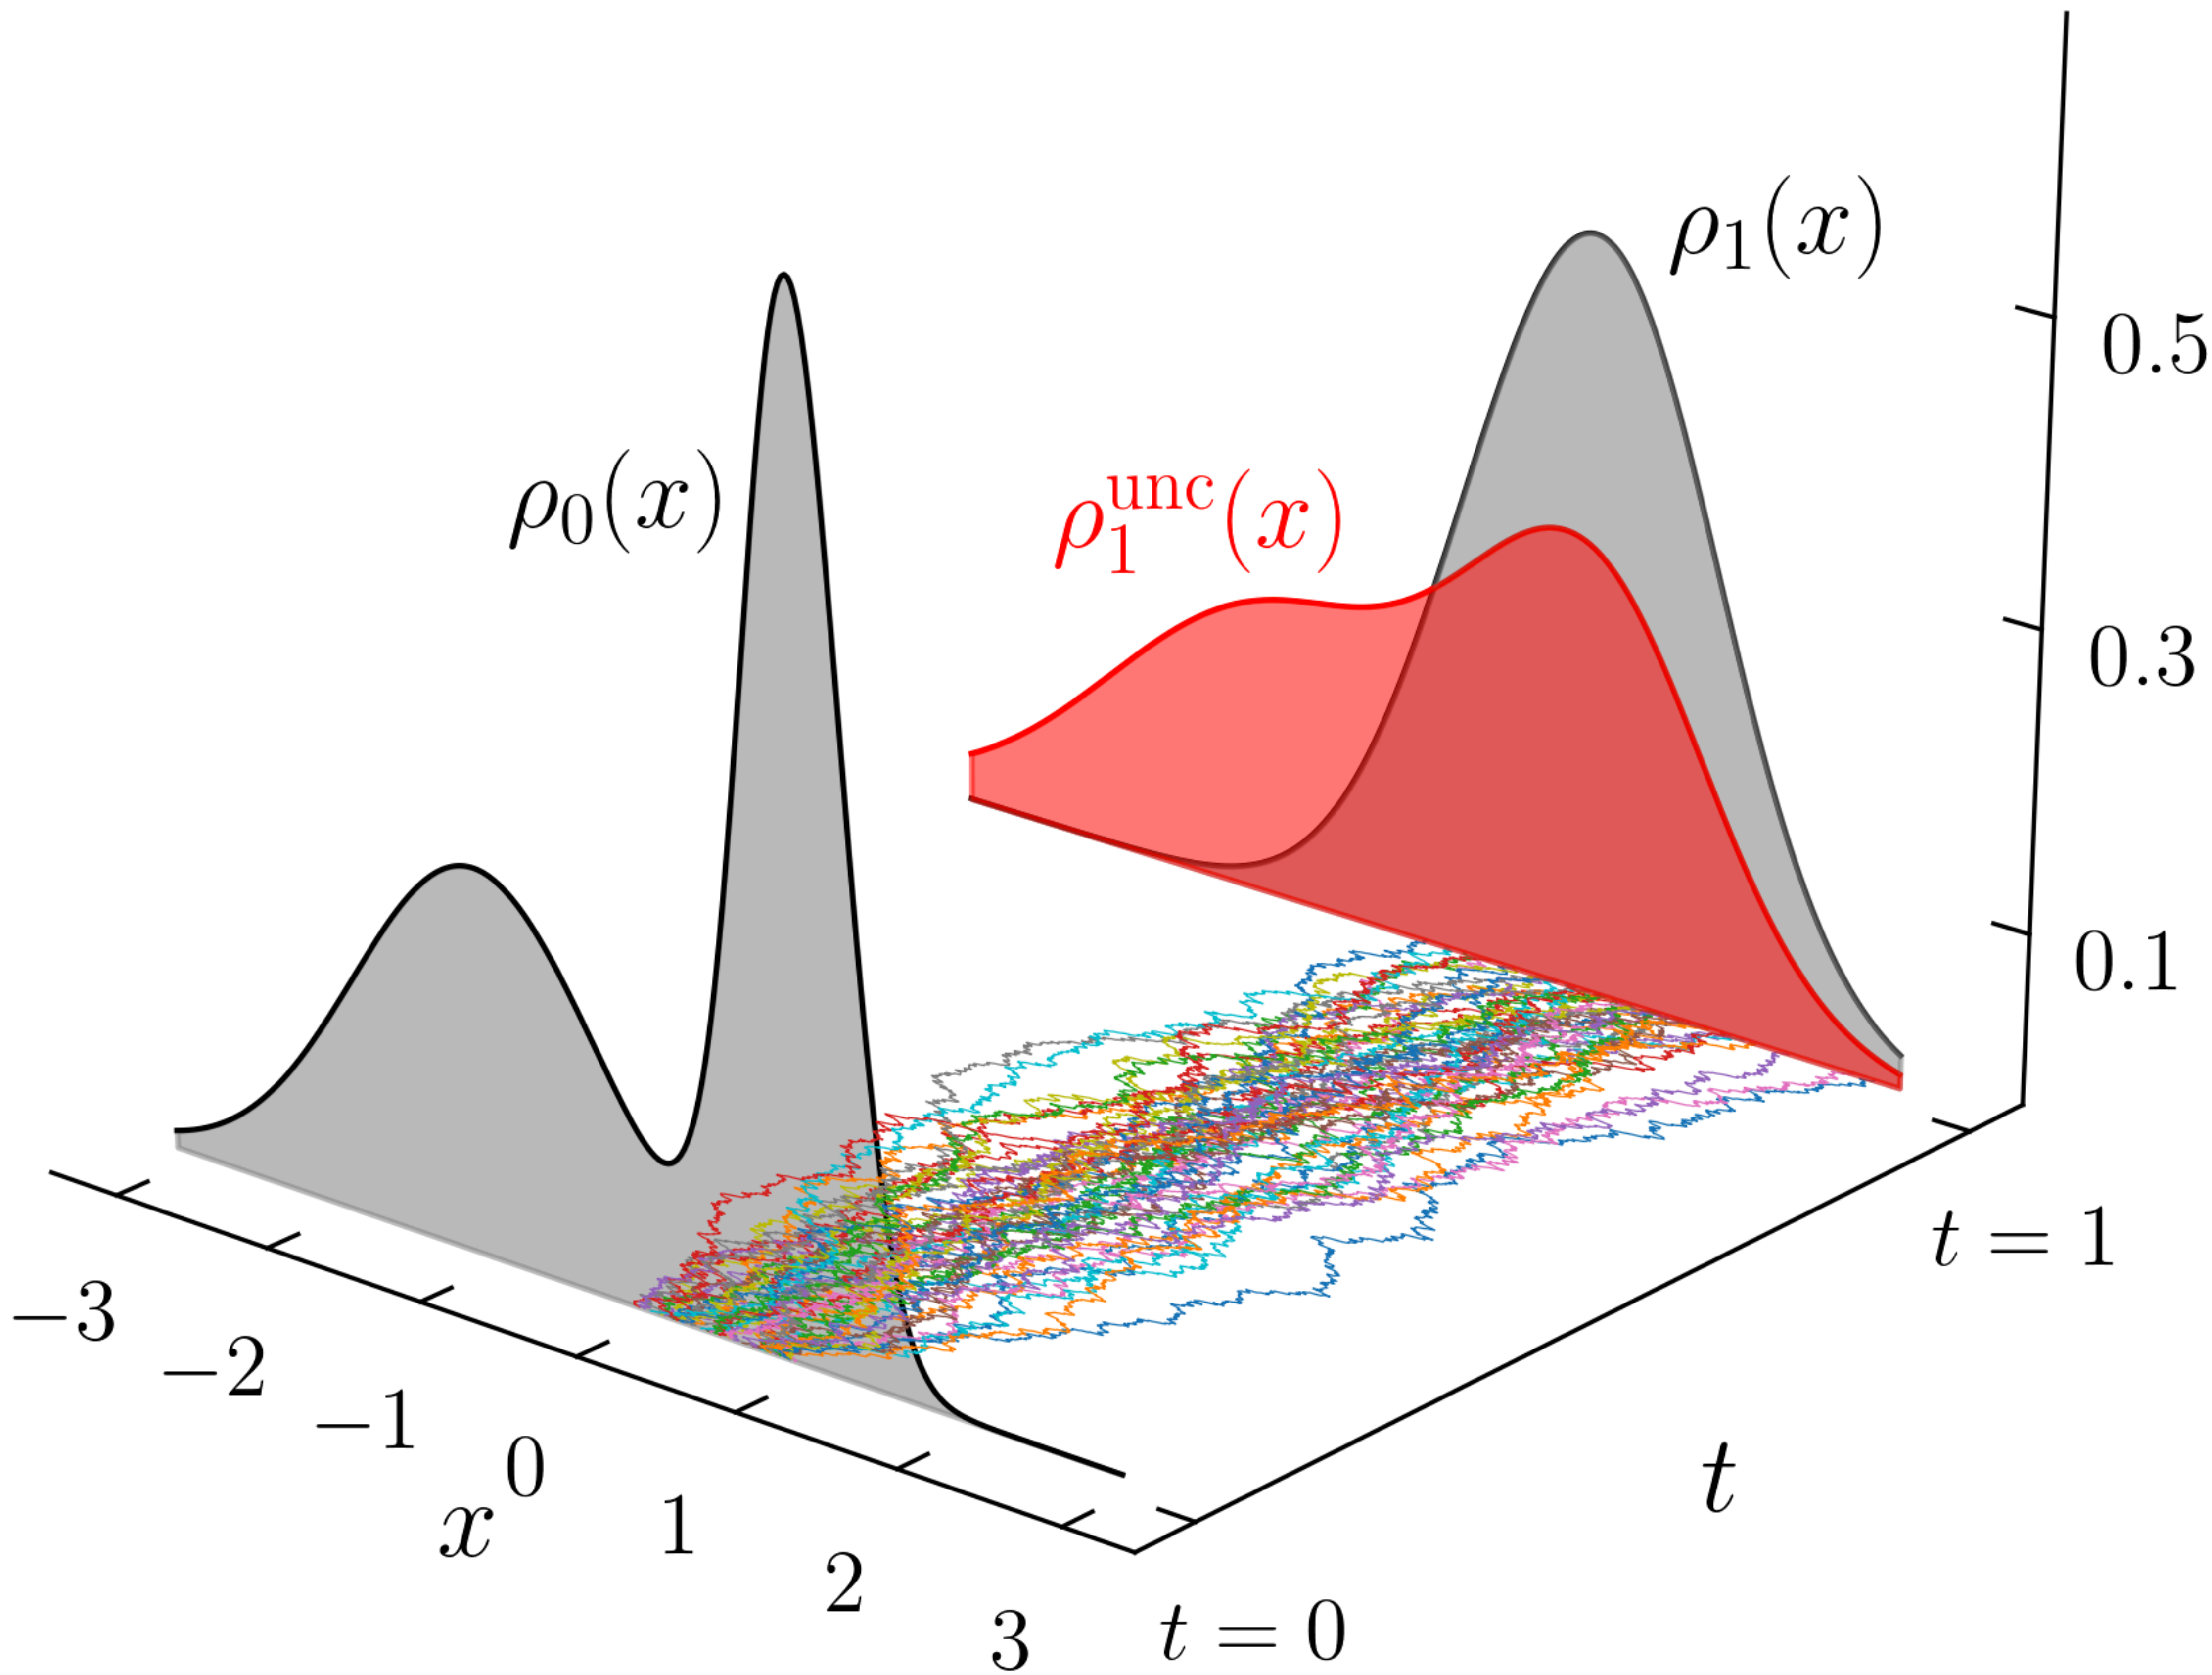

Feedback Density Control: LTI Prior Dynamics

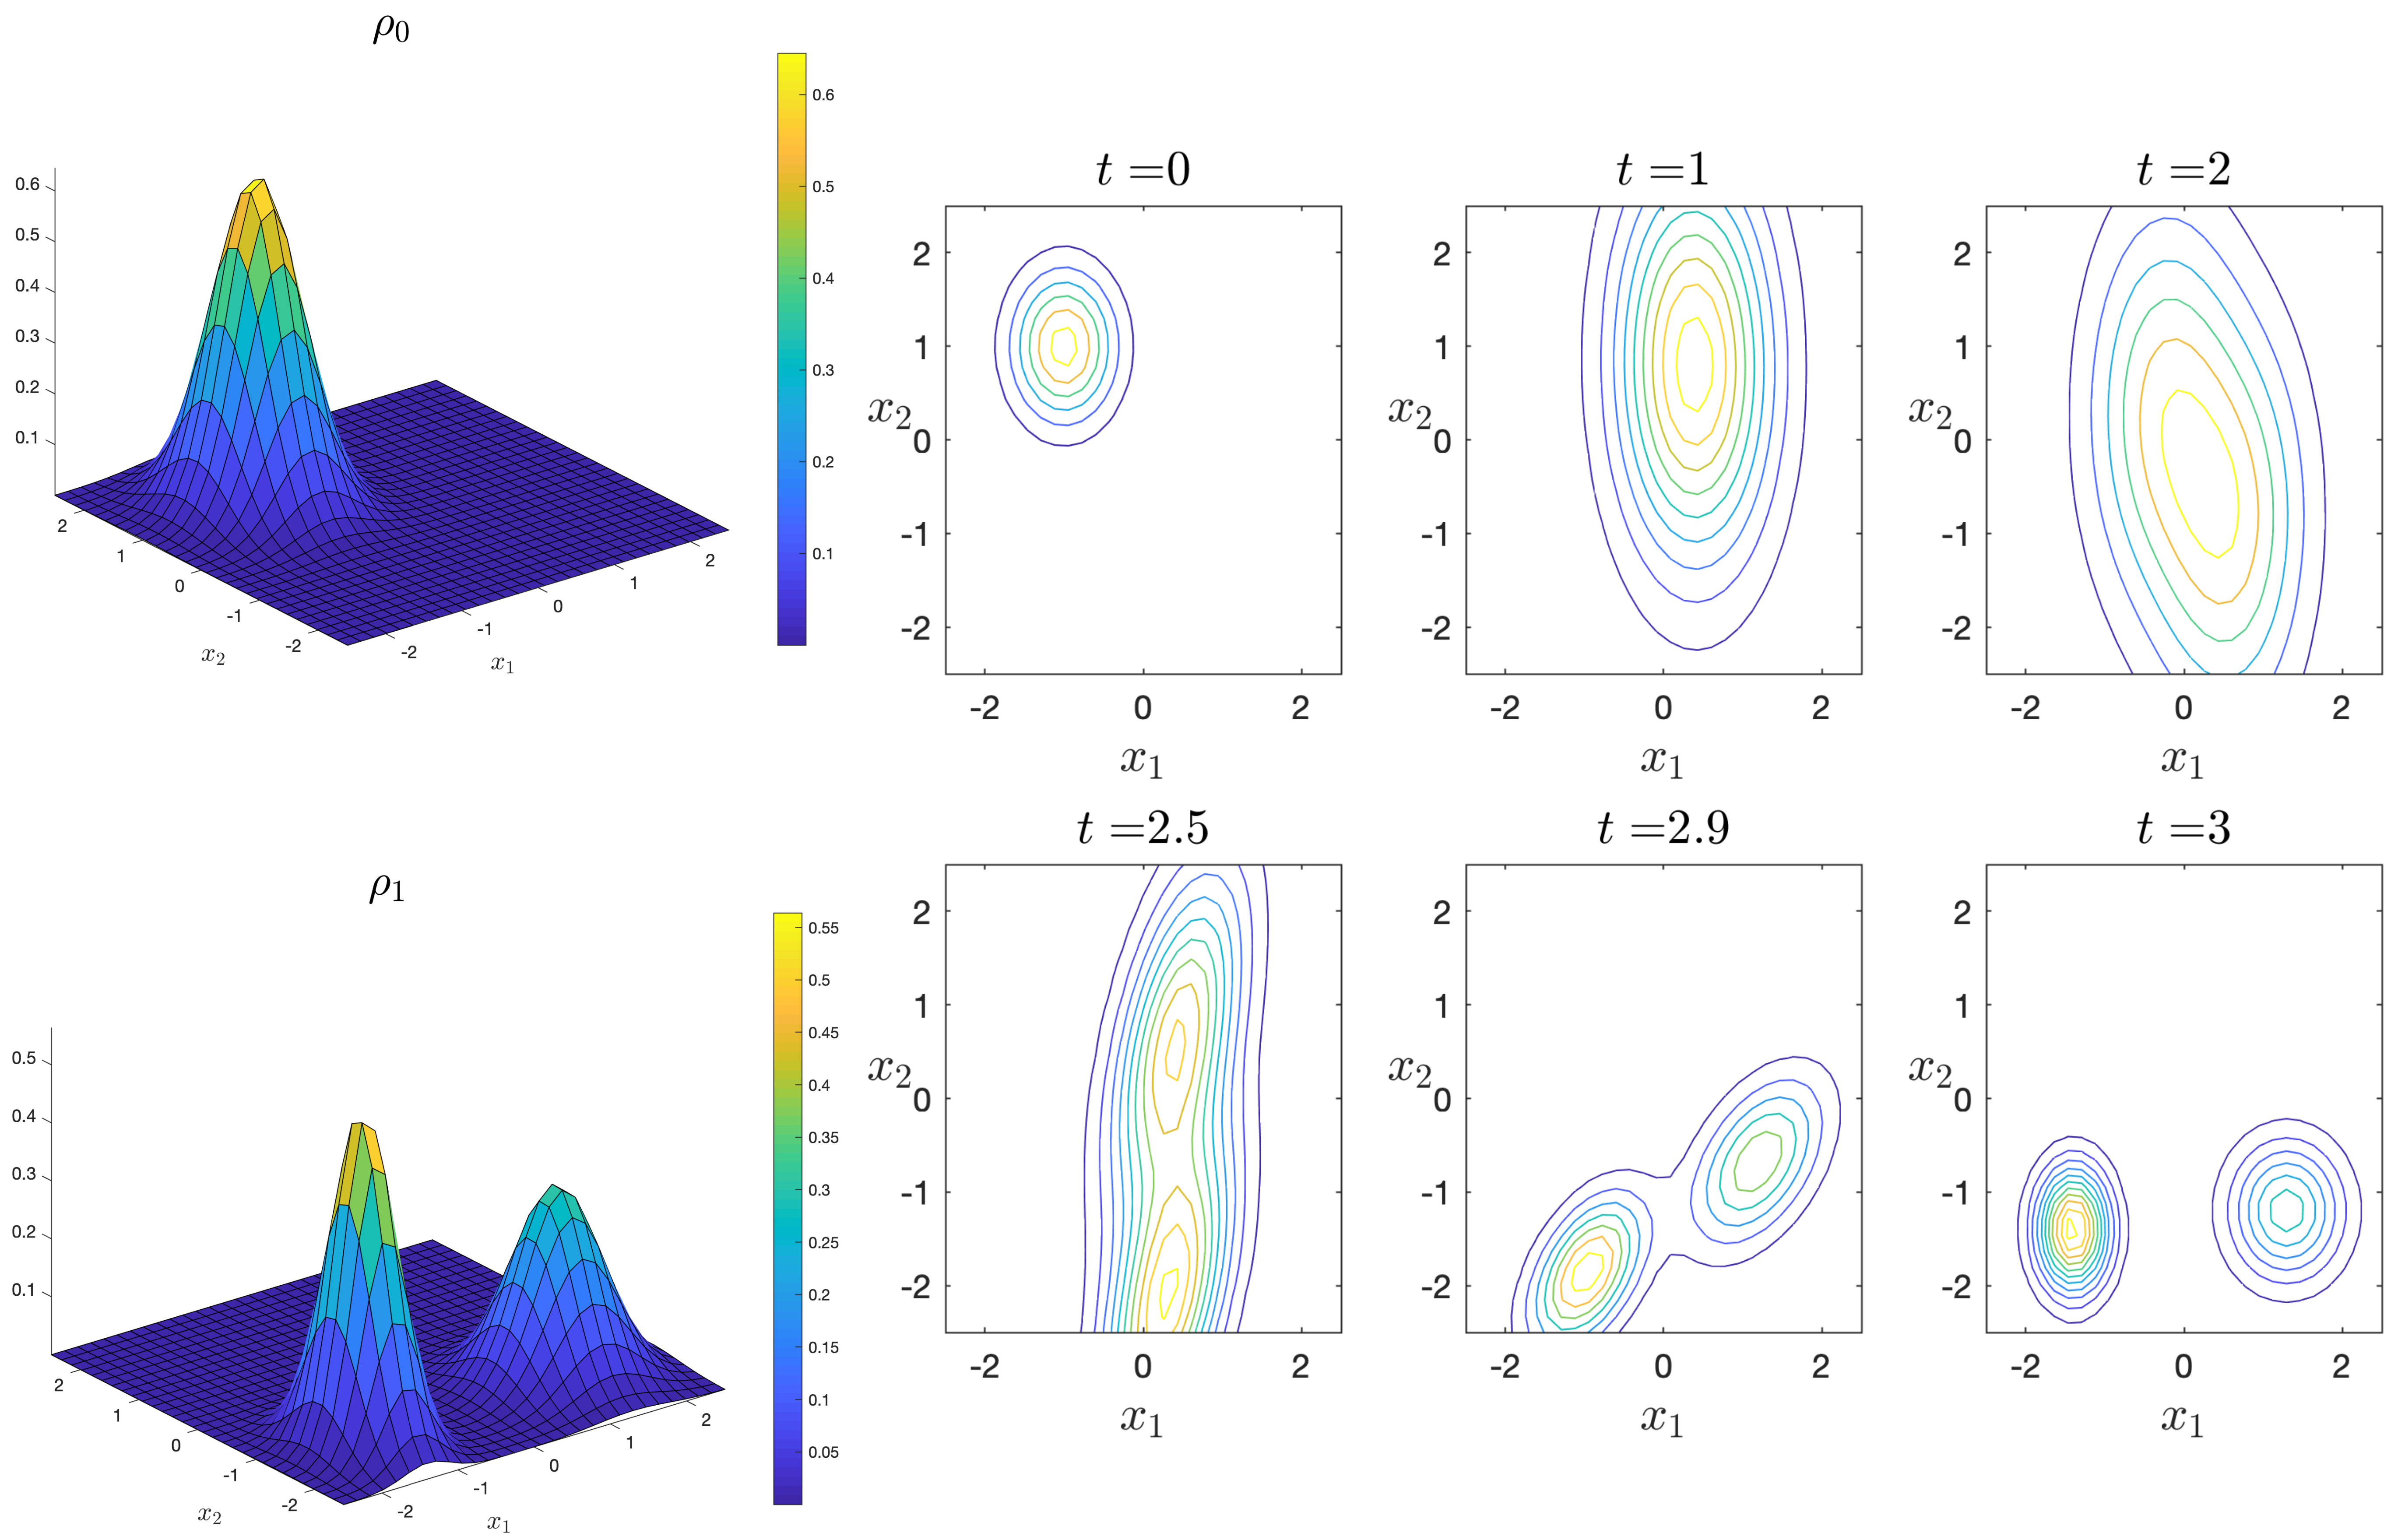

Feedback Density Control: Nonlinear Prior Dyn.

How to solve the Schrödinger System with nonlinear drift?

— No analytical handle on the transition kernel

— The backward Kolmogorov PDE cannot be written as Wasserstein gradient flow

Feedback Density Control: Nonlinear Prior Dyn.

How to solve the Schrödinger System with nonlinear drift?

— No analytical handle on the transition kernel

— The backward Kolmogorov PDE cannot be written as Wasserstein gradient flow

Can we exploit *some* structural nonlinearities in practice?

Gradient drift:

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

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

Mixed conservative -dissipative drift:

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

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

Feedback Density Control: Gradient Drift

Theorem

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

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

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

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

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

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

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

Feedback Density Control: Mixed Conservative-Dissipative Drift

Theorem

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

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

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

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

$$\frac{\partial\hat{\varphi}}{\partial t} = -\langle\boldsymbol{\eta},\nabla_{\boldsymbol{\xi}}\hat{\varphi}\rangle + \nabla_{\boldsymbol{\eta}}\cdot\left(\hat{\varphi}\left(\nabla_{\boldsymbol{\xi}}V\left(\boldsymbol{\xi}\right)+\kappa\boldsymbol{\eta}\right)\right) + \epsilon\kappa\Delta_{\boldsymbol{\eta}}\hat{\varphi},$$
$$\frac{\partial p}{\partial s} = -\langle\boldsymbol{\vartheta},\nabla_{\boldsymbol{\xi}}p\rangle + \nabla_{\boldsymbol{\vartheta}}\cdot\left(p\left(\nabla_{\boldsymbol{\xi}}V\left(\boldsymbol{\xi}\right)+\kappa\boldsymbol{\vartheta}\right)\right) + \epsilon\kappa\,\Delta_{\boldsymbol{\vartheta}}p,$$
$$\hat{\varphi}\left(\boldsymbol{\xi},\boldsymbol{\eta},0\right) = \hat{\varphi}_0(\boldsymbol{\xi},\boldsymbol{\eta}),$$
$$p(\boldsymbol{\xi},\boldsymbol{\vartheta},0) = \varphi_1(\boldsymbol{\xi},-\boldsymbol{\vartheta})\exp\left(-\tfrac{1}{\epsilon}\left(\tfrac{1}{2}\|\boldsymbol{\vartheta}\|_2^2 + V(\boldsymbol{\xi})\right)\right).$$

Feedback Density Control via Wasserstein prox.

Design proximal recursions over discrete time pair:

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

The recursions are of the form:

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

Consistency guarantees:

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

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

Details:

— K.F. Caluya, and A.H., Wasserstein Proximal Algorithms for the Schrödinger Bridge Problem: Density Control with Nonlinear Drift, *arXiv 1912.01244*.

Feedback Density Control: Gradient Drift

Uncontrolled joint PDF evolution:

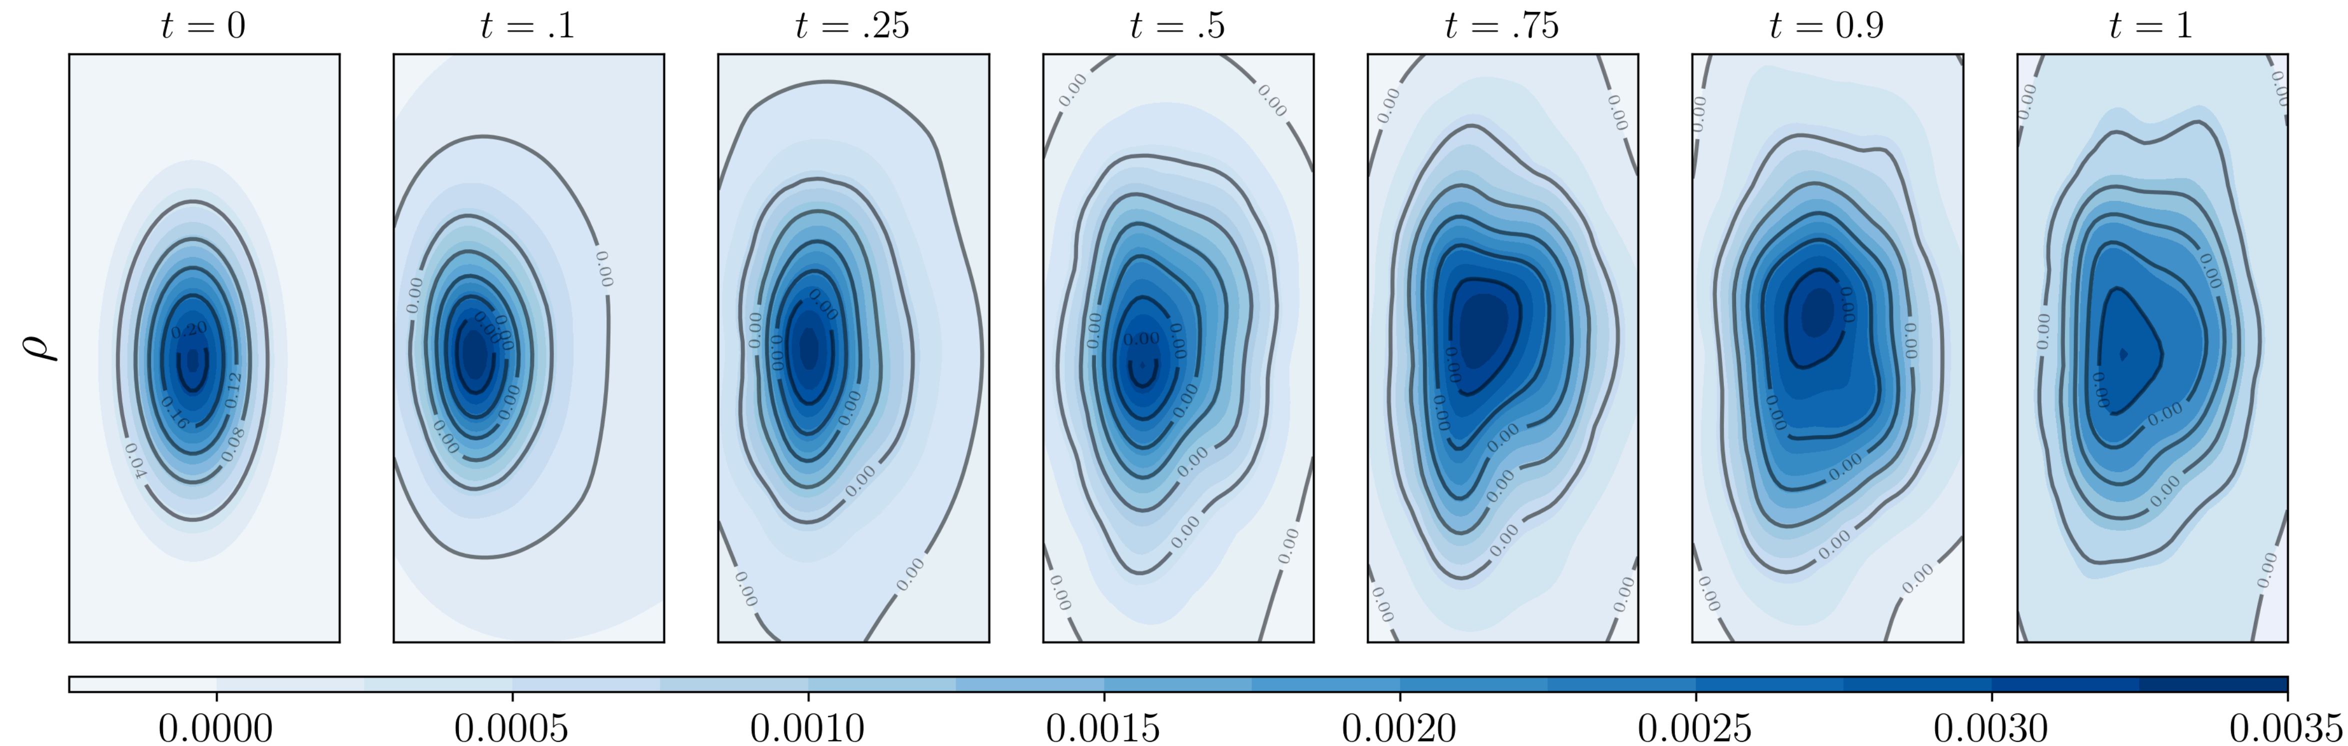

Optimal controlled joint PDF evolution:

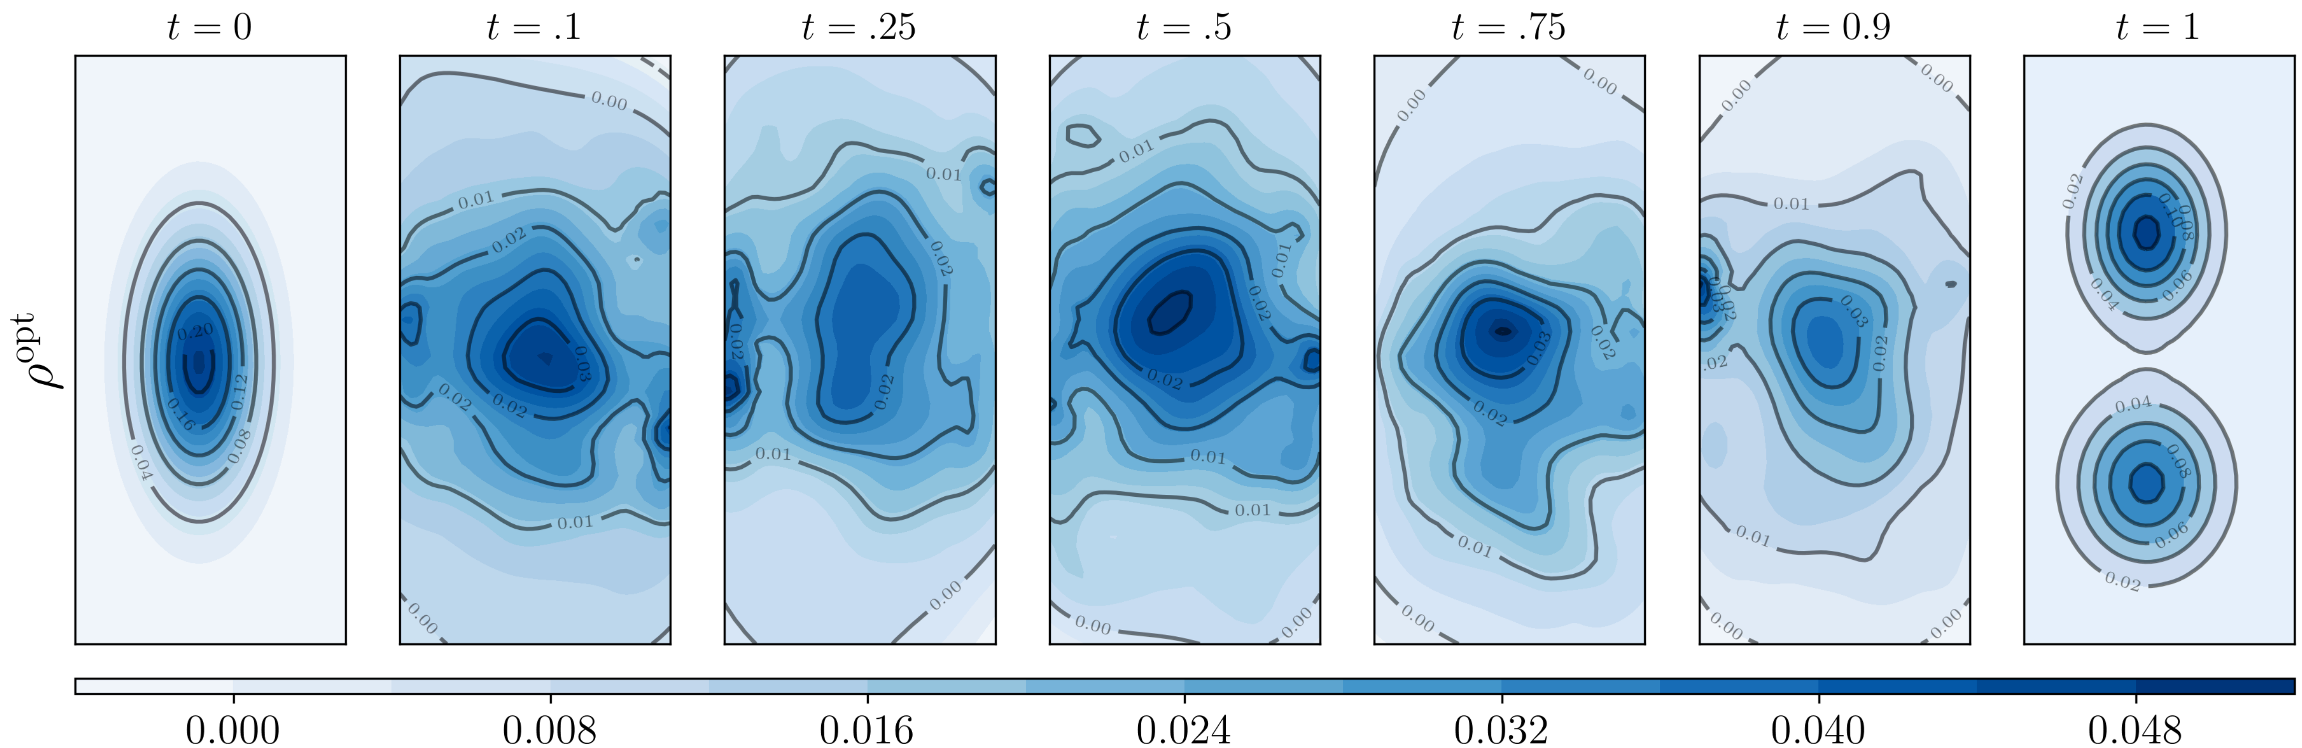

Feedback Density Control: Mixed Conservative-Dissipative Drift

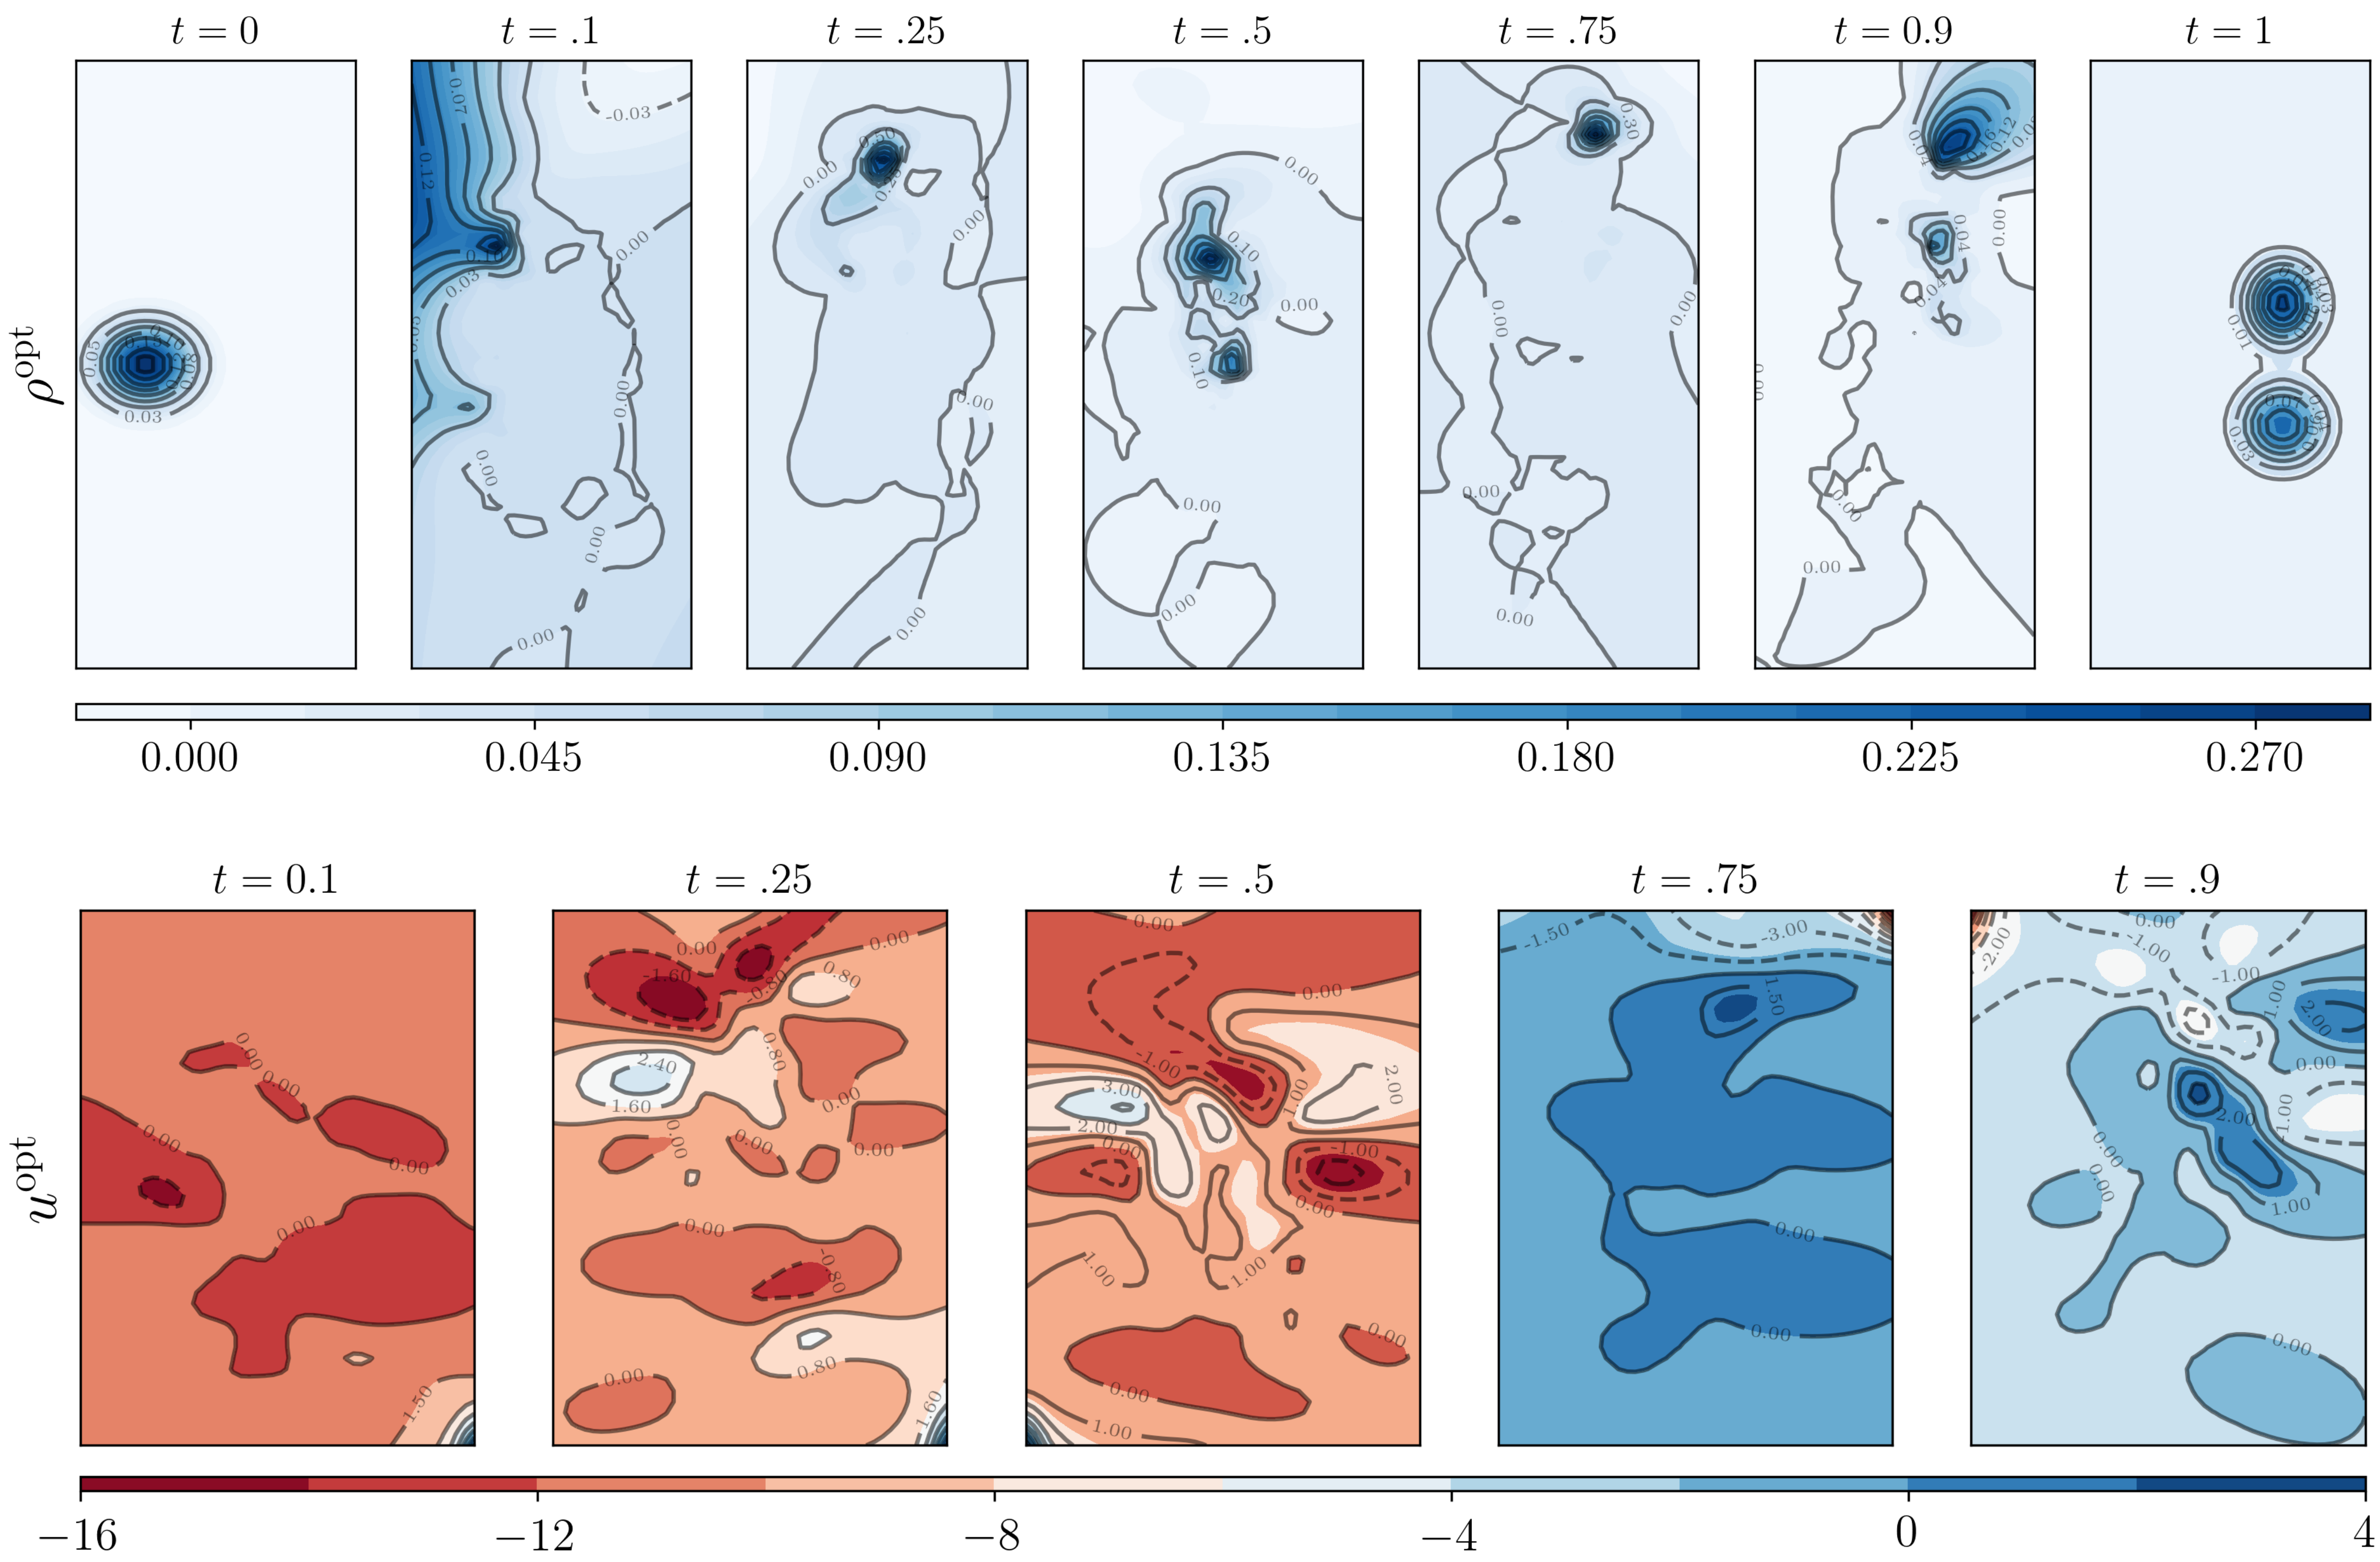

Density Control with Det. Path Constraints

Reflecting Schrödinger Bridge

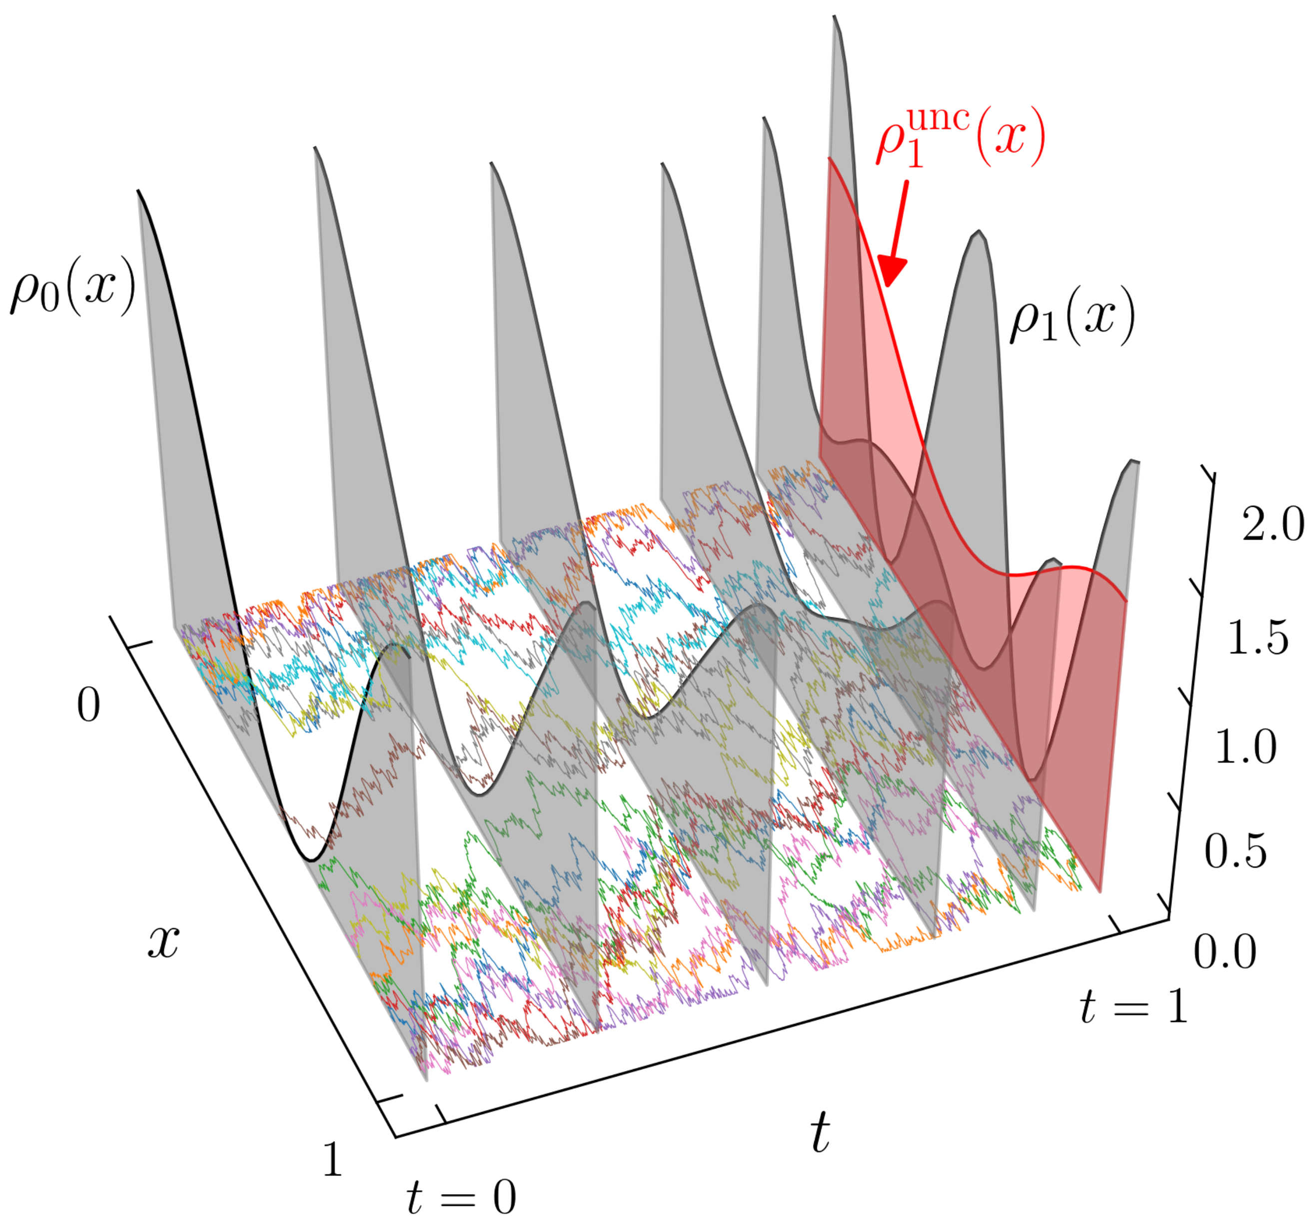

Contraction in the Hilbert metric

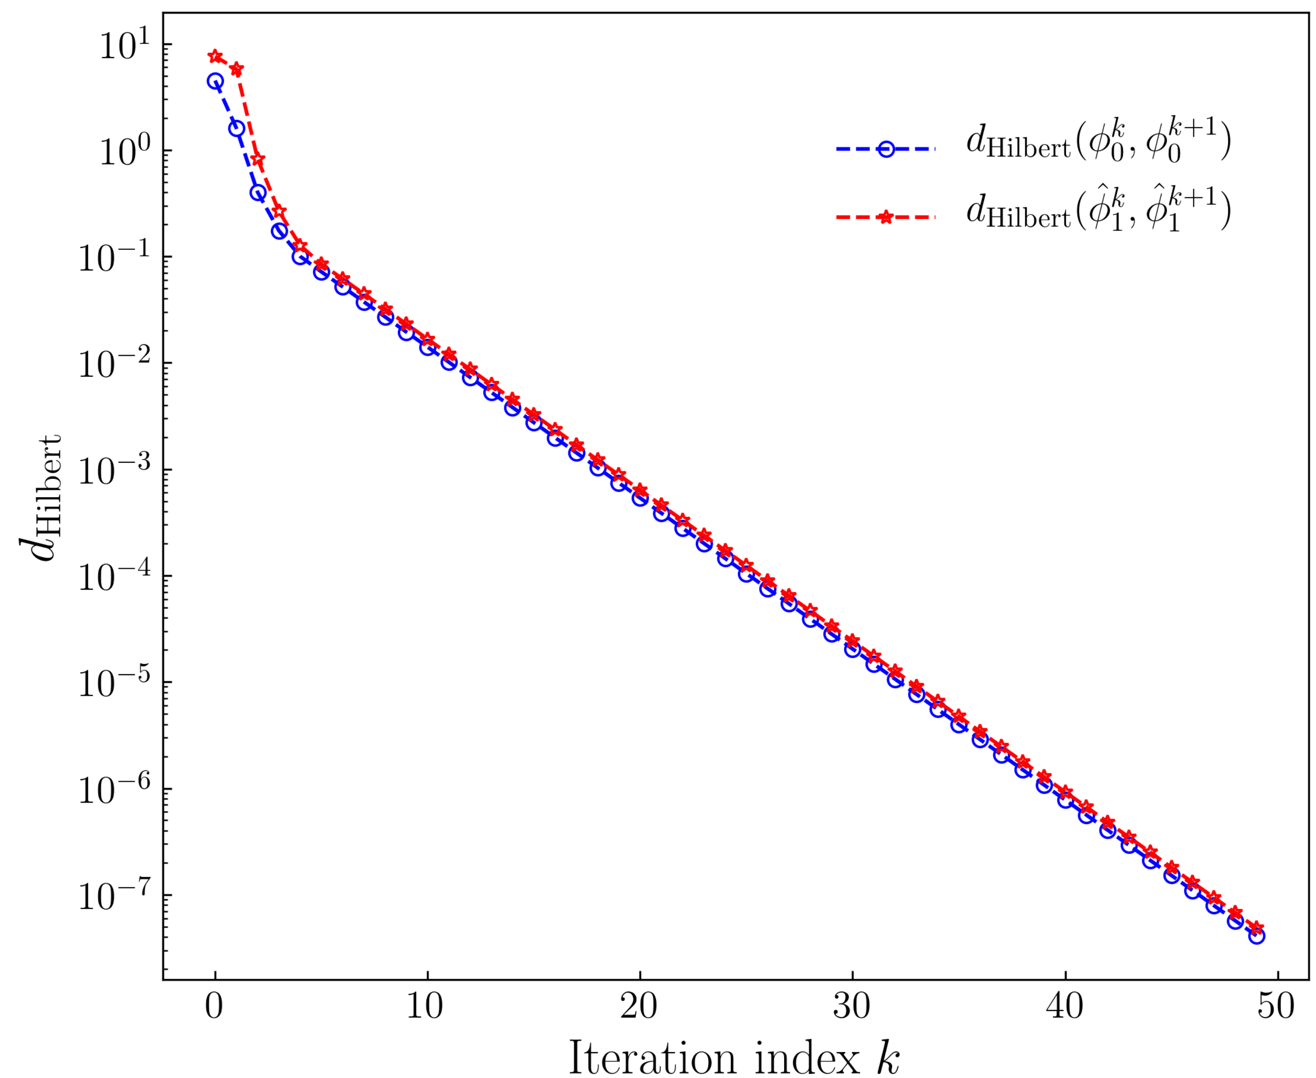

*** Ongoing work**

Density Control with Feedback Linearizable Dyn.

Setting:

For $\boldsymbol{x} \in \mathbb{R}^n$, $\boldsymbol{u} \in \mathbb{R}^m$, and given ρ_0, ρ_1, consider

$$
\begin{aligned}
\inf_{\boldsymbol{u} \in \mathcal{U}} \quad & \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2 \, \mathrm{d}t\right\}, \\
\text{subject to} \quad & \dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{u}, \\
& \boldsymbol{x}(t=0) \sim \rho_0(\boldsymbol{x}) \quad \boldsymbol{x}(t=1) \sim \rho_1(\boldsymbol{x}),
\end{aligned}
$$

with $(\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{G}(\boldsymbol{x}))$ feedback linearizable, i.e., there exists a triple $(\boldsymbol{\delta}(\boldsymbol{x}), \boldsymbol{\Gamma}(\boldsymbol{x}), \boldsymbol{\tau}(\boldsymbol{x}))$ such that

$$
\begin{aligned}
\left(\nabla \boldsymbol{\tau}\left(\boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{\delta}(\boldsymbol{x})\right)\right)_{\boldsymbol{x}=\boldsymbol{\tau}^{-1}(\boldsymbol{z})} &= \boldsymbol{A}\boldsymbol{z}, \\
\left(\nabla \boldsymbol{\tau}\left(\boldsymbol{G}(\boldsymbol{x})\boldsymbol{\Gamma}(\boldsymbol{x})\right)\right)_{\boldsymbol{x}=\boldsymbol{\tau}^{-1}(\boldsymbol{z})} &= \boldsymbol{B},
\end{aligned}
$$

where $(\boldsymbol{A}, \boldsymbol{B})$ is controllable. So, $(\boldsymbol{x}, \boldsymbol{u}) \mapsto (\boldsymbol{z}, \boldsymbol{v})$ with

$$
\dot{\boldsymbol{z}} = \boldsymbol{A}\boldsymbol{z} + \boldsymbol{B}\boldsymbol{v}, \quad \boldsymbol{u} = \boldsymbol{\delta}(\boldsymbol{x}) + \boldsymbol{\Gamma}(\boldsymbol{x})\boldsymbol{v}.
$$

Density Control with Feedback Linearizable Dyn.

Main idea:

Push-forward the endpoint PDFs via diffeomorphism $\boldsymbol{\tau} : \mathcal{X} \mapsto \mathcal{Z}$

$$\sigma_i(\boldsymbol{z}) := \boldsymbol{\tau}_\sharp \rho_i = \frac{\rho_i(\boldsymbol{\tau}^{-1}(\boldsymbol{z}))}{|\det(\nabla_{\boldsymbol{x}} \boldsymbol{\tau}_{\boldsymbol{x}=\boldsymbol{\tau}^{-1}(\boldsymbol{z})})|}, \quad i \in \{0, 1\}.$$

Define maps $\boldsymbol{\delta}_{\boldsymbol{\tau}} := \boldsymbol{\delta} \circ \boldsymbol{\tau}^{-1}, \boldsymbol{\Gamma}_{\boldsymbol{\tau}} := \boldsymbol{\Gamma} \circ \boldsymbol{\tau}^{-1}$

Rewrite the problem in feedback linearized coordinates as

$$\begin{aligned}
\underset{\sigma, \boldsymbol{v}}{\text{minimize}} \quad & \int_0^1 \int_{\mathcal{Z}} \frac{1}{2} \mathcal{L}(\boldsymbol{z}, \boldsymbol{v}) \sigma(\boldsymbol{z}, t) \, \mathrm{d}\boldsymbol{z}\mathrm{d}t, \\
\text{subject to} \quad & \frac{\partial \sigma}{\partial t} + \nabla_{\boldsymbol{z}} \cdot ((\boldsymbol{A}\boldsymbol{z} + \boldsymbol{B}\boldsymbol{v})\sigma) = 0 \\
& \sigma(\boldsymbol{z}, t = 0) = \sigma_0, \quad \sigma(\boldsymbol{z}, t = 1) = \sigma_1,
\end{aligned}$$

where $\mathcal{L}(\boldsymbol{z}, \boldsymbol{v}) := \|\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) + \boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z})\boldsymbol{v}\|_2^2.$

Density Control with Feedback Linearizable Dyn.

Optimality:

Optimal control: $v^{\mathrm{opt}}(z,t) = (\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{\top}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}(z))^{-1}\boldsymbol{B}^{\top}\nabla_{z}\psi - \boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1}(z)\delta_{\boldsymbol{\tau}}(z)$

HJB:

$$\frac{\partial\psi}{\partial t} + \langle\nabla_{z}\psi, \boldsymbol{A}z\rangle - \langle\nabla_{z}\psi, \boldsymbol{B}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}{}^{-1}(z)\delta_{\boldsymbol{\tau}}(z)\rangle + \frac{1}{2}\langle\nabla_{z}\psi, \boldsymbol{B}\left(\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{\top}(z)\boldsymbol{\Gamma}_{\boldsymbol{\tau}}(z)\right)^{-1}\boldsymbol{B}^{\top}\nabla_{z}\psi\rangle = 0.$$

Solve by dynamic stochastic regularization $\rightsquigarrow$ SBP $\rightsquigarrow$ fixed point recursion

Details:

— K.F. Caluya, and A.H., Finite Horizon Density Control for Static State Feedback Linearizable Systems, *arXiv 1904.02272*.

— K.F. Caluya, and A.H., Finite Horizon Density Steering for Multi-input State Feedback Linearizable Systems, *arXiv 1909.12511*.

Take Home Message

Emerging system-control theory for densities

Wasserstein gradient flow: one unifying framework for the prediction, estimation, and feedback control

Feedback density control theory: many recent progress, much remains to be done

Several applications: controlling biological and robotic swarm, process control

Thank You

Support:

CITRIS PEOPLE AND ROBOTS