UNIVERSITY OF CALIFORNIA

SANTA CRUZ

GENERALIZED GRADIENT FLOWS FOR DENSITY PREDICTION, CONTROL AND LEARNING

A dissertation submitted in partial satisfaction of the

requirements for the degree of

DOCTOR OF PHILOSOPHY

in

APPLIED MATHEMATICS

by

Kenneth Caluya

December 2022

The Dissertation of Kenneth Caluya
is approved:

Qi Gong, Chair

Abhishek Halder

Hongyun Wang

Peter F. Biehl
Vice Provost and Dean of Graduate Studies

Contents

List of Figures

Abstract

Generalized Gradient Flows for Density Prediction, Control and Learning

by

Kenneth Caluya

The need to predict, estimate and control density functions arise across engineering applications such as controlling biological and robotic swarms, vehicle guidance-control in uncertain dynamic environments, forecasting and demand response of loads in power systems, and active shaping of chemical concentrations in process control. Notwithstanding this recurring theme in practical applications, there does not exist a systems-control theory of densities. The perspective of this work is to close this gap by developing the theory and algorithms for prediction and control of densities subject to trajectory-level stochastic nonlinear dynamics.

We present theory and algorithms that leverage an emerging geometric interpretation of the equations of density propagation and steering. The governing partial differential equations for density propagation can be viewed as gradient flow of certain Lyapunov functionals with respect to the Wasserstein metric arising from the theory of optimal transport. This metric induces a Riemannian-like geometric structure on the infnite dimensional manifold of joint probability density functions (PDFs) supported on the state space. We leverage this geometric structure to design weighted scattered point cloud-based gradient descent algorithms via recursive evaluation of infinite dimensional proximal operators on the manifold of joint state PDFs. The resulting numerical algorithms avoid function approximation or spatial discretization, and enjoy fast computational speed due to certain conic

contraction property that we establish. We provide several numerical examples to elucidate our algorithms.

We show that the Wasserstein proximal recursions can also be leveraged to solve the minimum energy finite horizon density steering, also known as the Schrödinger Bridge Problem (SBP), which allows density regulation via feedback synthesis. This is a problem of minimum effort steering of a given joint state PDF to another over a finite time horizon, subject to a controlled stochastic differential evolution of the state vector. The same theory also arises in the study of mean-field dynamics of neural networks. We leverage the same theory to study second-order algorithms to prove their consistency and global convergence.

Notation

$\mathbb{N}$	Set of Natural Numbers
$\mathbb{R}$	Set of Real Numbers
$\mathbb{R}^n$	Euclidean Space
$\langle \cdot, \cdot \rangle$	Euclidean inner product
$\langle \boldsymbol{A}, \boldsymbol{B} \rangle$	Frobenius inner product between matrices
∇	Gradient Operator
$\nabla \cdot$	Gradient Operator
Δ	Laplacian operator
$\mathbf{Hess}(\cdot)$,	Hessian
$\delta(\boldsymbol{x} - \boldsymbol{y})$	Dirac delta located at $\boldsymbol{y}$
$\mathcal{P}_2(\mathbb{R}^n)$,	all joint PDFs supported on $\mathbb{R}^n$ with finite second moments
$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	Gaussian PDF with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$
$\mathbb{E}_\rho[\cdot]$	Expectation operator w.r.t. the PDF ρ
$\log(\cdot)$	Element-wise log
$\exp(\cdot)$	Element-wise exponential
$\odot$	Element-wise (Hadamard) product
$\oslash$	Element-wise (Hadamard) division
$\boldsymbol{I}_n$	$n \times n$ identity matrix
$\mathbf{1}$	Column vectors containing all ones
$\mathbf{0}$	Column vectors containing all zeros
$\mathrm{diag}(\cdot)$	Diagonal matrix of appropriate dimensions
$\mathrm{tr}(\cdot)$	Trace of a matrix

$\det(\cdot)$	Determinant of a matrix
$\operatorname{supp}(\cdot)$	Support of a function
$\Theta = \mathbb{R}^d$	Parameter space
$\mathcal{T}\Theta = \mathbb{R}^d \times \mathbb{R}^d$	Tangent Bundle of Parameter Space
$\mathcal{M}(\mathcal{T}\Theta)$	Space of Probability Measures on $\mathcal{T}\Theta$
$\mathcal{M}_{\text{ac}}(\mathcal{T}\Theta)$	$\mu \in \mathcal{M}(\mathcal{T}\Theta)$ Absolutely Continuous

To Lord Eddard

Acknowledgments

Thank you for all the friends and family who have supported me in this long and tumultuous journey. I don't think I need to name any names but everyone has truly been spectacular. Without your patience to hear my rants during random late nights and constant love and support I would not be in this position. This dissertation would not be complete without your help.

Chapter 1

Introduction

In this dissertation we develop theory and algorithms for density prediction, density control, and machine learning. The need to predict and control density functions arises from varying engineering applications such as the control of biological and robotic swarms, vehicle guidance-control in uncertain dynamic environments, forecasting and demand response of loads in power systems, and active shaping of chemical concentrations in process control. In recent years, the study of density functions has also been used to understand the global convergence and performance of neural networks. Despite this recurring theme in practical applications, there does not exist a unified framework for systems-control of densities. The goal of this work is to close this gap and develop theories and algorithms for density prediction, control, and machine learning.

The goal of this chapter is to collect some notation, tools and heuristic arguments that will be used repeatedly throughout. In this work, the term "density function" can represent the uncertainty of the state of a single dynamical system modeled

by a stochastic differential equation

$$\mathrm{d}\boldsymbol{x}_t = \boldsymbol{f}(\boldsymbol{x}_t, t)\mathrm{d}t + \boldsymbol{g}(\boldsymbol{x}_t, t)\mathrm{d}\boldsymbol{w}_t, \quad \boldsymbol{x}_0 = \boldsymbol{\xi}_0, \tag{1.1}$$

where $\boldsymbol{w}_t \in \mathcal{X}$ is a standard Wiener process, the initial condition $\boldsymbol{\xi}_0 \sim \rho_0$ is subject to a known initial probability density function (PDF) ρ_0, and the vector fields $\boldsymbol{f} : \mathcal{X} \times \mathbb{R}^+ \mapsto \mathcal{X}$, $\boldsymbol{g} : \mathcal{X} \times \mathbb{R}^+ \mapsto \mathcal{Y}$ are the drift and diffusion coefficients, respectively with $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^n$. Under mild assumptions on these coefficients, for instance that $\boldsymbol{f}$ and $\boldsymbol{g}$ are Lipschitz, then the PDF $\rho(\boldsymbol{x}, t)$ defined by

$$\mathbb{P}(\boldsymbol{x}_t \in \Omega) = \int_\Omega \rho(\boldsymbol{x}, t)\, \mathrm{d}x \tag{1.2}$$

exists for any measurable set Ω. Here, $\mathbb{P}(\boldsymbol{x}_t \in \Omega)$ denotes the probability of a state vector being in the set Ω at a fixed time t. Moreover, this density function satisfies an evolution equation given by

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \boldsymbol{f}) + \frac{1}{2} \sum_{i,j=1}^{n} \frac{\partial^2}{\partial x_i \partial x_j} (\rho \boldsymbol{g}\boldsymbol{g}^\top)_{i,j}. \tag{1.3}$$

This second order Partial Differential Equation (PDE) is known as the Fokker-Planck Equation or Kolmogorov's Forward equation (FPK) [186]. A derivation of this equation is provided in Appendix A.1 for the convenience of the reader.

A density function may also well represent the state of a collection of systems modeled by an interacting particle system given by

$$\mathrm{d}\boldsymbol{x}_t^i = \boldsymbol{f}(\boldsymbol{x}_t^i, \rho^n(\cdot, t), t)\mathrm{d}t + \boldsymbol{g}(\boldsymbol{x}_t^i, \rho^n(\cdot, t), t)\mathrm{d}\boldsymbol{w}^i, \quad \boldsymbol{x}_0^i = \boldsymbol{\xi}_0^i, \tag{1.4}$$

$$\rho^n(\cdot, t) = \frac{1}{n} \sum_{k=1}^{n} \delta_{\boldsymbol{x}_t^i}(\cdot), \tag{1.5}$$

where $i = 1, \ldots, n$, and $\rho^n(\cdot, t)$ is the population density function (PDF) of the state of the n particle system. The population density function gives the total fraction of particles or the frequency of particles in a desired volume Ω of the state space. Here, we emphasize that $\boldsymbol{x}_t^i$ are all defined on the same state space $\mathcal{X}$ and $\boldsymbol{f}, \boldsymbol{g}$ are the same for each particle. The arguments of $\boldsymbol{f}$ and $\boldsymbol{g}$ are a state variable, a probability measure and time. To be more precise, the equations for each $\boldsymbol{x}_t^i$ are evaluated at the particle's own state $\boldsymbol{x}_t^i$ and their nonlocal interaction with other particles denoted by $\boldsymbol{x}_t^j$, is captured by $\rho^n(\cdot, t)$. The individual Wiener processes $\boldsymbol{w}_t^1, \boldsymbol{w}_t^2, \ldots, \boldsymbol{w}_t^n$ are also assumed to be independent of each other and the initial conditions are assumed to be $\boldsymbol{\xi}_0^i \sim \rho_0$.

This setup is a natural model for many scientific and engineering applications involving coupled multi-agent interaction such as crowd movement [69], opinion dynamics [8, 90], population biology [45, 208] and communication systems. Understanding the "mean field limit" of these particles, (i.e., when the number of particles $n- \to \infty$) is an important tool in the analysis of the n particle system as well as the individual particles. It is well-known that given an interacting particle system (1.4) with independent identically distributed (i.i.d.) initial states and that the vector fields $\boldsymbol{f}$ and $\boldsymbol{g}$ are sufficiently smooth, the limiting dynamics of the population density function are described by the Mckean-Vlasov or nonlinear Fokker-Planck Equation (MVFPK) given by

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \boldsymbol{f}(\boldsymbol{x}, \rho)) + \frac{1}{2} \sum_{i,j=1}^{n} \frac{\partial^2}{\partial x_i \partial x_j} (\rho \boldsymbol{g}(\boldsymbol{x}, \rho) \boldsymbol{g}^\top(\boldsymbol{x}, \rho))_{i,j}. \tag{1.6}$$

where $\rho(\cdot, t)$ is a suitable mean field limiting measure of $\rho^n(\cdot, t)$. The associated

Mckean-Vlasov SDE is given by

$$\mathrm{d}\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x}, \rho_t, t)\,\mathrm{d}t + \boldsymbol{g}(\boldsymbol{x}, \rho_t, t)\,\mathrm{d}\boldsymbol{w}_t. \tag{1.7}$$

Notice that the vector fields in this equation depend on the distribution of the solution itself which is a typical characteristic of Mckean-Vlasov SDEs. Deriving this limit rigorously involves weak convergence arguments from measure theory, but an outline of the proof is given in Appendix A.13.

Remark 1. *We remark that in either definition of density function, we can abbreviate either probability density function or population density function as PDF and will use it interchangeably depending on context all throughout this work.*

The goal of this dissertation is to seek a unified framework using an emerging geometric viewpoint to develop foundational theory, algorithms and scalable algorithms for

1. predicting the trajectory flow of $\rho(\cdot, t)$ generated by (1.3) or (1.6) in a fast and tractable manner

2. shaping and regulating the flow $\rho(\cdot, t)$ generated by (1.3) or (1.6) subject to initial and terminal time PDF constraints as well as state space constraints, and

3. studying the global convergence of neural networks and quantifying their performance by studying the dynamics of the density $\rho(\cdot, t)$ of a population of neurons in the mean-field limit

Although these problems are seemingly different questions, it turns out they are intimately linked via metric gradient flows. The core idea is to view the governing evolution equations (1.3) and (1.6) as gradient flows of certain Lyapunov

functionals w.r.t the Wasserstein metric arising from the theory of Optimal Transport [215]. There has been a burgeoning interest in the connections of optimal transport theory to density prediction, control, and learning but it is not clear whether this new viewpoint offers any practical benefits. In the upcoming subsections, we will expound more on the theory of gradient flows and uncover its connections with each of the three subproblems.

1.1 Introduction: Gradient Flows

In this section, we review the geometric notion of gradient flows. To conceptualize the main idea, we make an analogy between gradient flow in Euclidean space and gradient flow in the space of probability measures. The collected results below are only intended for a simple review of a rapidly evolving subject, but the main ideas are stated. We refer to the works [6, 215] for more detailed accounts of this topic.

We first review the metric viewpoint of gradient descent in finite dimension. Consider the continuous time gradient flow given by the initial value problem

$$\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}t} = -\nabla\varphi(\boldsymbol{x}), \quad \boldsymbol{x}(t=0) = \boldsymbol{x}_0. \tag{1.8}$$

where $\boldsymbol{x}, \boldsymbol{x}_0 \in \mathbb{R}^n$ and φ is a continuously differentiable function. Recall that the fixed points of this dynamical system are just the stationary points of the function φ since they must satisfy $\nabla\varphi = 0$. When φ is convex, we have a unique fixed point $\boldsymbol{x}^*$ where φ achieves its minima. An interesting feature of this dynamical system is revealed by writing its discretization in time. Let h be a small time step

and consider its (backward) Euler discretization

$$\boldsymbol{x}_k = \boldsymbol{x}_{k-1} - h\nabla\varphi(\boldsymbol{x}_k). \tag{1.9}$$

Rearranging yields

$$\frac{\boldsymbol{x}_k - \boldsymbol{x}_{k-1}}{h} + \nabla\varphi(\boldsymbol{x}_k) = 0 \Rightarrow \nabla\left(\frac{\|\boldsymbol{x} - \boldsymbol{x}_{k-1}\|_2^2}{2h} + \varphi(\boldsymbol{x})\right)_{\boldsymbol{x}=\boldsymbol{x}_k} = 0 \tag{1.10}$$

$$\Rightarrow \boldsymbol{x}_k = \underset{\boldsymbol{x}\in R^n}{arg\,min}\left\{\frac{\|\boldsymbol{x} - \boldsymbol{x}_{k-1}\|_2^2}{2h} + \varphi(\boldsymbol{x})\right\}. \tag{1.11}$$

This mapping from $\boldsymbol{x}_{k-1} \mapsto \boldsymbol{x}_k$ given by

$$\text{prox}_{h\varphi}^{\|\cdot\|_2}(\boldsymbol{x}_{k-1}) := \underset{\boldsymbol{x}\in R^n}{arg\,min}\left\{\frac{\|\boldsymbol{x} - \boldsymbol{x}_{k-1}\|_2^2}{2h} + \varphi(\boldsymbol{x})\right\} \tag{1.12}$$

is called the "proximal operator" [176] of $h\varphi$ w.r.t the standard Euclidean metric $\|\cdot\|_2$. This recursion generates a sequence of vectors $\{\boldsymbol{x}_k\}_{k=1}^{\infty}$ which converges pointwise to the solution of (1.8) in the sense that

$$\boldsymbol{x}_k \to \boldsymbol{x}(t = kh) \quad \text{as} \quad h \to 0. \tag{1.13}$$

Notice that φ also serves as a Lyapunov function since

$$\frac{\mathrm{d}}{\mathrm{d}t}\varphi = -\langle\nabla\varphi -, \nabla\varphi\rangle = -\|\varphi\|_2^2 \leq 0. \tag{1.14}$$

In other words, the value of the function φ decreases along the flow $\boldsymbol{x}(t)$ generated by this ODE and as $t \to \infty$ we get $\boldsymbol{x}(t) \to \boldsymbol{x}^*$ (see figure (1.1) for an illustration). Recasting in this form allows interpreting equation (1.8) as the steepest descent of the function φ measured w.r.t the Euclidean Distance $\|\cdot\|_2^2$. We want to extend this idea of steepest descent to dynamical systems in the space of measures.

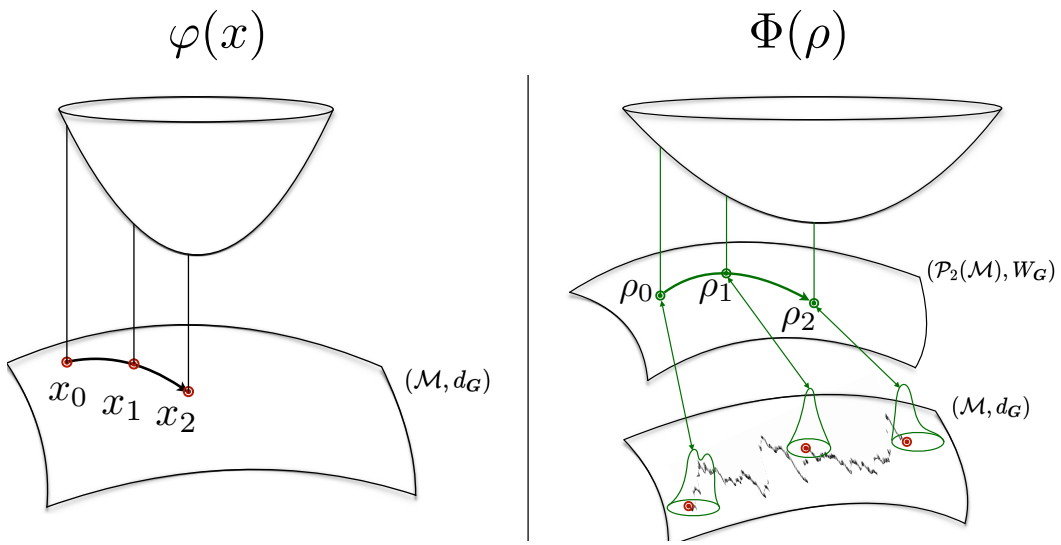

Figure 1.1: This figure depicts gradient flow in $\mathcal{M}$ (left) and in $P_2(\mathcal{M})$ (right). Here, we take $\mathcal{M} = \mathbb{R}^n$ and $G = I$ but the idea works for any compact manifold $\mathcal{M}$ with metric tensor G. We see the function decreases along the flow generated by the sequence of proximal evaluations. In finite dimensions, the flow reaches a fixed point given by the point where φ achieves its minima. Similarly, the flow in infinite dimensions approaches a fixed point given by the function that minimizes Φ which turns out to be given by the Gibbs distribution (1.29)

In the infinite dimensional setting, we are interested in computing $\rho(\cdot, t)$ generated by either (1.3) or (1.6) using gradient descent on the manifold of all PDFs with finite second moments denoted as

$$\mathcal{P}_2(\mathbb{R}^n) := \left\{ \rho : \mathbb{R}^n \mapsto \mathbb{R}^+ \middle| \int \rho \, \mathrm{d}\boldsymbol{x} = 1, \int \|x\|^2 \rho \, \mathrm{d}\boldsymbol{x} < \infty \right\}. \tag{1.15}$$

In other words, we want to design a proximal operator

$$\mathrm{prox}_{h\Phi}^{d(\cdot)}(\rho_{k-1}) := \underset{\rho \in \mathcal{P}_2(\mathbb{R}^n)}{\arg\inf} \frac{1}{2} d^2(\rho, \rho_{k-1}) + h\Phi(\rho), \tag{1.16}$$

where $h > 0$ is a time step and appropriately chosen distance metric $d(\cdot, \cdot)$ defined on $\mathcal{P}_2(\mathbb{R}^n)$ and functional $\Phi : \mathcal{P}_2(\mathbb{R}^n) \mapsto \mathbb{R}_{\geq 0}$. In the seminal work of Jordan, Kinderlehrer and Otto (JKO) [125], it was shown that if the Ito Stochastic Dif-

ferential Equation (1.1) is in "JKO Canonical Form" given by

$$\mathrm{d}\boldsymbol{x} = -\nabla U(\boldsymbol{x})\,\mathrm{d}t + \sqrt{2\beta^{-1}}\,\mathrm{d}\boldsymbol{w}, \tag{1.17}$$

that is to say, the vector field $\boldsymbol{f} = -\nabla U$ is the gradient of a smooth potential function U and $\boldsymbol{g} = \sqrt{2\beta^{-1}}\boldsymbol{I}$ is a scalar multiple of the identity matrix with $\beta > 0$. Then the associated Fokker-Planck equation

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \nabla U) + \beta^{-1}\Delta\rho, \quad \beta > 0, \quad U : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0}, \tag{1.18}$$

is gradient descent of the free energy functional

$$\Phi(\rho) := \int U\rho\,\mathrm{d}\boldsymbol{x} + \beta^{-1}\int \rho\log\rho\,\mathrm{d}\boldsymbol{x} \tag{1.19}$$

where the distance metric $d(\cdot,\cdot)$ is taken to be the 2-Wasserstein metric given by

$$W_2(\pi_1,\pi_2) := \left(\underset{d\pi\in\Pi(\pi_1,\pi_2)}{inf}\int_{\mathcal{X}\times\mathcal{Y}} \parallel \boldsymbol{x}-\boldsymbol{y}\parallel_2^2\ \mathrm{d}\pi\left(\boldsymbol{x},\boldsymbol{y}\right)\right)^{\frac{1}{2}}. \tag{1.20}$$

Notice that this metric is the solution of an optimization problem in itself. Here, $\mathrm{d}\pi_1(\boldsymbol{x}) = \rho_1(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$ and $\mathrm{d}\pi_2(\boldsymbol{y}) = \rho_2(\boldsymbol{y})\mathrm{d}\boldsymbol{y}$ are two probability measures supported respectively on $\mathcal{X},\mathcal{Y}\subseteq\mathbb{R}^n$. The Wasserstein Metric denoted as $W(\pi_1,\pi_2)$ (equivalently, $W\left(\rho_1,\rho_2\right)$ whenever π_1,π_2 are absolutely continuous so that the PDFs ρ_1,ρ_2 exist) arises in the theory of optimal mass transport [210]. $\Pi\left(\pi_1,\pi_2\right)$ denotes the collection of all probability measures on the product space $\mathcal{X}\times\mathcal{Y}$ having finite second moments, with marginals π_1 and π_2, respectively. Its square, $W^2(\pi_1,\pi_2)$ equals [21] the minimum amount of work required to transport π_1 to π_2 (or equivalently, ρ_1 to ρ_2). It is well-known [210] that $W(\pi_1,\pi_2)$ defines a metric on the manifold $\mathcal{P}_2(\mathbb{R}^n)$.

We consider the JKO scheme which is given by the discrete time sequence defined recursively by the sequence of optimizations

$$\rho_k = \underset{\rho \in P_2(R^n)}{arg\,min} \quad \frac{1}{2} W_2^2(\rho, \rho_{k-1}) + h\Phi(\rho). \tag{1.21}$$

The sequence of minimizers $\{\rho_k\}_{k=1}^{\infty}$ are well-defined (i.e., the optimization (1.21) has a solution each every step k) and we can define a piecewise, continuous time interpolation on the interval $t \in [hk, h(k+1))$ given by

$$\rho_t^h := \rho_k, \quad \text{for } t \in [hk, h(k+1)), \tag{1.22}$$

to approximate the solution $\rho(\cdot, t)$ of (1.18). The following theorem summarizes the main result in [125]:

Theorem 1 (JKO 1998)**.** *Let $\rho_0 \in \mathcal{P}_2(\mathbb{R}^n)$ be well-defined. Then, as $h \to 0$, $\rho_t^h \rightharpoonup \rho_t$ in $L^1(\mathbb{R}^n)$ where ρ_t satisfies the initial value problem*

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \nabla U) + \beta^{-1} \Delta \rho, \quad \rho(x, t = 0) = \rho_0. \tag{1.23}$$

This theorem states that the iterative scheme (1.21) converges to the solution of the Fokker-Planck Equation in the weak L^1 sense. Furthermore, observe that (1.18) can be rewritten as

$$\begin{aligned} \frac{\partial \rho}{\partial t} &= \nabla \cdot (\rho \nabla U) + \beta^{-1} \Delta \rho \\ &= \nabla \cdot (\rho \nabla (U + \beta^{-1} \log \rho)) \\ &= \nabla \cdot \left(\rho \nabla \frac{\delta \Phi}{\delta \rho} \right). \end{aligned} \tag{1.24}$$

It was shown by Otto in [174] that by formally interpreting $\mathcal{P}_2(\mathbb{R}^n)$ as a (infinite

dimensional) Riemannian Manifold and using the dynamical formulation of the Wasserstein metric W_2 as the Riemannian geodesic distance then the induced gradient operator is given by

$$\nabla^{W_2} := \nabla \cdot \left(\rho \nabla \frac{\delta}{\delta \rho} \right). \tag{1.25}$$

Therefore, we may formally write (1.24) as

$$\frac{\partial \rho}{\partial t} = \nabla^{W_2} \Phi(\rho). \tag{1.26}$$

In our context, (1.26) is the infinite dimensional analogue of (1.8). In addition, the Free Energy functional $\Phi(\cdot)$ serves as a Lyapunov functional by computing

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}\Phi(\rho) &= \int \frac{\partial \rho}{\partial t}(U + \beta^{-1} \log \rho + \beta^{-1}) \, \mathrm{d}\boldsymbol{x} \\ &= \int \nabla \cdot (\nabla (U + \beta^{-1} \log \rho)\rho)(U + \beta^{-1} \log \rho + \beta^{-1}) \, \mathrm{d}\boldsymbol{x} \\ &= -\int \|\nabla (U + \beta^{-1} \log \rho)\|_2^2 \rho \, \mathrm{d}\boldsymbol{x} \leq 0 \end{aligned} \tag{1.27}$$

along a solution trajectory $\rho(\cdot, t)$. Here, we used (1.18) in the first line and the integration by parts formula

$$\int u(\nabla \cdot v) \, \mathrm{d}\boldsymbol{x} = -\int \langle \nabla u, v \rangle \, \mathrm{d}\boldsymbol{x} \tag{1.28}$$

to go from the second to the last line. Moreover, equality is achieved at the

stationary density obtained by computing

$$\begin{aligned} 0 \equiv \frac{\mathrm{d}}{\mathrm{d}t}\Phi &\equiv -\int \|\nabla(U + \beta^{-1}\log\rho)\|_2^2 \rho \,\mathrm{d}\boldsymbol{x} \\ &\Rightarrow \nabla(U + \beta^{-1}\log\rho) \equiv 0 \\ &\Rightarrow U + \beta^{-1}\log\rho \equiv 0 \\ &\Rightarrow \rho_\infty = \frac{1}{Z}\exp(-\beta U), \quad Z = \int \exp(-\beta U)\,\mathrm{d}\boldsymbol{x}, \end{aligned} \tag{1.29}$$

where ρ_∞ is the well-known Gibbs distribution and Z is a normalizing constant known as the partition function.

In the presence of nonlocal interaction given by an interaction potential, one may consider the sample path dynamics with PDF dependent drift given by Mckean-Vlasov SDE

$$\mathrm{d}\boldsymbol{x} = -\nabla U(\boldsymbol{x}) - \nabla\left(V * \rho(\boldsymbol{x})\right)\,\mathrm{d}t + \sqrt{2\beta^{-1}}\mathrm{d}\boldsymbol{w}, \tag{1.30}$$

where U is as before, the interaction potential V is symmetric, $V(-\boldsymbol{x}) = V(\boldsymbol{x})$ for all $\boldsymbol{x} \in \mathbb{R}^n$ and $*$ denotes the convolution operator in $\mathbb{R}^n$. Then, the associated MVFPK equation is given by

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \nabla U + \nabla(V * \rho)) + \beta^{-1}\Delta\rho. \tag{1.31}$$

This PDE is a gradient flow of the generalized Free Energy functional [211]

$$F(\rho) := \int U\rho \,\mathrm{d}\boldsymbol{x} + \beta^{-1}\int \rho\log\rho \,\mathrm{d}\boldsymbol{x} + \int\int V(\boldsymbol{x}-\boldsymbol{y})\rho(\boldsymbol{x})\rho(\boldsymbol{y})\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{y}. \tag{1.32}$$

As before, the proximal operator with the distance metric $d \equiv W_2$ but now $\Phi \equiv F$ where F is given by (1.32) approximates the flow (1.31).

In the upcoming, subsections we describe some of the outstanding challenges in prediction, control, and learning, and how the theory of gradient flow and its generalizations provides a new perspective in tackling these challenges.

1.2 Contributions of this Work

1.2.1 Density Prediction

The problem of density or belief propagation, i.e., the problem of computing the transient joint PDF $\rho(\boldsymbol{x}, t)$ that solves a PDE of the form (1.3) or (1.6) is motivated by two types of problems. The *first* is dispersion analysis, where one is interested in predicting or analyzing the uncertainty evolution over time, e.g., in meteorological forecasting [83], spacecraft entry-descent-landing [94, 95], orientation density evolution for liquid crystals in chemical physics [111,126,164], and in motion planning [105,177,178]. In these applications, the quantity of interest is the joint PDF $\rho(\cdot, t)$ and its statistics (e.g., transient moments and marginal PDFs). The *second* type of applications require $\rho(\cdot, t)$ as an intermediate step toward finding other quantities of interest. For example, in nonlinear filtering [50, 77], the joint PDF $\rho(\cdot, t)$ serves as the prior in calculating the posterior (i.e., conditional state) PDF. In probabilistic model validation [96–98] and controller verification [103], computing $\rho(\boldsymbol{x}, t)$ helps in quantifying the density-level prediction-observation mismatch. Rather than developing algorithms on a case-by-case basis, these applications require fast computation of $\rho(\cdot, t)$ in a scalable and unified manner.

Given its widespread applications, density prediction has received sustained attention from the scientific computing community, but there remain several technical

challenges to achieve scalable computation. The dynamics of realistic systems are often nonlinear and evolve in high-dimensional state space. This means that we need algorithms that avoid spatial discretization or meshing of the state space because the predominant solution approaches like spatial discretization and function approximation or interpolation suffer from the "curse of dimensionality" [20]. The purpose of this work is to pursue the variational viewpoint that we discussed for computing $\rho(\cdot, t)$ without generating any discretization and/or meshing the state space. The idea is to seek a representation of $\rho(\boldsymbol{x}, t)$ as a linear combination of Dirac distributions

$$\rho(\boldsymbol{x}, t) \approx \sum_{i=1}^{N} \varrho^i(t)\delta(\boldsymbol{x} - \boldsymbol{x}^i(t)), \tag{1.33}$$

where the each points $\boldsymbol{x}^i(t) \in \mathcal{X}$ are independently evolved according to the SDE (1.17)

$$\mathrm{d}\boldsymbol{x}_t^i = -\nabla U(\boldsymbol{x}_t^i)\,\mathrm{d}t + \sqrt{2\beta^{-1}}\,\mathrm{d}\boldsymbol{w}_t^i. \tag{1.34}$$

The coefficients $\varrho^i(t)$ are called the particle weights and are obtained by solving a particle version of the JKO scheme (1.21). These coefficients represent the amount of mass carried by ith particle located at $\boldsymbol{x}^i(t)$ at time t, and N is the total number of particles. We refer to these types of solutions as particle solutions, and the Dirac delta distributions introduced above are called the particles. These types of schemes have several advantages compared to traditional methods like finite-differences and finite-element methods. Particle schemes are concentrated in the region of the state space in which we are interested thereby optimizing memory storage. They are also beneficial for problems with complicated geometries or problems with changing boundaries due to the dynamic nature of the particles.

In Chapter 2, we propose a novel computational framework that solves for the particle weights. We show that by introducing entropic regularization and taking the dual of the particle version of the objective (1.21) leads to a cone-preserving fixed point recursion that is proved to be contractive in an appropriate metric. A block co-ordinate iteration scheme is proposed to solve the resulting nonlinear recursions with guaranteed convergence. This approach enables remarkably fast computation for non-parametric transient joint PDF propagation.

At the same time, the JKO scheme in its original form has limited scope. Many stochastic systems of interest do not admit state dynamics that are in JKO canonical form. Even the linear-Gaussian setup ubiquitous in systems-control theory

$$\mathrm{d}\boldsymbol{x}_t = \boldsymbol{A}\boldsymbol{x}_t\,\mathrm{d}t + \sqrt{2}\boldsymbol{B}\,\mathrm{d}\boldsymbol{w}_t, \quad \boldsymbol{x}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0), \tag{1.35}$$

where the drift coefficient is a Hurwitz linear vector field, and the diffusion coefficient is a constant non-square real matrix seems to fall outside the scope of the JKO scheme. In many engineering applications, one also usually encounters SDE's where the drift vector fields have both a dissipative and conservative structure. For example, stochastic systems like Kramers' equation

$$\mathrm{d}\boldsymbol{q}_t = \boldsymbol{p}_t\,\mathrm{d}t \tag{1.36}$$

$$\mathrm{d}\boldsymbol{p}_t = -\nabla V(\boldsymbol{q}_t)\,\mathrm{d}t - \gamma\nabla F(\boldsymbol{p}_t)\,\mathrm{d}t + \sqrt{2\gamma\beta^{-1}}\,\mathrm{d}\boldsymbol{w}_t \tag{1.37}$$

arising from Newton's law in classical mechanics often have this structure. Motivated by the JKO scheme, the authors of [81,117], have established approximation schemes for Kramer's equation. The proposed proximal algorithm in Chapter 2 can be used to address problems of this type. The main challenge in the construction of such a variational scheme for Kramer's equation is finding the appropriate

cost function in the distance metric and a suitable free energy functional. This difficulty arises from the fact that Kramer's equation is not uniformly diffusive in the entire state space. Diffusion is only present in the velocity variables, not in the position variables and this has to be accounted for in the distance functional. This means that techniques in [125] cannot be applied directly although the schemes are the same form as (1.21)

1.2.2 Density Control

Building on our work in Chapter 2, we develop theory and numerical solvers for the density control problem i.e., the problem of steering the density $\rho(\cdot, t)$ subject to trajectory level dynamics while minimizing the control effort with two-point density constraints. The density control problem arises in modeling and control of the collective dynamics of an ensemble of *physical populations*. They arise in the study of self-organization and behaviors in organisms such as social insects. This work in biological systems has inspired numerous studies in the collective behavior of artificial systems like robotic swarms, [15] which have been increasingly garnered attention from large scale applications like environmental exploration, reconnaissance and surveillance and disaster response. Depending on the mission, it is critical to find the control laws that shape the behavior of the swarm in response to possibly unknown changes in the environment. In addition, density control also arises in the study of ensembles of neurons [163], shaping the bulk magnetization distribution for NMR spectroscopy [146] and density of highway traffic [63]. In all these applications, the population distribution is actively controlled over time while preserving the physical mass. The conservation of mass allows us to reformulate the problem as steering a single system with probabilistic

uncertainty in its initial and terminal state modeled via prescribed initial and terminal joint state probability instead of steering a large number of systems with identical dynamics.

Mathematically, this can be set up as non-standard stochastic optimal control problem of the form

$$\inf_{\boldsymbol{u}\in U} \quad \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2 \,\mathrm{d}t\right\}, \tag{1.38a}$$

$$\text{subject to} \quad \mathrm{d}\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x},t)\,\mathrm{d}t + \boldsymbol{B}(t)\boldsymbol{u}(\boldsymbol{x},t)\,\mathrm{d}t + \sqrt{2\epsilon}\boldsymbol{B}(t)\,\mathrm{d}\boldsymbol{w}(t), \tag{1.38b}$$

$$\boldsymbol{x}(t=0)\sim\rho_0(\boldsymbol{x}), \quad \boldsymbol{x}(t=1)\sim\rho_1(\boldsymbol{x}), \tag{1.38c}$$

where $\boldsymbol{x}\in\mathbb{R}^n$, and the set $\mathcal{U}$ comprises of all finite energy inputs, as before. Given $\boldsymbol{f}(\boldsymbol{x},t)$, $\boldsymbol{B}(t)$ and ϵ, our objective is to steer the joint state PDF $\rho(\boldsymbol{x},t)$ from a prescribed initial PDF ρ_0 at $t=0$ to another prescribed terminal PDF ρ_1 at $t=1$ while minimizing the expected control effort.

Define the diffusion tensor $\boldsymbol{D}(t) := \boldsymbol{B}(t)\boldsymbol{B}(t)^\top$. Problem (1.38) can be formally recast into a "fluid dynamics" version [21] given by:

$$\inf_{(\rho,\boldsymbol{u})} \quad \frac{1}{2}\int_0^1\int_{\mathbb{R}^n}\|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2\,\rho(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}t \tag{1.39a}$$

$$\text{subject to} \quad \frac{\partial\rho}{\partial t} + \nabla\cdot\left(\rho\left(\boldsymbol{f}+\boldsymbol{B}(t)\boldsymbol{u}\right)\right) = \epsilon\mathbf{1}^\top(\boldsymbol{D}(t)\odot\mathbf{Hess}\,(\rho))\,\mathbf{1}, \tag{1.39b}$$

$$\rho(\boldsymbol{x},0)=\rho_0(\boldsymbol{x}), \quad \rho(\boldsymbol{x},1)=\rho_1(\boldsymbol{x}), \tag{1.39c}$$

where the infimum is taken over all pairs $(\rho,\boldsymbol{u})\in\mathcal{P}_2(\mathbb{R}^n)\times\mathcal{U}$ satisfying (1.39b)-(1.39c). We note that (3.15b) is the controlled FPK PDE which governs the flow of the joint PDF associated with the SDE (1.38b). In 1931-32, Erwin Schrödinger published two papers [191,192] for the case $\boldsymbol{f}\equiv 0$ and $\boldsymbol{B}\equiv\boldsymbol{I}$ which we refer to as

the classical *Schrödinger Bridge Problem* (SBP). Hereafter, we refer to (1.39a) as the generalized *Schrödinger Bridge Problem* (SBP). Recently, the classical SBP has been extended [60] to the case when the prior dynamics is a linear time-varying (LTV) system, i.e.,

$$\mathrm{d}\boldsymbol{x}(t) = \boldsymbol{A}(t)\boldsymbol{x}(t)\,\mathrm{d}t + \boldsymbol{B}(t)\boldsymbol{u}(\boldsymbol{x},t)\,\mathrm{d}t + \sqrt{2\epsilon}\boldsymbol{B}(t)\,\mathrm{d}\boldsymbol{w}(t), \tag{1.40}$$

where the system matrices $\boldsymbol{A}(t) \in \mathbb{R}^{n\times n}, \boldsymbol{B}(t) \in \mathbb{R}^{n\times m}$, $m \leq n$, and the pair $(\boldsymbol{A}(t), \boldsymbol{B}(t))$ is assumed to be controllable for all t. We refer the readers to [60, Sec. 4] for the details.

In Chapter 3, we recap the basics of the SBP and use this as foundation to solve the SBP with generic nonlinear prior dynamics. The idea is to deduce the necessary conditions of optimality to obtain the optimal state feedback policy in terms of the solution of a certain Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE) that has a one-way coupling with the optimally PDF given by a controlled FPK equation. We show that we can transform this nonlinear system into a system of boundary-coupled linear PDE using a Hopf-Cole transformation which we refer to as the *Schrödinger System* whose solutions recover the optimal state feedback and the optimally controlled joint state PDF. To numerically solve the Schrödinger System, we will utilize the Wasserstein proximal recursions described in Chapter 2. We will apply the prior nonlinear dynamics given in (1.38b) to two different cases: gradient drift, and mixed conservative-dissipative drift.

In Chapter 4, we extend these results to generalized SBP with hard deterministic state constraints. This containment can be achieved by reflecting the sample paths from the ϵ-inner boundary layer of $\partial\mathcal{X}$ for ϵ small enough at all times. We show that this restriction results in Neumann Boundary conditions in (1.39a). In

applications, state constraints may encode safety requirements such as obstacle avoidance. We show how the same algorithms from Chapter 2 can be leveraged to tackle these challenges.

In Chapter 5, we derive theoretical results for generic SBP's subject to multi-input state feedback linearizable dynamics. We show that it is possible to exploit the structural nonlinearities to derive a *Schrödinger System* via adding an artificial diffusion coefficient. We show that the diffeomorphism induced by the feedback linearizable structure can be used to transform into a problem with simplified state space dynamics. This is at the cost of a modified Lagrangian cost function. We envision that the theoretical developments in this work will help design algorithms that solve the feedback density steering over a large class of nonlinear dynamical systems.

1.3 Machine Learning

In Chapter 6, we study the global convergence of neural networks from the perspective of gradient flows. In recent years, neural networks have become a staple for engineering applications, but despite their empirical success, there remains a critical gap in understanding their performance and convergence guarantees. The purpose of this study is to shed some light and explain the empirical success by exploring the dynamics in the space of distributions supported over the parameter space. In particular, one can view shallow neural networks as interacting particle systems that admit a mean-field limit. Under certain assumptions, their training dynamics can be viewed as gradient flows with global convergence guarantees.

In the case of fully connected networks, universal approximation results such

as [16, 74] explain this empirical success. They demonstrate that a large enough network can approximate any continuous function on a compact set. The focus of this work is the dynamical version of the universal approximation theorem wherein we study whether local search algorithms such as stochastic gradient descent (SGD) can find global solutions. Recent works [65, 158, 188, 194] have addressed this question for two-layer networks, showing convergence to global solutions by studying the dynamics in the space of distributions over parameters that is given by some nonlinear PDE with gradient flow structure w.r.t. Wasserstein metric. This work seeks to extend these results to second-order momentum methods such as Heavy Ball and Nesterov's methods. A description of the problem setting is presented first, followed by a summary of the results.

The basic problem of machine learning can be described as follows. We are given data $\{(\boldsymbol{y}_i, \boldsymbol{x}_i)\}_{i=1}^n$ which are identically distributed from a distribution $\mathcal{D}$. The data $\boldsymbol{x}_i \in \mathcal{X}$ is a feature vector (one can think of it as a descriptor of an image) and $\boldsymbol{y}_i \in \mathcal{Y}$ is a response or label (we can think of this as the object described by the image). The goal of the learning problem is to find a function $\hat{y} : \mathcal{X} \mapsto \mathcal{Y}$ that models the dependency of $\boldsymbol{y}_i$ on $\boldsymbol{x}_i$. In most cases, the function $\hat{y}(\boldsymbol{x}; \boldsymbol{\theta})$ is parameterized by the vector $\boldsymbol{\theta}$ which makes learning this function equivalent to finding the parameters such that our predictor $\hat{y}$ is accurate when measured with respect to a certain loss or metric.

In this work, we consider the simplest case of a single hidden layer where the predictor is of the form

$$\hat{y}(\boldsymbol{x}, \boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^{N} \sigma_*(\boldsymbol{x}, \boldsymbol{\theta}), \tag{1.41}$$

where N is the number of neurons (also referred to as the width the network).

For the sake of convenience, we take $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{Y} = \mathbb{R}$ (See Figure (1.2)). In the two-layer case, we set

$$\sigma_*(\boldsymbol{x}, \boldsymbol{\theta}^i) = a^i \sigma(\langle \boldsymbol{x}, \boldsymbol{b}^i \rangle), \tag{1.42}$$

with $\boldsymbol{\theta}^i = (a^i, \boldsymbol{b}^i) \in \mathbb{R} \times \mathbb{R}^d$ and $\boldsymbol{b}^i, a^i$ are the the weights of the first and second layer, respectively, and $\sigma : \mathbb{R} \mapsto \mathbb{R}$ is an activation function. Here, we also collectively denote $\boldsymbol{\theta} = (\boldsymbol{\theta}^1, \boldsymbol{\theta}^2, \ldots, \boldsymbol{\theta}^n) \in \mathbb{R}^{(d+1)N}$. Notice that this neural network is written as an average of activation functions parameterized by $\boldsymbol{\theta}^i$ and the factor $1/N$ is introduced for convenience. Although, this is an unusual way to describe the output of the network, this is a favorable setting to do mean field analysis i.e., taking the width $N \to \infty$.

For the sake of exposition, we consider the case of quadratic loss

$$\begin{aligned} R_N(\boldsymbol{\theta}) &= \mathbb{E}_{\mathcal{D}}\{(y - \hat{y})^2\} \\ &= \mathbb{E}_{\mathcal{D}} \left\{ \left(y - \frac{1}{N} \sum_{i=1}^{N} \sigma_*(x, \boldsymbol{\theta}^i) \right)^2 \right\}, \end{aligned} \tag{1.43}$$

but this analysis applies to more general loss functions. The goal of the learning problem is to make this expected loss as small as possible given a choice of the parameter $\boldsymbol{\theta}$. In other words, the learning problem can be set up as the optimization problem

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^{(d+1)N}} \quad R_N(\boldsymbol{\theta}). \tag{1.44}$$

In practice, we don't have access to the distribution $\mathcal{D}$ so we approximate using

the samples which gives

$$\mathbb{E}_{\mathcal{D}}\left\{\left(y - \frac{1}{N}\sum_{i=1}^{N}\sigma_*(\boldsymbol{x},\boldsymbol{\theta}^i)\right)^2\right\} \approx \sum_{k=1}^{L}\left\{\left(y - \frac{1}{N}\sum_{i=1}^{N}\sigma_*(\boldsymbol{x}_k,\boldsymbol{\theta}_i)\right)^2\right\}. \tag{1.45}$$

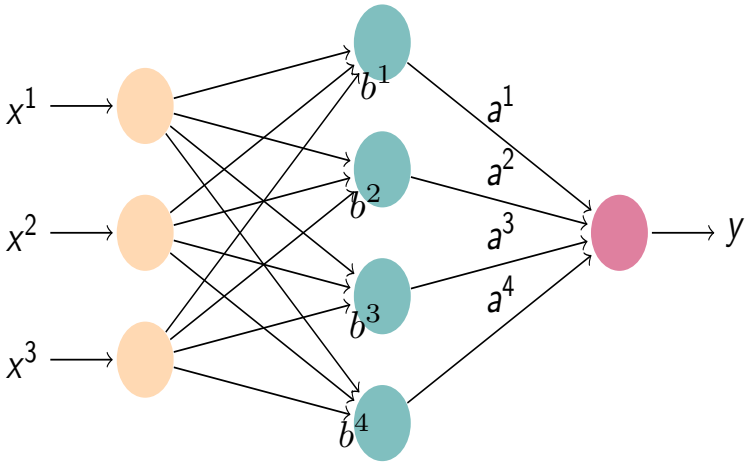

Figure 1.2: A neural network with $N = 4$ number of neurons in the hidden layer and $d = 3$ in the input layer.

Despite many advances in learning, a fundamental theory for explaining its effectiveness (and possible limitations) are only beginning to emerge. We bring to light some of the major difficulties:

1. **Non-Convexity**: The dependence on the input weights $\boldsymbol{b}^i$'s is non-linear because of the activation function. This typically leads to non-convexity.

2. **Overparameterization**: The number N of hidden units is very large and exceeds the number of observations. This leads to a very high dimensional optimization problem since $\boldsymbol{\theta} \in \mathbb{R}^{(d+1)N}$.

Overparametrization of neural networks does not typically result in a degradation of the network's performance. In light of this, mean-field analysis is a viable strat-

egy for analysing optimization dynamics; see [5, 7, 17]. This overparameterization actually enables us to obtain some theoretical results for width $N \to \infty$.

In practice, Stochastic Gradient Descent (SGD) remains one of the most effective algorithms for machine learning. The basic idea behind SGD is to obtain the current iterate by updating the previous iterate in the direction of the gradient. In mathematical terms, we can think of the parameters of each neurons as "particles" evolving according to the recursion

$$\boldsymbol{\theta}_k^i = \boldsymbol{\theta}_{k-1}^i - \eta N \nabla_{\boldsymbol{\theta}^i} R_N(\boldsymbol{\theta}_{k-1}^i) \tag{1.46}$$

for $i = 1, 2, \ldots, N$ and η is a stepsize. Notice that the scaling N does not affect the objective and we put it there for convenience. We run this recursion until we get reasonably close to a minimum. As $\eta \to 0$, this algorithm converges to the ODE

$$\frac{\mathrm{d}\boldsymbol{\theta}_t^i}{\mathrm{d}t} = -N \nabla R_N(\boldsymbol{\theta}_t), \tag{1.47}$$

which is called gradient flow. One other interesting variant of SGD is Noisy SGD (also referrred to as the Langevin Algorithm in the literature) where in each step we add a Gaussian Noise $\boldsymbol{w}_{k-1} \sim \mathcal{N}(0, \boldsymbol{I}\eta)$ giving

$$\boldsymbol{\theta}_k^i = \boldsymbol{\theta}_{k-1}^i - \eta N \nabla R_N(\boldsymbol{\theta}_{k-1}^i) + \sqrt{2\eta\beta^{-1}}\boldsymbol{w}_{k-1}^i. \tag{1.48}$$

Similarly, when the time step $\eta \to 0$, we also have a continuous time analogue given by the Ito Stochastic Differential Equation (SDE)

$$\mathrm{d}\boldsymbol{\theta}_t^i = -N \nabla R_N(\boldsymbol{\theta}_t)\mathrm{d}t + \sqrt{2\beta^{-1}}\mathrm{d}\boldsymbol{w}_t. \tag{1.49}$$

Here β^{-1} is called the inverse temperature, and we will see this parameter play a big role in our analysis. Throughout this work, we focus on noisy SGD because of its nice regularization properties that are useful for the analysis. We will prove results where we take $\beta^{-1} \to 0$ to obtain some qualitative results about SGD which is our main problem of interest. At this point, we can use the theory of gradient flows and optimal transport to collect the relevant results. We refocus our attention to the learning problem. We can start by expanding the quadratic risk (1.43)

$$\begin{aligned} R_N(\boldsymbol{\theta}) &= \mathbb{E}_{\mathcal{D}}\{(y-\hat{y})^2\} \\ &= \mathbb{E}_{\mathcal{D}}\{y^2 - y\hat{y}^2 + \hat{y}^2\} \\ &= \underbrace{\mathbb{E}_{\mathcal{D}}\{y^2\}}_{R_\sharp} + \frac{1}{N}\sum_{i=1}^{N} \underbrace{\mathbb{E}_{\mathcal{D}}\{-y\sigma^*(x,\boldsymbol{\theta}^i)\}}_{:=V(\boldsymbol{\theta}^i)} + \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} \underbrace{\mathbb{E}_{\mathcal{D}}\{\sigma_*(x,\boldsymbol{\theta}^i)\sigma_*(x,\widetilde{\boldsymbol{\theta}}^j)\}}_{:=U(\boldsymbol{\theta}^i,\widetilde{\boldsymbol{\theta}}^j)} \\ &:= R_\sharp + \frac{1}{N}\sum_{i=1}^{N} V(\boldsymbol{\theta}^i) + \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N} U(\boldsymbol{\theta}^i,\widetilde{\boldsymbol{\theta}}^j). \end{aligned} \tag{1.50}$$

For large N, it makes sense to view (1.50) as an empirical average in the parameter space so we have

$$R(\rho) := R_\sharp + \int V(\boldsymbol{\theta})\rho(\boldsymbol{\theta})\mathrm{d}\boldsymbol{\theta} + \int\int U(\boldsymbol{\theta},\widetilde{\boldsymbol{\theta}})\rho(\boldsymbol{\theta})\rho(\widetilde{\boldsymbol{\theta}})\mathrm{d}\boldsymbol{\theta}\mathrm{d}\widetilde{\boldsymbol{\theta}} \tag{1.51}$$

for some density $\rho \in \mathcal{P}_2(\mathbb{R}^{(d+1)})$. To aid our analysis, it is useful to consider the regularized version

$$R_\lambda(\rho) = R(\rho) + \lambda \int g(\boldsymbol{\theta})\rho(\boldsymbol{\theta})\,\mathrm{d}\theta, \tag{1.52}$$

where $g(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|_2^2$ and $\lambda > 0$ is some regularization parameter. In general, we can take any $g(\boldsymbol{\theta}) \geq \|\boldsymbol{\theta}\|_2^2$. It was shown in [158] that for some constant C that

depends only on U, we have

$$\left| \inf_{\theta \in \mathbb{R}^{(d+1)N}} \; R_{N,\lambda}(\boldsymbol{\theta}) \; - \inf_{\rho \in \mathcal{P}_2(\mathbb{R}^{(d+1)})} \; R_\lambda(\rho) \right| \leq \frac{K}{N}, \tag{1.53}$$

where

$$R_{N,\lambda}(\boldsymbol{\theta}) = R_N(\boldsymbol{\theta}) + \frac{\lambda}{N} \sum_{i=1}^{N} g(\boldsymbol{\theta}^i). \tag{1.54}$$

So instead of optimizing a non-convex problem over parameter space we lift the problem to a problem with a strongly convex objective (since U is positive semi-definite) over $\mathcal{P}_2(\mathbb{R}^{d+1})$ which is an infinite dimensional manifold. Recall, that we are trying to solve the optimization problem (1.44) with quadratic regularization and we update each particle via the continuous time noisy SGD dynamics given by

$$\mathrm{d}\boldsymbol{\theta}_t^i = -N \nabla_{\boldsymbol{\theta}^i} R_{N,\lambda}(\boldsymbol{\theta}_t) \mathrm{d}t + \sqrt{2\beta^{-1}} \mathrm{d}\boldsymbol{w}_t, \tag{1.55}$$

where $i = 1, 2, \ldots, N$. Since we want to take $N \to \infty$ it is intractable to consider the dynamics of each particle so instead we consider the collective dynamics. We can formalize this idea by consider the population density function

$$\rho_t(\cdot) := \lim_{N \to \infty} \frac{1}{N} \sum_{i=1}^{N} \delta_{\boldsymbol{\theta}_t^i}(\cdot), \tag{1.56}$$

which evolves according to

$$\frac{\partial \rho_t}{\partial t} = \nabla \cdot \left(\nabla \frac{\delta R_\lambda}{\delta \rho} \rho_t \right) + \beta^{-1} \Delta \rho_t \tag{1.57}$$

$$= \nabla \cdot \left(\nabla \frac{\delta \Phi_\lambda}{\delta \rho} \right) = \nabla^{W_2} \Phi_\lambda(\rho), \tag{1.58}$$

where the free energy Φ is the the sum of the terms

$$\Phi_\lambda(\rho) = R_\lambda(\rho) + \beta^{-1} S(\rho), \quad S(\rho) = \int \rho \log \rho. \tag{1.59}$$

Moreover, this PDE is gradient flow of w.r.t the Wasseretein metric of the functional R_λ plus some entropic regularization term. Because of this geometric structure, all solution trajectories ρ_t of (1.58) with some initial condition ρ_0 will converge to a unique stationary solution ρ_∞ which is also the unique minimizer of R_λ which is given by the fixed-point Gibbs density

$$\rho_\infty(\boldsymbol{\theta}) = \frac{1}{Z} \exp\left(-\beta \Phi_\lambda\left(\rho_\infty(\boldsymbol{\theta}).\right)\right) \tag{1.60}$$

Notice that this is different from the Gibbs distribution in the linear Fokker-Planck equation because we have an implicit equation, i.e., ρ_∞ shows up on both sides of (1.60).

In Chapter 6, we extend these results to optimization algorithms with second-order dynamics like the Heavy Ball Method [185] given by

$$\mathrm{d}\boldsymbol{\theta}_t^i = \boldsymbol{r}_t^i \, \mathrm{d}t, \tag{1.61}$$

$$\mathrm{d}\boldsymbol{r}_t^i = -N \nabla_{\theta_t^i} R_{N,\lambda}(\boldsymbol{\theta}_t) - \gamma \boldsymbol{r}_t^i + \sqrt{2\gamma\beta^{-1}} \, \mathrm{d}\boldsymbol{w}_t^i. \tag{1.62}$$

We study the distribution dynamics for a two-layer neural network and derive the distribution evolution given by a nonlinear Kinetic Fokker-Planck equation

$$\frac{\partial \rho}{\partial t} = -\langle \boldsymbol{r}, \nabla_{\boldsymbol{\theta}} \rho \rangle + \nabla_{\boldsymbol{r}} \cdot \left(\rho \left(\gamma \boldsymbol{r} + \lambda \boldsymbol{\theta} + \nabla_{\boldsymbol{\theta}} \Phi'([\rho]^{\boldsymbol{\theta}}) \right) \right) + \gamma \beta^{-1} \Delta_{\boldsymbol{r}} \rho, \tag{1.63}$$

where ρ is now supported over the parameter space and velocity space and $[\rho]^{\boldsymbol{\theta}}$

denotes the θ marginal PDF of the ρ. Unlike gradient dynamics, we do not have the luxury of using Wasserstein gradient flow to demonstrate convergence. The main results of this work are to characterize the convergence of the global minimum and establish its existence and uniqueness. The approach we take in this work will involve previous methods from the SGD case combined with tools from the study of the (linear) Kinetic Fokker-Planck equation. Our goal is to extend this work beyond second-order algorithms to a larger class of algorithms.

Finally, in Chapter 7 we state the conclusions of this work as well as future directions of research. Hereafter, we present the journal publications, conference publications and pre-prints that have resulted from this work:

Journal Papers

- **K.F. Caluya**, and A. Halder. Gradient Flow Algorithms for Density Propagation in Stochastic Systems. IEEE Transactions on Automatic Control
- A. Halder, **K.F. Caluya**, B. Travacca, and S.J. Moura. Hopfield Neural Network Flow: A Geometric Viewpoint. IEEE Transactions on Neural Networks and Learning Systems
- **K.F. Caluya**, and A. Halder. Wasserstein Proximal Algorithms for the Schrödinger Bridge Problem: Density Control with Nonlinear Drift. IEEE Transactions on Automatic Control
- A, Halder, **K.F. Caluya**, P. Ojaghi, X. Geng. Stochastic Uncertainty Propagation in Power System Dynamics using Measured-value Proximal Recursions. IEEE Transactions on Power Systems arXiv:2108.13405

Conference Papers

- **K.F. Caluya**, and A. Halder. Proximal Recursion for Solving the Fokker-Planck Equation. *American Control Conference*, Philadelphia, 2019.

- **K.F. Caluya**, and A. Halder. Finite Horizon Density Steering for Multi-input State Feedback Linearizable Systems. *American Control Conference*, Denver, 2020.

- **K.F. Caluya**, and A. Halder. Reflected Schrödinger Bridge: Density Control with Path Constraints *American Control Conference*, New Orleans, 2021.

- S Haddad, **K.F. Caluya**, A Halder, B Singh. Prediction and Optimal Feedback Steering of Probability Density Functions for Safe Automated Driving *American Control Conference*, New Orleans, 2021.

Preprints

- W.Krichene, **K.F. Caluya**, and A. Halder. Global Convergence of Second-order Dynamics in Two-layer Neural Networks. arXiv:2007.06852

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

Chapter 2

Gradient Flow Algorithms for Density Propagation in Stochastic Systems

2.1 Introduction

Consider the continuous-time dynamics of the state vector $\boldsymbol{x}(t) \in \mathbb{R}^n$ governed by an Itô stochastic differential equation (SDE)

$$\mathrm{d}\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x}, t)\,\mathrm{d}t \;+\; \boldsymbol{g}(\boldsymbol{x}, t)\,\mathrm{d}\boldsymbol{w}, \quad \boldsymbol{x}(t=0) = \boldsymbol{x}_0, \tag{2.1}$$

where the joint probability density function (PDF) for the initial condition $\boldsymbol{x}_0$ is a known function ρ_0; we use the notation $\boldsymbol{x}_0 \sim \rho_0$. The process noise $\boldsymbol{w}(t) \in \mathbb{R}^m$ is Wiener and satisfy $\mathbb{E}\left[\mathrm{d}w_i \mathrm{d}w_j\right] = \delta_{ij}\mathrm{d}t$ for all $i, j = 1, \ldots, n$, where $\delta_{ij} = 1$ for $i = j$, and zero otherwise. Then, the flow of the joint PDF $\rho(\boldsymbol{x}, t)$ for the state

vector $\boldsymbol{x}(t)$ (i.e., $\boldsymbol{x} \sim \rho$) is governed by the partial differential equation (PDE) initial value problem:

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \boldsymbol{f}) + \frac{1}{2} \sum_{i,j=1}^{n} \frac{\partial^2}{\partial x_i \partial x_j} (\rho \boldsymbol{g}\boldsymbol{g}^\top)_{ij}, \tag{2.2a}$$

$$\rho(\boldsymbol{x}, t = 0) = \rho_0(\boldsymbol{x}) \quad \text{(given)}. \tag{2.2b}$$

The *second order* transport PDE (2.2a) is known as the *Fokker-Planck* or *Kolmogorov's forward equation* [186]. Hereafter, we will refer it as the Fokker-Planck-Kolmogorov (FPK) PDE.

In this paper, we consider the problem of density or belief propagation, i.e., the problem of computing the transient joint PDF $\rho(\boldsymbol{x}, t)$ that solves a PDE initial value problem of the form (2.2). From an application standpoint, the need for computing $\rho(\boldsymbol{x}, t)$ can be motivated by two types of problems. The *first* is dispersion analysis, where one is interested in predicting or analyzing the uncertainty evolution over time, e.g., in meteorological forecasting [83], spacecraft entry-descent-landing [94, 95], orientation density evolution for liquid crystals in chemical physics [111, 126, 164], and in motion planning [105, 177, 178]. In these applications, the quantity of interest is the joint PDF $\rho(\boldsymbol{x}, t)$ and its statistics (e.g., transient moments and marginal PDFs). The *second* type of applications require $\rho(\boldsymbol{x}, t)$ as an intermediate step toward computing other quantities of interest. For example, in nonlinear filtering [50, 77], the joint PDF $\rho(\boldsymbol{x}, t)$ serves as the prior in computing the posterior (i.e., conditional state) PDF. In probabilistic model validation [96–98] and controller verification [103], computing $\rho(\boldsymbol{x}, t)$ helps in quantifying the density-level prediction-observation mismatch. All these applications require fast computation of $\rho(\boldsymbol{x}, t)$ in a scalable and unified manner, as opposed to developing algorithms in a case-by-case basis.

Given its widespread applications, problem (2.2) has received sustained attention from the scientific computing community where the predominant solution approaches have been spatial discretization and function approximation – both of which, in general, suffer from the "curse of dimensionality" [20]. The purpose of this paper is to pursue a systems-theoretic variational viewpoint for computing $\rho(\boldsymbol{x}, t)$ that breaks away from the "solve PDE as a PDE" philosophy, and instead solves (2.2) as a gradient descent on the manifold of joint PDFs. This emerging geometric viewpoint for uncertainty propagation and filtering has been reported in our recent work [100, 101], but it remained unclear whether this viewpoint can offer computational benefit over the standard PDE solvers. It is not at all obvious whether and how an infinite-dimensional gradient descent can numerically obviate function approximation or spatial discretization. The contribution of this paper is to demonstrate that not only this is possible, but also that the same enables fast computation.

To conceptualize the main idea, we appeal to the *metric viewpoint* of gradient descent, where a continuous-time gradient flow is realized by small time-step recursions of a proximal operator with respect to (w.r.t.) a suitable metric. For example, consider the finite dimensional gradient flow

$$\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}t} = -\nabla\varphi\left(\boldsymbol{x}\right), \quad \boldsymbol{x}(0) = \boldsymbol{x}_0, \tag{2.3}$$

where $\boldsymbol{x}, \boldsymbol{x}_0 \in \mathbb{R}^n$, and the continuously differentiable function $\varphi : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$. The flow $\boldsymbol{x}(t)$ generated by (2.3) can be realized via variational recursion

$$\boldsymbol{x}_k = \underset{\boldsymbol{x}\in R^n}{arg\,min}\ \frac{1}{2} \parallel \boldsymbol{x} - \boldsymbol{x}_{k-1} \parallel_2^2 + h\,\varphi(\boldsymbol{x}) + o(h), \quad k \in \mathbb{N}, \tag{2.4}$$

in the sense that as the step-size $h \downarrow 0$, we have $\boldsymbol{x}_k \to \boldsymbol{x}(t = kh)$, i.e., the sequence

$\{\boldsymbol{x}_k\}$ converges pointwise to the flow $\boldsymbol{x}(t)$. This can be verified by rewriting the Euler discretization of (2.3), given by

$$\boldsymbol{x}_k - \boldsymbol{x}_{k-1} = -h\nabla\varphi(\boldsymbol{x}_{k-1}),$$

as

$$\begin{aligned}\boldsymbol{x}_k &= \underset{\boldsymbol{x}\in R^n}{arg\,min}\ \frac{1}{2} \parallel \boldsymbol{x} - (\boldsymbol{x}_{k-1} - h\nabla\varphi(\boldsymbol{x}_{k-1})) \parallel_2^2 \\ &= \underset{\boldsymbol{x}\in R^n}{arg\,min}\ \frac{1}{2} \parallel \boldsymbol{x} - \boldsymbol{x}_{k-1} \parallel_2^2 + \langle \boldsymbol{x} - \boldsymbol{x}_{k-1}, h\nabla\varphi(\boldsymbol{x}_{k-1})\rangle \\ &\qquad\qquad + h\varphi(\boldsymbol{x}_{k-1}), \end{aligned} \tag{2.5}$$

where we used the fact that adding and omitting constant terms do not change the arg min. From (2.5), one can arrive at (2.4) by invoking first order approximation of $\varphi(\boldsymbol{x})$ at $\boldsymbol{x} = \boldsymbol{x}_k$. In (2.4), the mapping $\boldsymbol{x}_{k-1} \mapsto \boldsymbol{x}_k$ given by

$$\text{prox}_{h\varphi}^{\|\cdot\|_2}(\boldsymbol{x}_{k-1}) := \underset{\boldsymbol{x}\in R^n}{arg\,min}\ \frac{1}{2} \parallel \boldsymbol{x} - \boldsymbol{x}_{k-1} \parallel_2^2 + h\,\varphi(\boldsymbol{x}), \tag{2.6}$$

is called the "proximal operator" [176, p. 142] of $h\varphi$ w.r.t. the standard Euclidean metric $\|\cdot\|_2$. Notice that $\varphi(\cdot)$ serves as a Lyapunov function since the quantity

$$\frac{\mathrm{d}}{\mathrm{d}t}\varphi = \langle \nabla\varphi, -\nabla\varphi\rangle = -\parallel \nabla\varphi \parallel_2^2 \tag{2.7}$$

equals 0 at the stationary point of (2.3), and < 0 otherwise.

In the infinite dimensional setting, we are interested in computing the flow generated by (2.2) via gradient descent on the manifold of joint PDFs with finite second

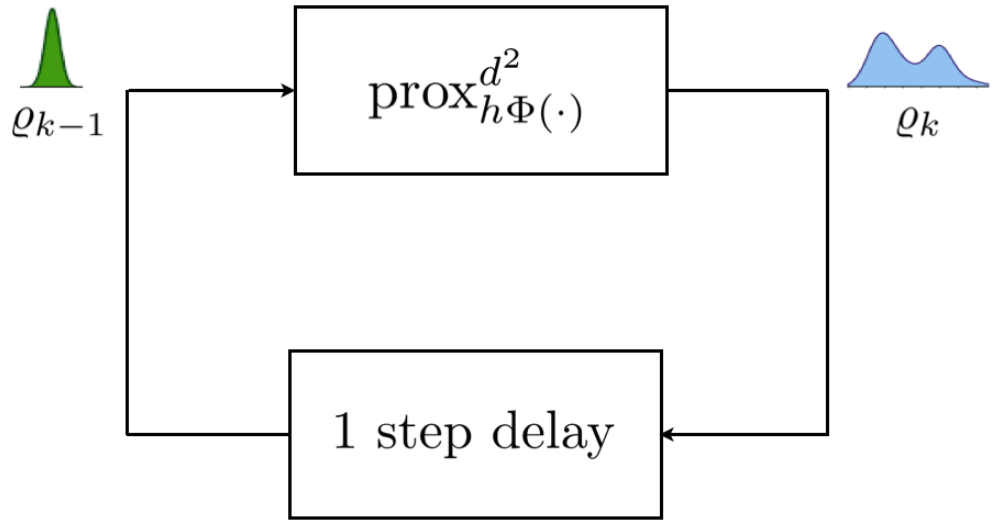

Figure 2.1: The gradient descent on the manifold of PDFs can be described by successive evaluation of proximal operators to recursively update PDFs from time $t = (k-1)h$ to $t = kh$ for $k \in \mathbb{N}$, and time-step $h > 0$.

(raw) moments, denoted as*

$$\mathscr{D}_2 := \{\rho : \mathbb{R}^n \mapsto \mathbb{R} \mid \rho \geq 0, \int_{\mathbb{R}^n} \rho = 1,\ \mathbb{E}_\rho[\boldsymbol{x}^\top \boldsymbol{x}] < \infty\}.$$

Specifically, let $d(\cdot,\cdot)$ be a distance metric on the manifold $\mathscr{D}_2$, and let the functional $\Phi : \mathscr{D}_2 \mapsto \mathbb{R}_{\geq 0}$. Then, for some chosen step-size $h > 0$, the infinite dimensional proximal operator of $h\Phi$ w.r.t. the distance metric $d(\cdot,\cdot)$, given by

$$\text{prox}^{d(\cdot)}_{h\Phi}(\varrho_{k-1}) := \underset{\varrho \in D_2}{arg\,inf}\ \frac{1}{2} d\left(\varrho, \varrho_{k-1}\right)^2 + h\,\Phi(\varrho), \tag{2.8}$$

can be used to define a proximal recursion (Fig. 2.1):

$$\varrho_k = \text{prox}^{d(\cdot)}_{h\Phi}(\varrho_{k-1}), \quad k \in \mathbb{N}, \quad \varrho_0(\boldsymbol{x}) := \rho_0(\boldsymbol{x}). \tag{2.9}$$

Just as the proximal recursion (2.4) approximates the finite dimensional flow (2.3), similarly it is possible to design $d(\cdot,\cdot)$ and $\Phi(\cdot)$ in (2.8) as function of the drift

*We denote the expectation operator w.r.t. the measure $\rho(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$ as $\mathbb{E}_\rho[\cdot]$.

$\boldsymbol{f}$ and diffusion $\boldsymbol{g}$ in (2.2a) such that the proximal recursion (2.9) approximates the infinite dimensional flow (2.2). This idea was first proposed in [125], showing that when $\boldsymbol{f}$ is a gradient vector field and $\boldsymbol{g}$ is a scalar multiple of identity matrix, then the distance $d(\cdot,\cdot)$ can be taken as the Wasserstein-2 metric [210] with $\Phi(\cdot)$ as the free energy functional. In particular, the solution of (2.9) was shown to converge to the flow of (2.2), i.e., $\varrho_k(\boldsymbol{x}) \to \rho(\boldsymbol{x}, t = kh)$ in strong $L_1(\mathbb{R}^n)$ sense, as $h \downarrow 0$. The resulting variational recursion (2.9) has since been known as the *Jordan-Kinderlehrer-Otto (JKO) scheme* [6], and we will refer to the FPK operator (2.2a) with such assumptions on $\boldsymbol{f}$ and $\boldsymbol{g}$ to be in "JKO canonical form". Similar gradient descent schemes have been derived for many other PDEs; see e.g., [190] for a recent survey.

The remaining of this paper is organized as follows. Section 2.2 explains the JKO canonical form and the corresponding free energy functional $\Phi(\cdot)$. Our algorithms and convergence results are presented in Section 2.3, followed by numerical simulation results in Section 2.4. Section 2.5 provides various extensions of the basic algorithm showing how the framework proposed here can be applied to systems not in JKO canonical form. Section 2.6 concludes the paper.

We remark here that a preliminary version [42] of this work appeared in the 2019 American Control Conference. This paper significantly expands [42] by incorporating additional results for the so-called McKean-Vlasov flow (Sections 2.2.2 and 2.4.3), and by providing various extensions of the basic algorithm (Section 2.5).

Preliminaries

We next collect definitions and some properties of the 2-Wasserstein metric and the Kullback-Leibler divergence, which will be useful in the development below.

Definition 1. *(**2-Wasserstein metric**) The 2-Wasserstein metric between two probability measures* $\mathrm{d}\pi_1(\boldsymbol{x}) := \rho_1(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$ *and* $\mathrm{d}\pi_2(\boldsymbol{y}) := \rho_2(\boldsymbol{y})\mathrm{d}\boldsymbol{y}$*, supported respectively on* $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^n$*, is denoted as* $W(\pi_1, \pi_2)$ *(equivalently,* $W(\rho_1, \rho_2)$ *whenever the measures* π_1, π_2 *are absolutely continuous so that the respective PDFs* ρ_1, ρ_2 *exist); it is defined as*

$$W(\pi_1, \pi_2) := \left(\underset{d\pi \in \Pi(\pi_1, \pi_2)}{inf} \int_{\mathcal{X}\times\mathcal{Y}} s(\boldsymbol{x}; \boldsymbol{y}) \, \mathrm{d}\pi \, (\boldsymbol{x}, \boldsymbol{y}) \right)^{\frac{1}{2}}, \tag{2.10}$$

where $s(\boldsymbol{x}; \boldsymbol{y}) :=\| \boldsymbol{x} - \boldsymbol{y} \|_2^2$ *is the squared Euclidean distance in* $\mathbb{R}^n$*, and* $\Pi(\pi_1, \pi_2)$ *denotes the collection of all joint probability measures on* $\mathcal{X} \times \mathcal{Y}$ *having finite second moments, with marginals* π_1 *and* π_2*, respectively.*

The existence and uniqueness of the minimizer in (2.10) is guaranteed. It is well-known [210, Ch. 7] that $W(\pi_1, \pi_2)$ defines a metric on the manifold $\mathscr{D}_2$. This means that $W(\pi_1, \pi_2) \geq 0$ with equality if and only if $\pi_1 = \pi_2$, the symmetry: $W(\pi_1, \pi_2) = W(\pi_2, \pi_1)$, and that $W(\pi_1, \pi_2)$ satisfies the triangle inequality. Its square, $W^2(\pi_1, \pi_2)$ equals [21] the minimum amount of work required to transport π_1 to π_2 (or equivalently, ρ_1 to ρ_2), and vice versa. For any PDF ν, the function $\rho \mapsto W^2(\rho, \nu)$ is convex on $\mathscr{D}_2$, i.e., for any $\rho_1, \rho_2 \in \mathscr{D}_2$, and $0 \leq \tau \leq 1$, we have

$$W^2(\nu, (1-\tau)\rho_1 + \tau\rho_2) \;\leq\; (1-\tau)W^2(\nu, \rho_1) + \tau W^2(\nu, \rho_2). \tag{2.11}$$

Definition 2. *(**Kullback-Leibler divergence**) The Kullback-Leibler divergence, also known as relative entropy, between two probability measures* $\mathrm{d}\pi_1(\boldsymbol{x}) := \rho_1(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$

and $\mathrm{d}\pi_2(\boldsymbol{y}) := \rho_2(\boldsymbol{y})\mathrm{d}\boldsymbol{y}$*, denoted as* $D_{\mathrm{KL}}(\pi_1||\pi_2)$*, is defined as*

$$\begin{aligned} D_{\mathrm{KL}}(\pi_1||\pi_2) &= \int \left(\frac{\mathrm{d}\pi_1}{\mathrm{d}\pi_2}\right) \log \left(\frac{\mathrm{d}\pi_1}{\mathrm{d}\pi_2}\right) \mathrm{d}\pi_2 \\ &= \int_{\mathbb{R}^n} \rho_1(\boldsymbol{x}) \log \frac{\rho_1(\boldsymbol{x})}{\rho_2(\boldsymbol{x})} \mathrm{d}\boldsymbol{x}, \end{aligned} \tag{2.12}$$

where $\frac{\mathrm{d}\pi_1}{\mathrm{d}\pi_2}$ *denotes the Radon-Nikodym derivative. Henceforth, we use the notational equivalence* $D_{\mathrm{KL}}(\pi_1||\pi_2) \equiv D_{\mathrm{KL}}(\rho_1||\rho_2)$*.*

From Jensen's inequality, $D_{\mathrm{KL}}(\rho_1||\rho_2) \geq 0$; however, D_{KL} is not a metric since it is neither symmetric, nor does it satisfy the triangle inequality. The Kullback-Leibler divergence (2.12) is jointly convex in ρ_1 and ρ_2.

2.2 JKO Canonical Form

The JKO canonical form refers to a continuous-time stochastic dynamics where the drift vector field is the gradient of a potential function. The potential can be state dependent, or can depend on both the state $\boldsymbol{x}(t)$ and its joint PDF $\rho(\boldsymbol{x}, t)$. In the former case, the associated flow of the joint state PDF is governed by a FPK PDE of the form (2.2a) whereas in the latter case, the same is governed by the McKean-Vlasov integro-PDE. For both cases, the proposed gradient flow algorithms (Section 2.3) will be able to compute the transient state PDFs $\rho(\boldsymbol{x}, t)$. We next describe these canonical forms; generalization of our framework to systems *not* in JKO canonical form will be given in Sections 2.4 and 2.5.

2.2.1 FPK Gradient Flow

Consider the Itô SDE

$$\mathrm{d}\boldsymbol{x} = -\nabla\psi\left(\boldsymbol{x}\right)\mathrm{d}t \;+\; \sqrt{2\beta^{-1}}\,\mathrm{d}\boldsymbol{w}, \quad \boldsymbol{x}(0) = \boldsymbol{x}_0, \tag{2.13}$$

where the the drift potential $\psi : \mathbb{R}^n \mapsto (0, \infty)$, the diffusion coefficient $\beta > 0$, and the initial condition $\boldsymbol{x}_0 \sim \rho_0(\boldsymbol{x})$. For the sample path $\boldsymbol{x}(t)$ dynamics given by the SDE (2.13), the flow of the joint PDF $\rho\left(\boldsymbol{x}, t\right)$ is governed by the FPK PDE initial value problem

$$\frac{\partial\rho}{\partial t} = \nabla\cdot\left(\rho\nabla\psi\right) \;+\; \beta^{-1}\Delta\rho, \quad \rho(\boldsymbol{x}, 0) = \rho_0(\boldsymbol{x}). \tag{2.14}$$

It is easy to verify that the unique stationary solution of (2.14) is the Gibbs PDF $\rho_\infty(\boldsymbol{x}) = \kappa\exp\left(-\beta\psi(\boldsymbol{x})\right)$, where the normalizing constant κ is known as the *partition function*.

A Lyapunov functional associated with (2.14) is the *free energy*

$$\begin{aligned} F(\rho) \;&:=\; \mathbb{E}_\rho\left[\psi + \beta^{-1}\log\rho\right] \\ &=\; \beta^{-1}D_{\mathrm{KL}}\left(\rho \parallel \exp\left(-\beta\psi(\boldsymbol{x})\right)\right) \geq 0, \end{aligned} \tag{2.15}$$

that decays [125] along the solution trajectory of (2.14). This follows from re-writing (2.14) as

$$\frac{\partial\rho}{\partial t} = \nabla\cdot\left(\rho\nabla\zeta\right), \quad \text{where} \quad \zeta := \beta^{-1}\left(1 + \log\rho\right) + \psi, \tag{2.16}$$

and consequently

$$\frac{\mathrm{d}}{\mathrm{d}t}F = -\mathbb{E}_\rho\left[\| \nabla\zeta \|_2^2\right], \tag{2.17}$$

which is < 0 for the transient solution $\rho(\boldsymbol{x}, t)$, and $= 0$ at the stationary solution $\rho_\infty(\boldsymbol{x}) = \kappa\exp(-\beta\psi(\boldsymbol{x}))$. In our context, (2.17) is an analogue of (2.7). Notice that the free energy (2.15) can be seen as the sum of the *potential energy* $\int_{\mathbb{R}^n}\psi(\boldsymbol{x})\rho\,\mathrm{d}\boldsymbol{x}$ and the *internal energy* $\beta^{-1}\int_{\mathbb{R}^n}\rho\log\rho\,\mathrm{d}\boldsymbol{x}$. If $\psi \equiv 0$, the PDE (2.14) reduces to the heat equation, which by (2.15), can then be interpreted as an entropy maximizing flow.

The seminal result of [125] was that the transient solution of (2.14) can be computed via the proximal recursion (2.9) with the distance metric $d \equiv W$ (i.e., the 2-Wasserstein metric in (2.10)) and the functional $\Phi \equiv F$ (i.e., the free energy (2.15)). Just as (2.3) and (2.4) form a gradient flow-proximal recursion pair, likewise (2.14) and (2.9) form the same with the stated choices of d and Φ. From (2.10), notice that the distance metric W itself is defined as the solution of an optimization problem, hence it is not apparent how to numerically implement the recursion (2.9) in a scalable manner.

2.2.2 McKean-Vlasov Gradient Flow

In addition to the drift and diffusion, if one has non-local interaction, then the corresponding PDF evolution equation becomes the McKean-Vlasov integro-PDE

$$\frac{\partial\rho}{\partial t} = \nabla\cdot\left(\rho\nabla\left(\psi + \rho * \phi\right)\right) + \beta^{-1}\Delta\rho, \quad \rho(\boldsymbol{x}, 0) = \rho_0(\boldsymbol{x}), \tag{2.18}$$

where $*$ denotes the convolution in $\mathbb{R}^n$, the interaction potential $\phi : \mathbb{R}^n \mapsto (0, \infty)$ and is symmetric, i.e., $\phi(-\boldsymbol{v}) = \phi(\boldsymbol{v})$ for $\boldsymbol{v} \in \mathbb{R}^n$. The associated sample path $\boldsymbol{x}(t)$ dynamics has PDF-dependent drift:

$$\mathrm{d}\boldsymbol{x} = -\left(\nabla\psi\left(\boldsymbol{x}\right) + \nabla\left(\rho * \phi\right)\right)\mathrm{d}t + \sqrt{2\beta^{-1}}\,\mathrm{d}\boldsymbol{w}, \quad \boldsymbol{x}(0) = \boldsymbol{x}_0. \tag{2.19}$$

As an example, when $\phi(\boldsymbol{v}) := \frac{1}{2} \parallel \boldsymbol{v} \parallel^2$, then $\nabla(\rho * \phi)(\boldsymbol{x}) = \boldsymbol{x} - \mathbb{E}_\rho[\boldsymbol{x}]$. Clearly, (2.18) reduces to (2.14) in the absence of interaction ($\phi \equiv 0$). The McKean-Vlasov equation serves as a model for coupled multi-agent interaction in applications such as crowd movement [69], opinion dynamics [8,90], population biology [45,208] and communication systems.

A Lyapunov functional for (2.18) can be obtained [211] by generalizing the free energy (2.15) as

$$F(\rho) := \mathbb{E}_\rho\left[\psi \;+\; \beta^{-1}\log\rho \;+\; \rho * \phi\right], \tag{2.20}$$

which is a sum of the potential energy (as before), the internal energy (as before), and the *interaction energy* $\frac{1}{2}\int_{\mathbb{R}^{2n}} \phi(\boldsymbol{x}-\boldsymbol{y})\rho(\boldsymbol{x})\rho(\boldsymbol{y})\mathrm{d}\boldsymbol{x}\mathrm{d}\boldsymbol{y}$. In this case, (2.17) holds with

$$\zeta := \beta^{-1}\left(1 + \log\rho\right) + \psi + \rho * \phi, \tag{2.21}$$

which is often referred to as the *entropy dissipation functional* [46]. As before, the proximal recursion (2.9) with $d \equiv W$ and $\Phi \equiv F$ (now F given by (2.20)), approximates the flow (2.18).

2.3 Framework

We now describe our computational framework to solve the proximal recursion

$$\varrho_k = \operatorname{prox}^{W}_{hF(\cdot)}(\varrho_{k-1}) \tag{2.22a}$$

$$= \underset{\varrho \in D_2}{arg\,inf}\ \frac{1}{2}W^2(\varrho_{k-1}, \varrho) + h\, F(\varrho), \quad k \in \mathbb{N}, \tag{2.22b}$$

with $\varrho_0 \equiv \rho_0(\boldsymbol{x})$ (the initial joint PDF) for some small step-size h. For maximal clarity, we develop the framework with the free energy $F(\cdot)$ as in (2.15), i.e., for FPK gradient flow. In Section 2.4.C, we will show how the same can be generalized when $F(\cdot)$ is of the form (2.20). As per the convexity properties mentioned in Section 3.1.2, problem (2.22) involves (recursively) minimizing sum of two convex functionals, and hence is a convex problem for each $k \in \mathbb{N}$.

We discretize time as $t = 0, h, 2h, \ldots$, and develop an algorithm to solve (2.22) without making any spatial discretization. Specifically, we would like to perform the recursion (2.22) on weighted scattered point cloud $\{\boldsymbol{x}_k^i, \varrho_k^i\}_{i=1}^N$ of cardinality N at $t_k = kh$, $k \in \mathbb{N}$, where the location of the i^{th} point $\boldsymbol{x}_k^i \in \mathbb{R}^n$ denotes its state-space coordinate, and the corresponding weight $\varrho_k^i \in \mathbb{R}_{\geq 0}$ denotes the value of the joint PDF evaluated at that point at time t_k. Such weighted point cloud representation of (2.22) results in the following problem:

$$\boldsymbol{\varrho}_k = \underset{\boldsymbol{\varrho}}{arg\,min}\Big\{ \min_{\boldsymbol{M} \in \Pi(\boldsymbol{\varrho}_{k-1}, \boldsymbol{\varrho})} \frac{1}{2}\langle \boldsymbol{C}_k, \boldsymbol{M}\rangle + h\,\langle \boldsymbol{\psi}_{k-1} + \beta^{-1}\log \boldsymbol{\varrho}, \boldsymbol{\varrho}\rangle \Big\}, \quad k \in \mathbb{N}, \tag{2.23}$$

where the drift potential vector $\boldsymbol{\psi}_{k-1} \in \mathbb{R}^N$ is given by

$$\boldsymbol{\psi}_{k-1}(i) := \psi\left(\boldsymbol{x}_{k-1}^i\right), \quad i = 1, \ldots, N. \tag{2.24}$$

Here, the probability vectors $\boldsymbol{\varrho}, \boldsymbol{\varrho}_{k-1} \in \mathcal{S}_{N-1}$, the probability simplex in $\mathbb{R}^N$. Furthermore, for each $k \in \mathbb{N}$, the matrix $\boldsymbol{C}_k \in \mathbb{R}^{N\times N}$ is given by

$$\boldsymbol{C}_k(i,j) := s(\boldsymbol{x}_k^i; \boldsymbol{x}_{k-1}^j) = \| \boldsymbol{x}_k^i - \boldsymbol{x}_{k-1}^j \|_2^2, \ i,j = 1, \ldots, N, \tag{2.25}$$

and $\Pi(\boldsymbol{\varrho}_{k-1}, \boldsymbol{\varrho})$ stands for the set of all matrices $\boldsymbol{M} \in \mathbb{R}^{N\times N}$ such that

$$\boldsymbol{M} \geq 0, \quad \boldsymbol{M}\mathbf{1} = \boldsymbol{\varrho}_{k-1}, \quad \boldsymbol{M}^\top \mathbf{1} = \boldsymbol{\varrho}. \tag{2.26}$$

Notice that the inner minimization in (2.23) is a standard linear programming problem if it were to be solved for a given $\boldsymbol{\varrho} \in \mathcal{S}_{N-1}$, as in the Monge-Kantorovich optimal mass transport [210]. However, the outer minimization in (2.23) precludes a direct numerical approach.

To circumvent the aforesaid issues, following [130], we first regularize and then dualize (2.23). Specifically, adding an entropic regularization $H(\boldsymbol{M}) := \langle \boldsymbol{M}, \log \boldsymbol{M} \rangle$ in (2.23) yields

$$\begin{aligned} \boldsymbol{\varrho}_k = \underset{\boldsymbol{\varrho}}{arg\,min} \Big\{ \underset{\boldsymbol{M}\in\Pi(\boldsymbol{\varrho}_{k-1},\boldsymbol{\varrho})}{min} & \frac{1}{2}\langle \boldsymbol{C}_k, \boldsymbol{M}\rangle + \epsilon H(\boldsymbol{M}) \\ & + h\left\langle \boldsymbol{\psi}_{k-1} + \beta^{-1}\log\boldsymbol{\varrho}, \boldsymbol{\varrho}\right\rangle \Big\}, \end{aligned} \tag{2.27}$$

where $\epsilon > 0$ is a regularization parameter. The entropic regularization is standard in optimal mass transport literature [22, 73] and leads to efficient Sinkhorn iteration for the inner minimization. Here we point out that there has been par-

allel work [53, 59] on the connection between optimal mass transport and the so called Schrödinger bridge problem which is a dynamic version of this type of regularization.

In our context, the entropic regularization “algebrizes" the inner minimization in the sense if $\boldsymbol{\lambda}_0, \boldsymbol{\lambda}_1 \in \mathbb{R}^N$ are Lagrange multipliers associated with the equality constraints in (2.26), then the optimal coupling matrix $\boldsymbol{M}^{\text{opt}} := [m^{\text{opt}}(i,j)]_{i,j=1}^N$ in (2.27) has the Sinkhorn form

$$m^{\text{opt}}(i,j) = \exp\left(\boldsymbol{\lambda}_0(i)h/\epsilon\right) \exp\left(-\boldsymbol{C}_k(i,j)/(2\epsilon)\right) \exp\left(\boldsymbol{\lambda}_1(j)h/\epsilon\right). \tag{2.28}$$

Since the objective in (2.27) is proper convex and lower semi-continuous in $\boldsymbol{\varrho}$, the *strong duality* holds, and we consider the Lagrange dual of (2.27) given by:

$$\left(\boldsymbol{\lambda}_0^{\text{opt}}, \boldsymbol{\lambda}_1^{\text{opt}}\right) = \underset{\boldsymbol{\lambda}_0, \boldsymbol{\lambda}_1 \in R^N}{arg\,max} \Big\{ \langle \boldsymbol{\lambda}_0, \boldsymbol{\varrho}_{k-1} \rangle - F^\star(-\boldsymbol{\lambda}_1) - \frac{\epsilon}{h} \Big(\exp(\boldsymbol{\lambda}_0^\top h/\epsilon) \exp(-\boldsymbol{C}_k/2\epsilon) \exp(\boldsymbol{\lambda}_1 h/\epsilon) \Big) \Big\}, \tag{2.29}$$

where

$$F^\star(\cdot) := \sup_{\vartheta} \left\{ \langle \cdot, \vartheta \rangle \; - \; F(\vartheta) \right\} \tag{2.30}$$

is the *Legendre-Fenchel conjugate* of the free energy F in (2.15). Next, we derive the first order optimality conditions for (2.29) resulting in the proximal updates, and then provide an algorithm to solve the same.

2.3.1 Proximal Recursions

Given the vectors $\boldsymbol{\varrho}_{k-1}, \boldsymbol{\psi}_{k-1}$, the matrix $\boldsymbol{C}_k$, and the positive scalars β, h, ϵ in (2.29), let

$$\boldsymbol{y} := \exp(\boldsymbol{\lambda}_0 h/\epsilon), \qquad \boldsymbol{z} := \exp(\boldsymbol{\lambda}_1 h/\epsilon), \tag{2.31}$$

$$\boldsymbol{\Gamma}_k := \exp(-\boldsymbol{C}_k/2\epsilon), \quad \boldsymbol{\xi}_{k-1} := \exp(-\beta\boldsymbol{\psi}_{k-1} - \mathbf{1}). \tag{2.32}$$

The following result establishes a system of nonlinear equations for computing $\boldsymbol{\lambda}_0^{\text{opt}}, \boldsymbol{\lambda}_1^{\text{opt}}$ in (2.29), and consequently $\boldsymbol{\varrho}_k$ in (2.27).

Theorem 1. *The vectors $\boldsymbol{\lambda}_0^{\text{opt}}, \boldsymbol{\lambda}_1^{\text{opt}}$ in (2.29) can be found by solving for $\boldsymbol{y}, \boldsymbol{z} \in \mathbb{R}^N$ from the following system of equations:*

$$\boldsymbol{y} \odot (\boldsymbol{\Gamma}_k \boldsymbol{z}) = \boldsymbol{\rho}_{k-1}, \tag{2.33a}$$

$$\boldsymbol{z} \odot \left(\boldsymbol{\Gamma}_k{}^{\top} \boldsymbol{y}\right) = \boldsymbol{\xi}_{k-1} \odot \boldsymbol{z}^{-\frac{\beta\epsilon}{h}}, \tag{2.33b}$$

and then inverting the maps (2.31). Let $(\boldsymbol{y}^{\text{opt}}, \boldsymbol{z}^{\text{opt}})$ be the solution of (2.33). The vector $\boldsymbol{\varrho}_k$ in (2.27), i.e., the proximal update (Fig. 2.1) can then be obtained as

$$\boldsymbol{\varrho}_k = \boldsymbol{z}^{\text{opt}} \odot \left(\boldsymbol{\Gamma}_k{}^{\top} \boldsymbol{y}^{\text{opt}}\right). \tag{2.34}$$

Proof. From (2.15), the "discrete free energy" is

$$F(\boldsymbol{\varrho}) = \langle \boldsymbol{\psi} + \beta^{-1} \log \boldsymbol{\varrho}, \boldsymbol{\varrho} \rangle.$$

Its Legendre-Fenchel conjugate, by (2.30), is given by

$$F^{\star}(\boldsymbol{\lambda}) = \sup_{\boldsymbol{\varrho}} \{\boldsymbol{\lambda}^{\top}\boldsymbol{\varrho} - \boldsymbol{\psi}^{\top}\boldsymbol{\varrho} - \beta^{-1}\boldsymbol{\varrho}^{\top}\log\boldsymbol{\varrho}\}. \tag{2.35}$$

Setting the gradient of the objective function in (2.35) w.r.t. $\boldsymbol{\varrho}$ to zero, and solving for $\boldsymbol{\varrho}$ yields

$$\boldsymbol{\varrho}_{\max} = \exp(\beta(\boldsymbol{\lambda} - \boldsymbol{\psi}) - \mathbf{1}). \tag{2.36}$$

Substituting (2.36) back into (2.35), results

$$F^{\star}(\boldsymbol{\lambda}) = \beta^{-1}\mathbf{1}^{\top}\exp(\beta(\boldsymbol{\lambda} - \boldsymbol{\psi}) - \mathbf{1}). \tag{2.37}$$

Fixing $\boldsymbol{\lambda}_1$, and taking the gradient of the objective in (2.29) w.r.t. $\boldsymbol{\lambda}_0$, gives (2.33a). Likewise, fixing $\boldsymbol{\lambda}_0$, and taking the gradient of the objective in (2.29) w.r.t. $\boldsymbol{\lambda}_1$ gives

$$\nabla_{\boldsymbol{\lambda}_1} F^{\star}(-\boldsymbol{\lambda}_1) = \boldsymbol{z} \odot \left(\boldsymbol{\Gamma}_k{}^{\top}\boldsymbol{y}\right). \tag{2.38}$$

Using (2.37) to simplify the left-hand-side of (2.38) results in (2.33b). To derive (2.34), notice that combining the last equality constraint in (2.26) with (2.28), (2.31) and (2.32) gives

$$\boldsymbol{\varrho}_k = (\boldsymbol{M}^{\mathrm{opt}})^{\top}\mathbf{1} = \sum_{j=1}^{N} m^{\mathrm{opt}}(j,i) = \boldsymbol{z}(i)\sum_{j=1}^{N}\boldsymbol{\Gamma}_k(j,i)\boldsymbol{y}(j),$$

which is equal to $\boldsymbol{z} \odot \boldsymbol{\Gamma}_k^{\top}\boldsymbol{y}$, as claimed. ■

In the following (Section 2.3.2), we propose an algorithm to solve (2.33) and (2.34),

and then outline the overall implementation of our computational framework. The convergence results for the proposed algorithm are given in Section 2.3.3.

2.3.2 Algorithm

Proximal algorithm

We now propose a block co-ordinate iteration scheme to solve (2.33). The proposed procedure, which we call ProxRecur, and detail in Algortihm 1, takes $\boldsymbol{\varrho}_{k-1}$ as input and returns the proximal update $\boldsymbol{\varrho}_k$ as output for $k \in \mathbb{N}$. In addition to the data $\boldsymbol{\varrho}_{k-1}, \boldsymbol{\psi}_{k-1}, \boldsymbol{C}_k, \beta, h, \epsilon, N$, the Algorithm 1 requires two parameters as user input: numerical tolerance δ, and maximum number of iterations L. The computation in Algorithm 1, as presented, involves making an initial guess for the vector $\boldsymbol{z}$ and then updating $\boldsymbol{y}$ and $\boldsymbol{z}$ until convergence. The iteration over index $\ell \leq L$ is performed while keeping the physical time "frozen".

We next outline the overall algorithmic setup to implement the proximal recursion over probability weighted scattered data.

Overall algorithm

Samples from the *known* initial joint PDF ρ_0 are generated as point cloud $\{\boldsymbol{x}_0^i, \varrho_0^i\}_{i=1}^N$. Then for $k \in \mathbb{N}$, the point clouds $\{\boldsymbol{x}_k^i, \varrho_k^i\}_{i=1}^N$ are updated as shown in Fig. 2.2. Specifically, the state vectors are updated via Euler-Maruyama scheme applied to the underlying SDE (2.13) or (2.19). Explicitly, the Euler-Maruyama update for

Algorithm 1 Proximal recursion to compute $\boldsymbol{\varrho}_k$ from $\boldsymbol{\varrho}_{k-1}$

procedure ProxRecur($\boldsymbol{\varrho}_{k-1}$, $\boldsymbol{\psi}_{k-1}$, $\boldsymbol{C}_k$, β, h, ϵ, N, δ, L)
  $\boldsymbol{\Gamma}_k \leftarrow \exp(-\boldsymbol{C}_k/2\epsilon)$
  $\boldsymbol{\xi} \leftarrow \exp(-\beta\boldsymbol{\psi}_{k-1} - \mathbf{1})$
  $\boldsymbol{z}_0 \leftarrow \text{rand}_{N\times 1}$ ▷ initialize
  $\boldsymbol{z} \leftarrow \left[\boldsymbol{z}_0, \mathbf{0}_{N\times(L-1)}\right]$
  $\boldsymbol{y} \leftarrow \left[\boldsymbol{\varrho}_{k-1} \oslash (\boldsymbol{\Gamma}_k \boldsymbol{z}_0), \mathbf{0}_{N\times(L-1)}\right]$
  $\ell = 1$ ▷ iteration index
  while $\ell \leq L$ **do**
    $\boldsymbol{z}(:, \ell+1) \leftarrow \left(\boldsymbol{\xi}_{k-1} \oslash \left(\boldsymbol{\Gamma}_k^\top \boldsymbol{y}(:,\ell)\right)\right)^{\frac{1}{1+\beta\epsilon/h}}$
    $\boldsymbol{y}(:, \ell+1) \leftarrow \boldsymbol{\varrho}_{k-1} \oslash (\boldsymbol{\Gamma}_k \boldsymbol{z}(:,\ell+1))$
    if $\| \boldsymbol{y}(:,\ell+1) - \boldsymbol{y}(:,\ell) \| < \delta$ & $\| \boldsymbol{z}(:,\ell+1) - \boldsymbol{z}(:,\ell) \| < \delta$ **then** ▷ error within tolerance
      break
    else
      $\ell \leftarrow \ell + 1$
    end if
  end while
  return $\boldsymbol{\varrho}_k \leftarrow \boldsymbol{z}(:,\ell) \odot \left(\boldsymbol{\Gamma}_k^\top \boldsymbol{y}(:,\ell)\right)$ ▷ proximal update
end procedure

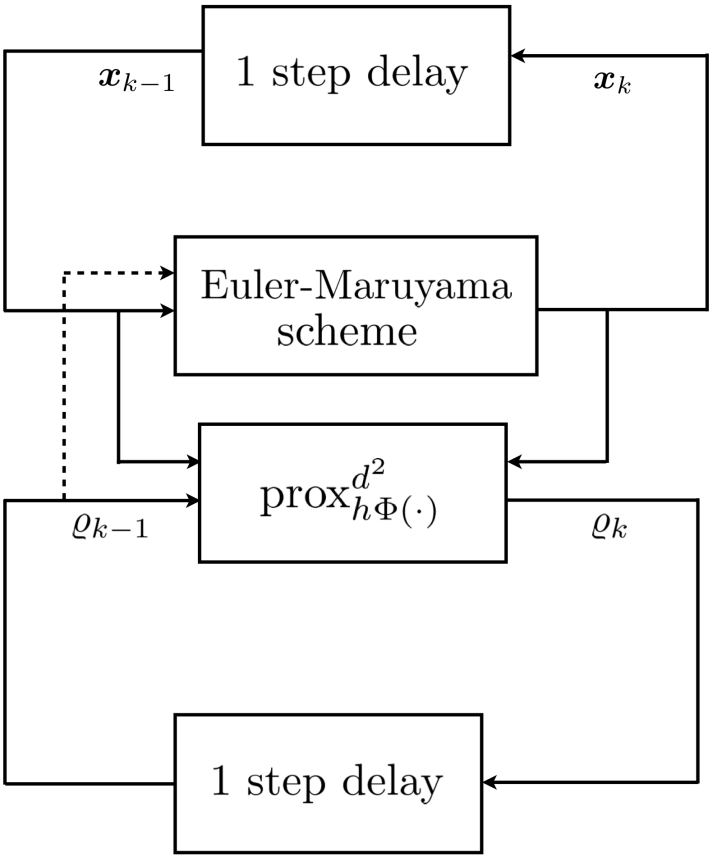

Figure 2.2: Schematic of the proposed algorithmic setup for propagating the joint state PDF as probability weighted scattered point cloud $\{\boldsymbol{x}_k^i, \varrho_k^i\}_{i=1}^N$. The location of the points $\{\boldsymbol{x}_k^i\}_{i=1}^N$ are updated via Euler-Maruyama scheme; the corresponding probability weights are updated via Algorithm 1. The dashed arrow shown above is present only when the state dynamics is density dependent, as in (2.19).

(2.19) is

$$\begin{aligned}\boldsymbol{x}_k^i = \boldsymbol{x}_{k-1}^i - h\nabla\left(\psi(\boldsymbol{x}_{k-1}^i) + \omega(\boldsymbol{x}_{k-1}^i)\right)\\ +\sqrt{2\beta^{-1}}\left(\boldsymbol{w}_k^i - \boldsymbol{w}_{k-1}^i\right),\end{aligned} \tag{2.39}$$

where $\omega(\cdot) := \int \phi(\cdot - \boldsymbol{y})\rho(\boldsymbol{y})\mathrm{d}\boldsymbol{y}$, and $\boldsymbol{w}_k^i := \boldsymbol{w}^i(t = kh)$, $k \in \mathbb{N}$. The same for (2.13) is obtained by setting $\phi \equiv \omega \equiv 0$ in (2.39).

The corresponding probability weights $\{\boldsymbol{\varrho}_k^i\}_{i=1}^N$ are updated via Algorithm 1. No-

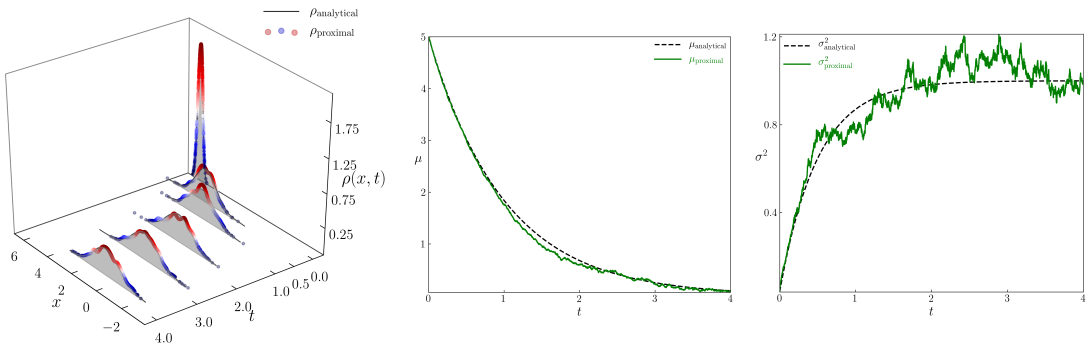

Figure 2.3: Comparison of the analytical and proximal solutions of the FPK PDE for (2.45) with time step $h = 10^{-3}$, and with parameters $a = 1$, $\beta = 1$, $\epsilon = 5 \times 10^{-2}$. Shown above are the time evolution of the (*left*) PDFs, (*middle*) means, and (*right*) variances.

tice that computing $\boldsymbol{C}_k$ requires both $\{\boldsymbol{x}_{k-1}^i\}_{i=1}^N$ and $\{\boldsymbol{x}_k^i\}_{i=1}^N$, and that $\boldsymbol{C}_k$ needs to be passed as input to Algorithm 1. Thus, the execution of Euler-Maruyama scheme precedes that of Algorithm 1.

Remark 2. *Our choice of the (explicit) Euler-Maruyama scheme for updating the location of the points in state space was motivated by its simplicity and ease of implementation. Since the diffusion coefficient of (2.13) or (2.19) is constant, the Euler-Maruyama scheme is guaranteed to converge strongly to the true solution of the corresponding SDE provided the drift coefficient is globally Lipschitz; see e.g., [133, Ch. 10.2]. If the drift coefficient $-\nabla\psi$ in (2.13), or $-(\nabla\psi + \nabla(\rho * \phi))$ in (2.19), is non-globally Lipschitz with superlinear growth, then the Euler-Maruyama scheme is known [119] to diverge in mean-squared error sense. In such cases, one could replace the explicit Euler-Maruyama scheme with the "tamed" Euler-Maruyama scheme [120], or with the partially implicit Euler scheme [118], or with the split-step backward Euler scheme [112].*

2.3.3 Convergence

Next, we will prove the convergence properties for Algorithm 1. To this end, the Definition 3 and Proposition 2 given below will be useful in establishing Theorem 3 that follows.

Definition 3. *(**Thompson metric**) Consider $\boldsymbol{z}, \widetilde{\boldsymbol{z}} \in \mathcal{K}$, where $\mathcal{K}$ is a non-empty open convex cone. Further, suppose that $\mathcal{K}$ is a normal cone, i.e., there exists constant α such that $\| \boldsymbol{z} \| \leq \alpha \| \widetilde{\boldsymbol{z}} \|$ for $\boldsymbol{z} \leq \widetilde{\boldsymbol{z}}$. Thompson [207] proved that $\mathcal{K}$ is a complete metric space w.r.t. the so-called* Thompson metric *given by*

$$d_{\mathrm{T}}(\boldsymbol{z}, \widetilde{\boldsymbol{z}}) := \max\{\log \gamma(\boldsymbol{z}/\widetilde{\boldsymbol{z}}), \log \gamma(\widetilde{\boldsymbol{z}}/\boldsymbol{z})\},$$

where $\gamma(\boldsymbol{z}/\widetilde{\boldsymbol{z}}) := \inf\{c > 0 \mid \boldsymbol{z} \leq c\widetilde{\boldsymbol{z}}\}$. In particular, if $\mathcal{K} \equiv \mathbb{R}^N_{>0}$ (positive orthant of $\mathbb{R}^N$), then

$$d_{\mathrm{T}}(\boldsymbol{z}, \widetilde{\boldsymbol{z}}) = \log \max \left\{ \max_{i=1,\ldots,N} \left(\frac{\boldsymbol{z}_i}{\widetilde{\boldsymbol{z}}_i} \right), \max_{i=1,\ldots,N} \left(\frac{\widetilde{\boldsymbol{z}}_i}{\boldsymbol{z}_i} \right) \right\}. \tag{2.40}$$

Proposition 2. *[149, Proposition 3.2], [167] Let $\mathcal{K}$ be an open, normal, convex cone, and let $\boldsymbol{p} : \mathcal{K} \mapsto \mathcal{K}$ be an order preserving homogeneous map of degree $r \geq 0$, i.e., $\boldsymbol{p}(c\boldsymbol{z}) = c^r \boldsymbol{p}(\boldsymbol{z})$ for any $c > 0$ and $\boldsymbol{z} \in \mathcal{K}$. Then, for all $\boldsymbol{z}, \widetilde{\boldsymbol{z}} \in \mathcal{K}$, we have*

$$d_{\mathrm{T}}(\boldsymbol{p}(\boldsymbol{z}), \boldsymbol{p}(\widetilde{\boldsymbol{z}})) \leq r d_{\mathrm{T}}(\boldsymbol{z}, \widetilde{\boldsymbol{z}}).$$

In particular, if $r \in [0, 1)$, then the map $\boldsymbol{p}(\cdot)$ is strictly contractive in the Thompson metric d_{T}, and admits unique fixed point in $\mathcal{K}$.

Using (2.40) and Proposition 2, we establish the convergence result below.

Theorem 3. *Consider the notations in (2.31)-(2.32), and those in Algorithm 1.*

The iteration

$$\begin{aligned}\boldsymbol{z}(:,\ell+1) &= \left(\boldsymbol{\xi}_{k-1} \oslash \left(\boldsymbol{\Gamma}_k^\top \boldsymbol{y}(:,\ell)\right)\right)^{\frac{1}{1+\beta\epsilon/h}} \\ &= \left(\boldsymbol{\xi}_{k-1} \oslash \left(\boldsymbol{\Gamma}_k^\top \boldsymbol{\varrho}_{k-1} \oslash (\boldsymbol{\Gamma}_k \boldsymbol{z}(:,\ell))\right)\right)^{\frac{1}{1+\beta\epsilon/h}} \end{aligned} \tag{2.41}$$

for $\ell = 1, 2, \ldots$, is strictly contractive in the Thompson metric (2.40) on $\mathbb{R}^N_{>0}$, and admits unique fixed point $\boldsymbol{z}^{\text{opt}} \in \mathbb{R}^N_{>0}$.

Proof. Rewriting (2.41) as

$$\boldsymbol{z}(:,\ell+1) = \left(\left(\boldsymbol{\xi}_{k-1} \oslash \left(\boldsymbol{\Gamma}_k^\top \boldsymbol{\varrho}_{k-1}\right)\right) \odot (\boldsymbol{\Gamma}_k \boldsymbol{z}(:,\ell))\right)^{\frac{1}{1+\beta\epsilon/h}},$$

and letting $\boldsymbol{\eta} \equiv \boldsymbol{\eta}_{k,k+1} := \boldsymbol{\xi}_{k-1} \oslash \left(\boldsymbol{\Gamma}_k^\top \boldsymbol{\varrho}_{k-1}\right)$, we notice that iteration (2.41) can be expressed as a cone preserving composite map $\boldsymbol{\theta} := \boldsymbol{\theta}_1 \circ \boldsymbol{\theta}_2 \circ \boldsymbol{\theta}_3$, where $\boldsymbol{\theta} : \mathbb{R}^N_{>0} \mapsto \mathbb{R}^N_{>0}$, given by

$$\boldsymbol{z}(:,\ell+1) = \boldsymbol{\theta}\left(\boldsymbol{z}(:,\ell)\right) = \boldsymbol{\theta}_1 \circ \boldsymbol{\theta}_2 \circ \boldsymbol{\theta}_3 \ \left(\boldsymbol{z}(:,\ell)\right), \tag{2.42}$$

and $\boldsymbol{\theta}_1(\boldsymbol{z}) := \boldsymbol{z}^{\frac{1}{1+\beta\epsilon/h}}$, $\boldsymbol{\theta}_2(\boldsymbol{z}) := \boldsymbol{\eta} \odot \boldsymbol{z}$, $\boldsymbol{\theta}_3 := \boldsymbol{\Gamma}_k \boldsymbol{z}$. Our strategy is to prove that the composite map $\boldsymbol{\theta}$ is contractive on $\mathbb{R}^N_{>0}$ w.r.t. the metric d_{T}.

From (2.25) and (2.32), $\boldsymbol{C}_k(i,j) \in [0,\infty)$ which implies $\boldsymbol{\Gamma}_k(i,j) \in (0,1]$; therefore, $\boldsymbol{\Gamma}_k$ is a positive linear map for each $k = 1, 2, \ldots$. Thus, by (linear) Perron-Frobenius theorem, the map $\boldsymbol{\theta}_3$ is contractive on $\mathbb{R}^N_{>0}$ w.r.t. d_{T}. The map $\boldsymbol{\theta}_2$ is an isometry by Definition 3. As for the map $\boldsymbol{\theta}_1$, notice that the quantity $r := 1/(1+\beta\epsilon/h) \in (0,1)$ since $\beta\epsilon/h > 0$. Therefore, the map $\boldsymbol{\theta}_1(\boldsymbol{z}) := \boldsymbol{z}^r$ (element-wise exponentiation) is monotone (order preserving) and homogeneous of degree $r \in (0,1)$ on $\mathbb{R}^N_{>0}$. By Proposition 2, the map $\boldsymbol{\theta}_1(\boldsymbol{z})$ is strictly contractive.

Thus, the composition

$$\boldsymbol{\theta} = \underbrace{\boldsymbol{\theta}_1}_{\text{strictly contractive}} \circ \underbrace{\boldsymbol{\theta}_2}_{\text{isometry}} \circ \underbrace{\boldsymbol{\theta}_3}_{\text{contractive}}$$

is strictly contractive w.r.t. d_{T}, and (by Banach contraction mapping theorem) admits unique fixed point $\boldsymbol{z}^{\mathrm{opt}}$ in $\mathbb{R}^N_{>0}$. ■

Corollary 4. *The Algorithm 1 converges to unique fixed point* $(\boldsymbol{y}^{\mathrm{opt}}, \boldsymbol{z}^{\mathrm{opt}}) \in \mathbb{R}^N_{>0} \times \mathbb{R}^N_{>0}$.

Proof. Since $\boldsymbol{y}(:, \ell+1) = \boldsymbol{\varrho}_{k-1} \oslash (\boldsymbol{\Gamma}_k \boldsymbol{z}(:, \ell+1))$, the $\boldsymbol{z}$ iterates converge to unique fixed point $\boldsymbol{z}^{\mathrm{opt}} \in \mathbb{R}^N_{>0}$ (by Theorem 3), and the linear maps $\boldsymbol{\Gamma}_k$ are contractive (by Perron-Frebenius theory, as before), consequently the $\boldsymbol{y}$ iterates also converge to unique fixed point $\boldsymbol{y}^{\mathrm{opt}} \in \mathbb{R}^N_{>0}$. Hence the statement. ■

2.4 Numerical Results

We now illustrate the computational framework proposed in Section 2.3 via numerical examples. Our examples involve systems already in JKO canonical form (Section 2.2), as well as those which can be transformed to such form by non-obvious change of coordinates.

2.4.1 Linear Gaussian Systems

For an Itô SDE of the form

$$\mathrm{d}\boldsymbol{x} = \boldsymbol{A}\boldsymbol{x}\,\mathrm{d}t + \boldsymbol{B}\,\mathrm{d}\boldsymbol{w}, \tag{2.43}$$

it is well known that if $\boldsymbol{x}_0 := \boldsymbol{x}(t=0) \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$, then the transient joint PDFs $\rho(\boldsymbol{x}, t) = \mathcal{N}(\boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t))$ where the vector-matrix pair $(\boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t))$ evolve according to the ODEs

$$\dot{\boldsymbol{\mu}}(t) = \boldsymbol{A}\boldsymbol{\mu}, \quad \boldsymbol{\mu}(0) = \boldsymbol{\mu}_0, \tag{2.44a}$$

$$\dot{\boldsymbol{\Sigma}}(t) = \boldsymbol{A}\boldsymbol{\Sigma}(\boldsymbol{t}) + \boldsymbol{A}\boldsymbol{\Sigma}(\boldsymbol{t})^\top + \boldsymbol{B}\boldsymbol{B}^\top, \quad \boldsymbol{\Sigma}(0) = \boldsymbol{\Sigma}_0. \tag{2.44b}$$

We benchmark the numerical results produced by the proposed proximal algorithm vis-à-vis the solutions of (2.44). We consider the following two sub-cases of (2.43).

Ornstein-Uhlenbeck Process

We consider the univariate system

$$\mathrm{d}x = -ax\,\mathrm{d}t + \sqrt{2\beta^{-1}}\mathrm{d}w, \quad a, \beta > 0, \tag{2.45}$$

which is in JKO form (2.13) with $\psi(x) = \frac{1}{2}ax^2$. We generate $N = 400$ samples from the initial PDF $\rho_0 = \mathcal{N}(\mu_0, \sigma_0^2)$ with $\mu_0 = 5$ and $\sigma_0^2 = 4 \times 10^{-2}$, and apply the proposed proximal recursion for (2.45) with time step $h = 10^{-3}$, and with parameters $a = 1$, $\beta = 1$, $\epsilon = 5 \times 10^{-2}$. For implementing Algorithm 1, we set the tolerance $\delta = 10^{-3}$, and the maximum number of iterations $L = 100$. Fig. 2.3 shows that the PDF point clouds generated by the proximal recursion match with the analytical PDFs $\mathcal{N}\left(\mu_0 \exp(-at), (\sigma_0^2 - \frac{1}{a\beta})\exp(-2at) + \frac{1}{a\beta}\right)$, and the mean-variance trajectories (computed from the numerical integration of the weighted scattered point cloud data) match with the corresponding analytical solutions.

Multivariate LTI

We next consider the multivariate case (2.43) where the pair $(\boldsymbol{A}, \boldsymbol{B})$ is assumed to be controllable, and the matrix $\boldsymbol{A}$ is Hurwitz (not necessarily symmetric). Under these assumptions, the stationary PDF is $\mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_\infty)$ where $\boldsymbol{\Sigma}_\infty$ is the unique stationary solution of (2.44b) that is guaranteed to be symmetric positive definite. However, it is not apparent whether (2.43) can be expressed in the form (2.13), since for non-symmetric $\boldsymbol{A}$, there does *not* exist constant symmetric positive definite matrix $\boldsymbol{\Psi}$ such that $\boldsymbol{A}\boldsymbol{x} = -\nabla \boldsymbol{x}^\top \boldsymbol{\Psi} \boldsymbol{x}$, i.e., the drift vector field does not admit a natural potential. Thus, implementing the JKO scheme for (2.43) is non-trivial in general.

In our recent work [100], two successive *time-varying* co-ordinate transformations were proposed which can bring (2.43) in the form (2.13), thus making it amenable to the JKO scheme. We apply these change-of-coordinates to (2.43) with

$$\boldsymbol{A} = \begin{pmatrix} -10 & 5 \\ -30 & 0 \end{pmatrix}, \quad \boldsymbol{B} = \begin{pmatrix} 2 \\ 2.5 \end{pmatrix},$$

which satisfy the stated assumptions for the pair $(\boldsymbol{A}, \boldsymbol{B})$, and implement the proposed proximal recursion on this transformed co-ordinates with $N = 400$ samples generated from the initial PDF $\rho_0 = \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$, where $\boldsymbol{\mu}_0 = (4, 4)^\top$ and $\boldsymbol{\Sigma}_0 = 4\boldsymbol{I}_2$. As before, we set $\delta = 10^{-3}, L = 100, h = 10^{-3}, \beta = 1, \epsilon = 5 \times 10^{-2}$. Once the proximal updates are done, we transform back the probability weighted scattered point cloud to the original state space co-ordinates via change-of-measure formula associated with the known co-ordinate transforms [100, Section III.B]. Fig. 2.4 shows the resulting point clouds superimposed with the contour plots for the analytical solutions $\mathcal{N}(\boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t))$ given by (2.44). Figs. 2.5 and 2.6 compare

the respective mean and covariance evolution. We point out that the change of co-ordinates in [100] requires implementing the JKO scheme in a time-varying rotating frame (defined via exponential of certain time varying skew-symmetric matrix) that depends on the stationary covariance $\mathbf{\Sigma}_\infty$. As a consequence, the stationary covariance resulting from the proximal recursion oscillates about the true stationary value.

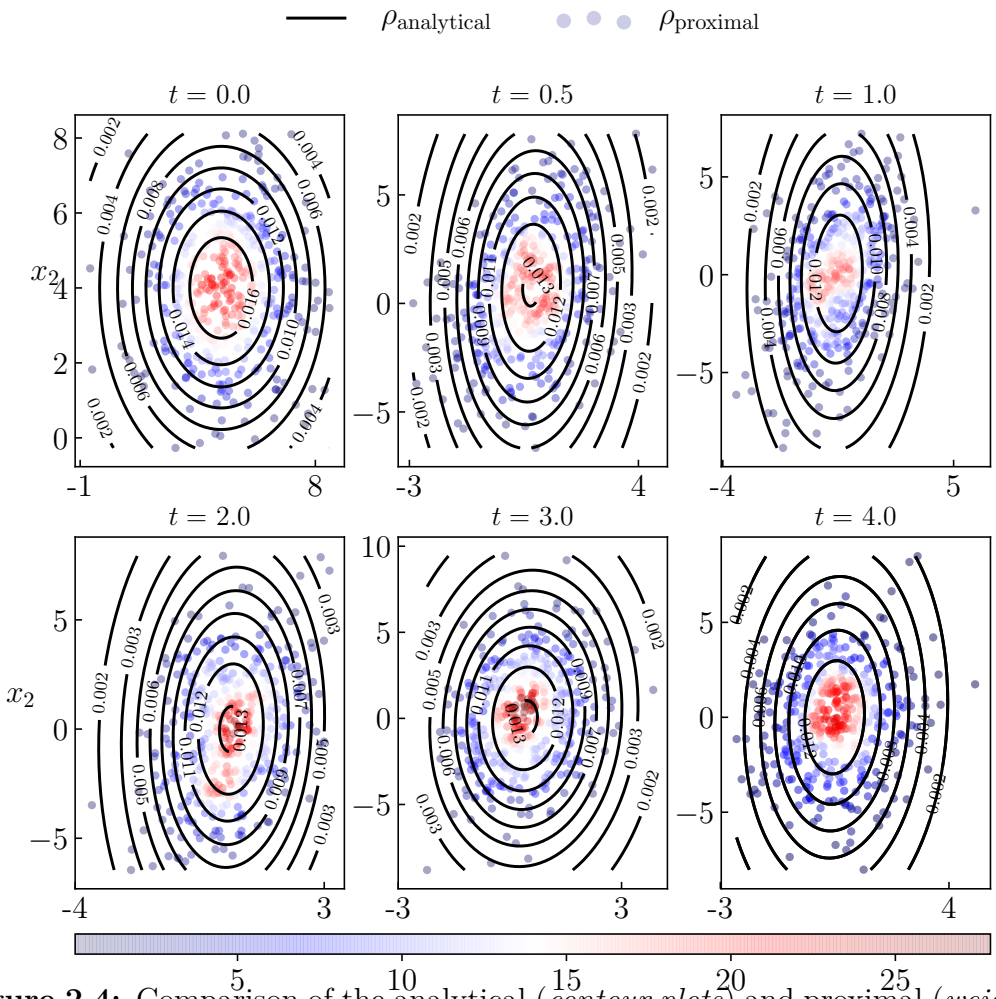

Figure 2.4: Comparison of the analytical (*contour plots*) and proximal (*weighted scattered point cloud*) joint PDFs of the FPK PDE for (2.43) with time step $h = 10^{-3}$, and with parameters $\beta = 1, \epsilon = 5 \times 10^{-2}$. Simulation details are given in Section 2.4.1. The color (*red = high, blue = low*) denotes the joint PDF value obtained via proximal recursion at a point at that time (see colorbar).

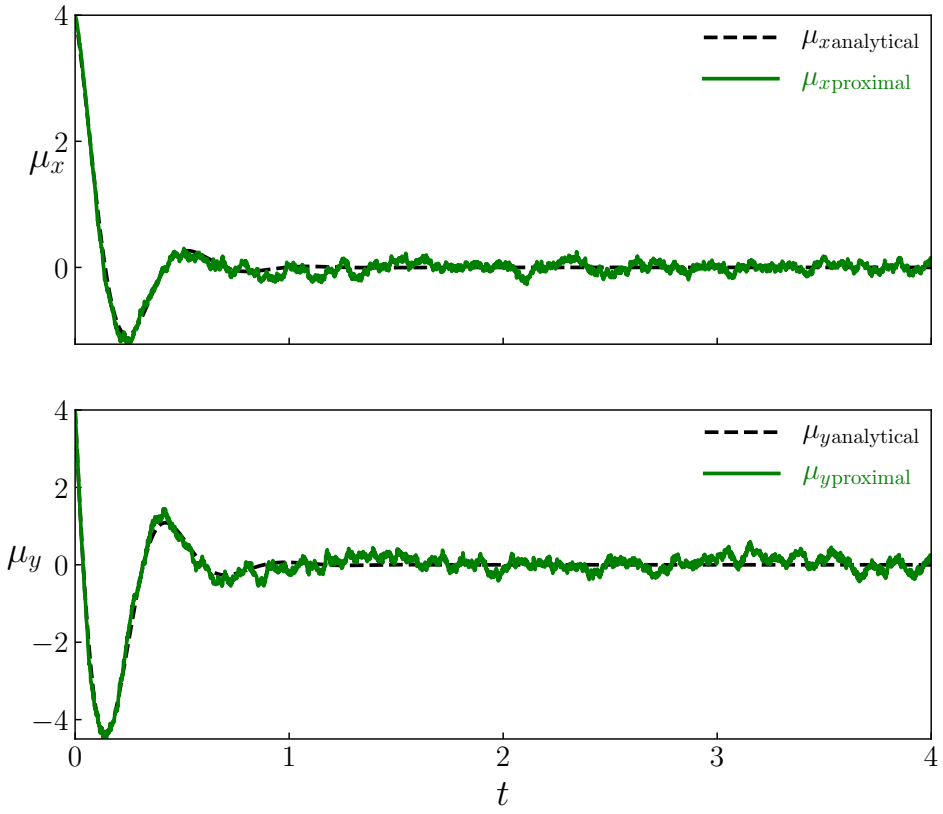

Figure 2.5: Comparison of the components of the mean vectors from analytical (*dashed*) and proximal (*solid*) computation of the joint PDFs for (2.43) with time step $h = 10^{-3}$, and with parameters $\beta = 1, \epsilon = 5 \times 10^{-2}$. Simulation details are given in Section 2.4.1.

2.4.2 Nonlinear Non-Gaussian System

Next we consider a planar nonlinear system of the form (2.13) with $\psi(x_1, x_2) = \frac{1}{4}(1 + x_1^4) + \frac{1}{2}(x_2^2 - x_1^2)$ (see Fig. 2.7). As mentioned in Section 2.2, the stationary PDF is $\rho_\infty(\boldsymbol{x}) = \kappa \exp(-\beta\psi(\boldsymbol{x}))$, which for our choice of ψ, is bimodal. In this case, the transient PDFs have no known analytical solution but can be computed using the proposed proximal recursion. For doing so, we generate $N = 400$ samples from the initial PDF $\rho_0 = \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ with $\boldsymbol{\mu}_0 = (2, 2)^\top$ and $\boldsymbol{\Sigma}_0 = 4\boldsymbol{I}_2$, and set $\delta = 10^{-3}, L = 100, h = 10^{-3}, \beta = 1, \epsilon = 5 \times 10^{-2}$, as before. The resulting weighted point clouds are shown in Fig. 2.8; it can be seen that as time progresses, the

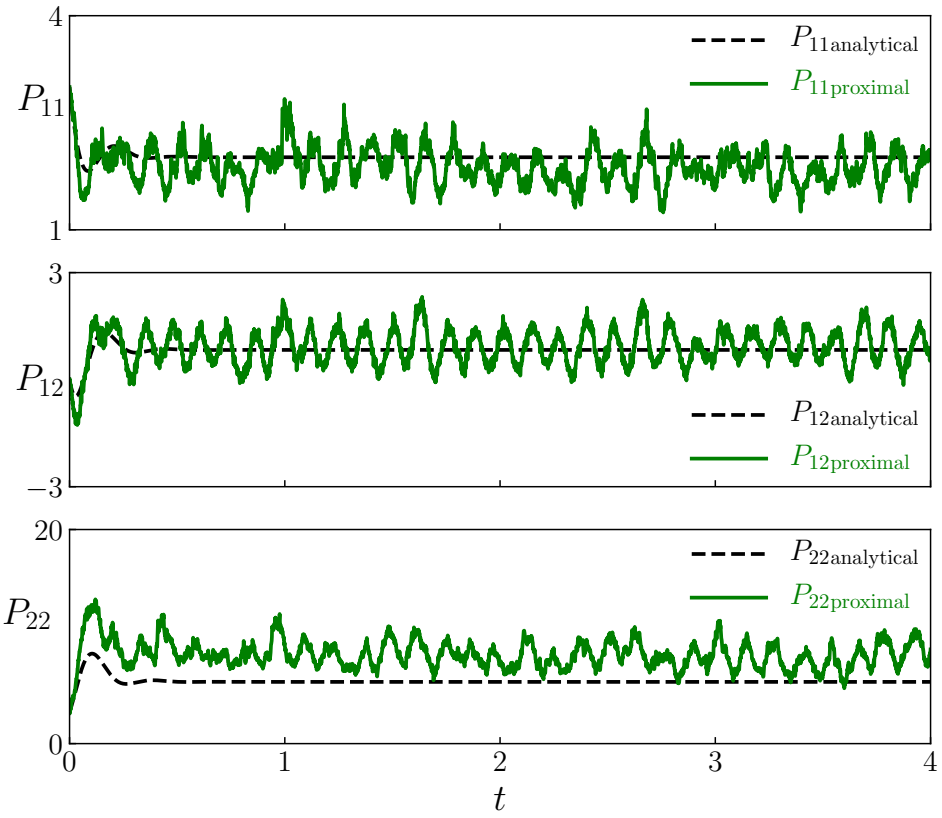

Figure 2.6: Comparison of the components of the covariance matrices from analytical (*dashed*) and proximal (*solid*) computation of the joint PDFs for (2.43) with time step $h = 10^{-3}$, and with parameters $\beta = 1, \epsilon = 5 \times 10^{-2}$. Simulation details are given in Section 2.4.1.

joint PDFs computed via the proximal recursion, tend to the known stationary solution ρ_∞ (contour plots in the right bottom sub-figure in Fig. 2.8).

Fig. 2.9 shows the computational times for the proposed proximal recursions applied to the above nonlinear non-Gaussian system. Since the proposed algorithm involves sub-iterations ("while loop" in Algorithm 1 over index $\ell \leq L$) while keeping the physical time "frozen", the convergence reported in Section 2.3.3 must be achieved at "sub-physical time step" level, i.e., must incur smaller than h (here, $h = 10^{-3}$ s) computational time. Indeed, Fig. 2.9 shows that each proximal update takes approx. 10^{-6} s, or $10^{-3}h$ computational time, which demonstrates the efficacy of the proposed framework.

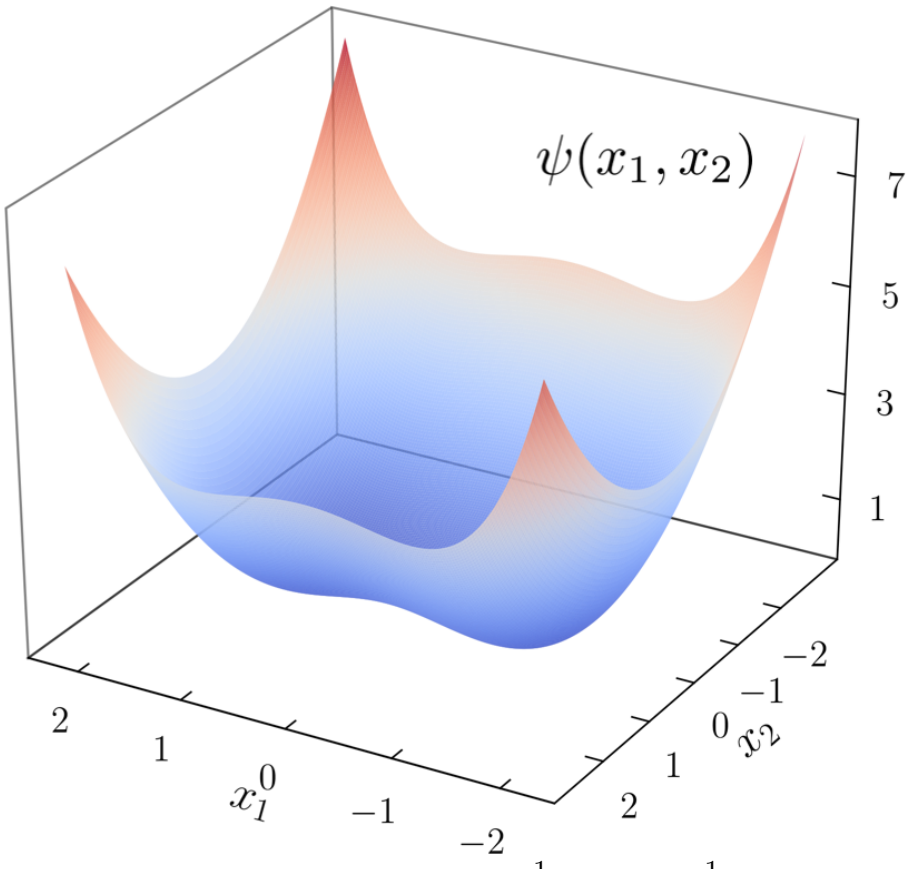

Figure 2.7: The drift potential $\psi(x_1, x_2) = \frac{1}{4}(1 + x_1^4) + \frac{1}{2}(x_2^2 - x_1^2)$ used in the numerical example given in Section 2.4.2.

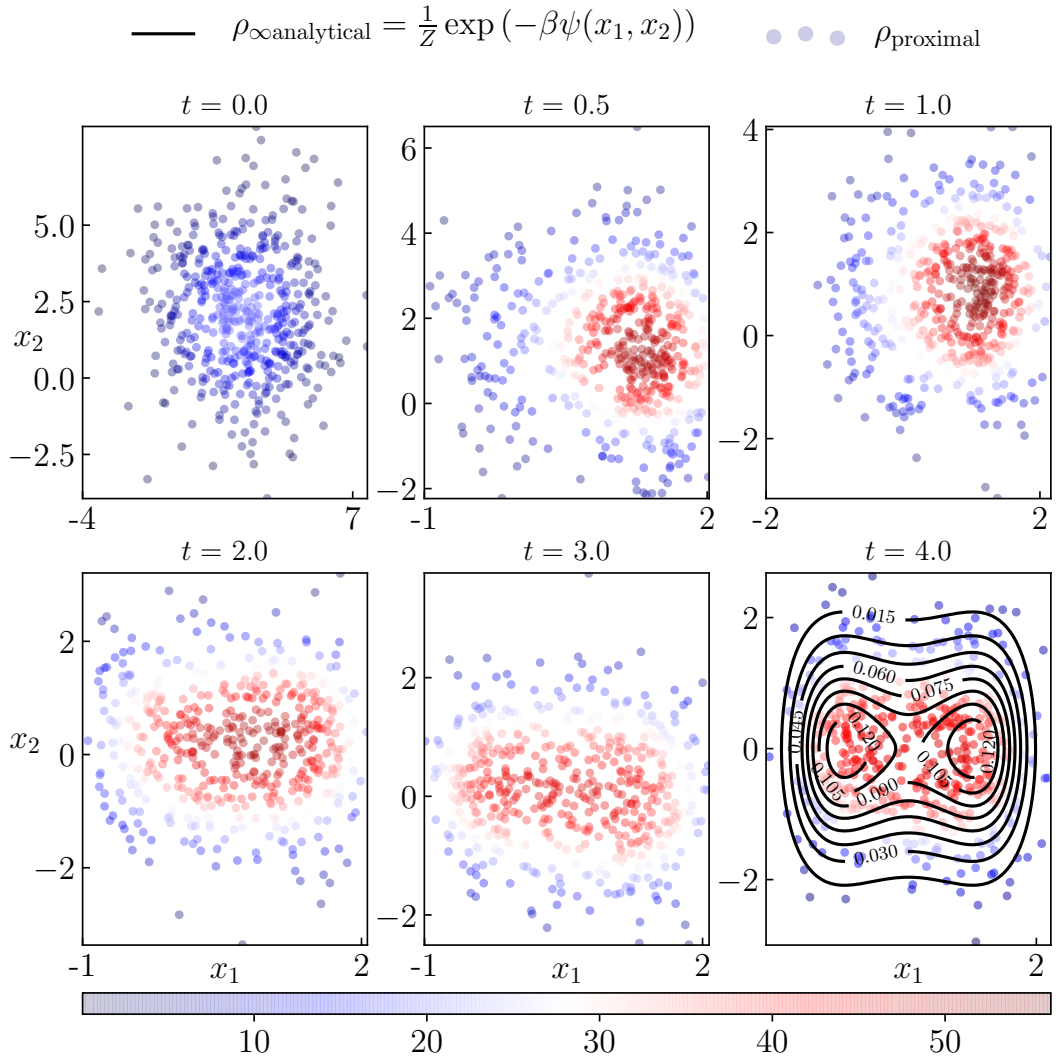

Figure 2.8: The proximal (*weighted scattered point cloud*) joint PDFs of the FPK PDE (2.14) with the drift potential shown in Fig. 2.7, time step $h = 10^{-3}$, and with parameters $\beta = 1, \epsilon = 5 \times 10^{-2}$. Simulation details are given in Section 2.4.2. The color (*red = high, blue = low*) denotes the joint PDF value obtained via proximal recursion at a point at that time (see colorbar). In the bottom right plot, the contour lines correspond to the analytical solution for the stationary PDF ρ_∞.

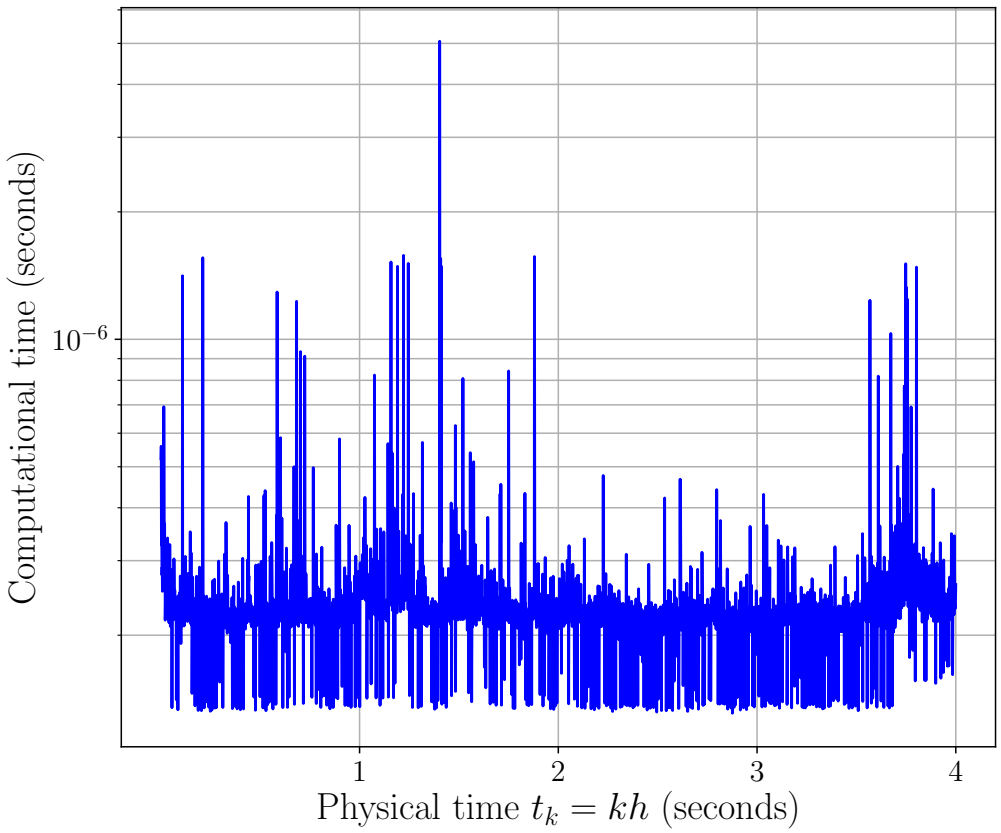

Figure 2.9: The computational times for proximal updates for the simulation in Section 2.4.2. Here, the physical time-step $h = 10^{-3}$ s, and $k \in \mathbb{N}$.

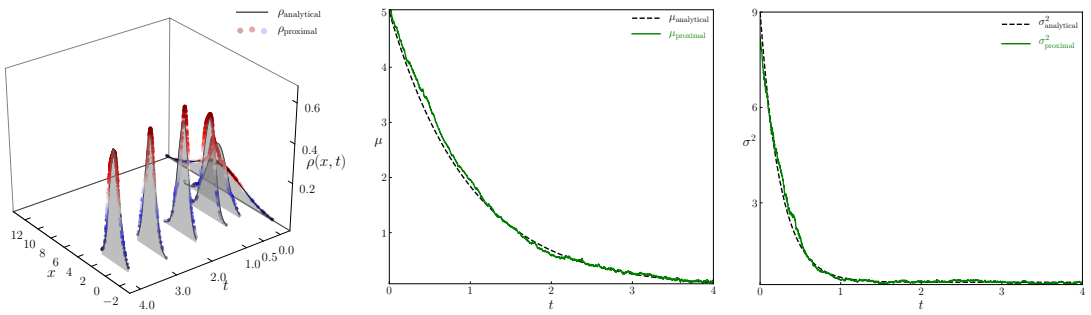

Figure 2.10: Comparison of the analytical and proximal solutions of the McKean-Vlasov flow for (2.46) with time step $h = 10^{-3}$, $\rho_0 = \mathcal{N}(5, 9)$, and with parameters $a = b = 1$, $\beta = 1$, $\epsilon = 5 \times 10^{-2}$. Shown above are the time evolution of the transient (*left*) PDFs, (*middle*) means, and (*right*) variances.

2.4.3 Non-local Interactions

We now consider a numerical example demonstrating our gradient flow framework for computing the transient PDFs generated by the McKean-Vlasov integro-PDE (2.18) associated with the density-dependent sample path dynamics (2.19). If both the potentials ψ and ϕ are convex functions, then (2.18) admits unique stationary density $\rho_\infty(\boldsymbol{x})$ (see e.g., [46], [154, Section 5.2]). To keep the exposition simple, we consider the univariate case $\psi(x) = \frac{1}{2}ax^2$, $\phi(x) = \frac{1}{2}bx^2$, $a, b > 0$. Performing the integration appearing in the convolution allows us to rewrite (2.19) as the mean-reverting process:

$$\mathrm{d}x = -\left((a+b)x - b\mu(t)\right)\,\mathrm{d}t \,+\, \sqrt{2\beta^{-1}}\,\mathrm{d}w, \quad x(0) = x_0, \tag{2.46}$$

where $\mu(t)$ is the mean of the transient PDF $\rho(x,t)$. Assuming $x_0 \sim \mathcal{N}(\mu_0, \sigma_0^2)$, and applying expectation operator to both sides of (2.46) yields $\mu(t) = \mu_0 \exp(-at)$. Consequently, the transient PDF for (2.46) at time t is $\rho(x,t) = \mathcal{N}\left(\mu(t), \sigma^2(t)\right)$,

with

$$\mu(t) = \mu_0 \exp(-at), \tag{2.47a}$$

$$\sigma^2(t) = \left(\sigma_0^2 - \frac{1}{(a+b)\beta}\right) \exp\left(-2(a+b)t\right) + \frac{1}{(a+b)\beta}. \tag{2.47b}$$

Clearly, the stationary PDF is $\rho_\infty(x) = \mathcal{N}\left(0, 1/(a+b)\beta\right)$.

To benchmark our algorithm with the analytical solution (2.47), we implement the proximal recursion for (2.46) with free energy (2.20). Following [23, Section 4], we replace the non-convex bilinear term $\int_{\mathbb{R}^n \times \mathbb{R}^n} \phi(\boldsymbol{x}-\boldsymbol{y})\varrho(\boldsymbol{x})\varrho(\boldsymbol{y})\mathrm{d}\boldsymbol{x}\mathrm{d}\boldsymbol{y}$ in (2.22b), with the linear term $\int_{\mathbb{R}^n} \phi(\boldsymbol{x}-\boldsymbol{y})\varrho(\boldsymbol{x})\varrho_{k-1}(\boldsymbol{y})\mathrm{d}\boldsymbol{x}\mathrm{d}\boldsymbol{y}$, $k \in \mathbb{N}$, resulting in a semi-implicit variant of (2.22), given by

$$\varrho_k = \underset{\varrho \in D_2}{arg\,inf}\ \frac{1}{2}W^2(\varrho_{k-1}, \varrho) + h\, F(\varrho_{k-1}, \varrho), \quad k \in \mathbb{N}, \tag{2.48}$$

with $\varrho_0 \equiv \rho_0(\boldsymbol{x})$ (the initial PDF). The proof for the fact that such a semi-implicit scheme guarantees $\varrho_k(\boldsymbol{x}) \to \rho(\boldsymbol{x}, t = kh)$ for $h \downarrow 0$, where $\rho(\boldsymbol{x}, t)$ is the flow generated by (2.18), can be found in [138, Section 12.3]. Notice that for the FPK gradient flow, the "discrete free energy" in (2.23) was $F(\boldsymbol{\varrho}) = \langle \boldsymbol{\psi}_{k-1} + \beta^{-1}\log \boldsymbol{\varrho}, \boldsymbol{\varrho}\rangle$. The scheme (2.48) allows us to write a similar expression in the McKean-Vlasov gradient flow case, as $F(\boldsymbol{\varrho}) = \langle \boldsymbol{\psi}_{k-1} + \boldsymbol{D}_{k-1}\boldsymbol{\varrho}_{k-1} + \beta^{-1}\log \boldsymbol{\varrho}, \boldsymbol{\varrho}\rangle$, where the symmetric matrix $\boldsymbol{D}_{k-1}$ is given by

$$\boldsymbol{D}_{k-1}(i,j) = \phi\left(\boldsymbol{x}_{k-1}^i - \boldsymbol{x}_{k-1}^j\right), \quad i,j = 1,\dots,N, \quad k \in \mathbb{N}.$$

For the particular choice $\phi(x) = \frac{1}{2}bx^2$, notice that $\boldsymbol{D}_{k-1}$ is a (scaled) Euclidean distance matrix on the Euler-Maruyama update. In this case, the associated

Euler-Maruyama scheme (2.39) is

$$\begin{aligned} \boldsymbol{x}_k^i &= \boldsymbol{x}_{k-1}^i - h\left(a\boldsymbol{x}_{k-1}^i + b\left(\boldsymbol{x}_{k-1}^i - \mathbf{1}^\top \boldsymbol{\varrho}_{k-1}\right)\right) \\ &+ \sqrt{2\beta^{-1}}\left(\boldsymbol{w}_k^i - \boldsymbol{w}_{k-1}^i\right), \quad k \in \mathbb{N},\, i = 1, \ldots, N, \end{aligned} \tag{2.49}$$

i.e., the dashed arrow in Fig. 2.2 becomes active. Fig. 2.10 shows that the weighted scattered point cloud solutions for (2.46) with $a = b = 1$, computed through our proximal algorithm match with the analytical solutions $\mathcal{N}(\mu(t), \sigma^2(t))$ given by (2.47). For Fig. 2.10, the parameter values used in our simulation are $h = 10^{-3}$, $\beta = 1$, $\epsilon = 5 \times 10^{-2}$, $\mu_0 = 5$, $\sigma_0^2 = 9$, $N = 400$, $\delta = 10^{-3}$, $L = 100$.

2.5 Extensions

In Section 2.4.1, we have already seen that systems not in JKO canonical form may be transformed to the same via suitable change-of-coordinates, thus making density propagation for such systems amenable via our framework. In this Section, we provide two extensions along these lines. First, we consider a case of state-dependent diffusion; thereafter, we consider a system with mixed conservative-dissipative drift. In both cases, we use specific examples (instead of general remarks) to help illustrate the extensions of the basic framework. These examples point out the broad scope of the algorithms proposed herein.

2.5.1 Multiplicative Noise

We consider the Itô SDE for the Cox-Ingersoll-Ross (CIR) model [72], given by

$$\mathrm{d}x = a(\theta - x)\,\mathrm{d}t \;+\; b\sqrt{x}\,\mathrm{d}w, \quad 2a > b^2 > 0, \quad \theta > 0. \tag{2.50}$$

Due to multiplicative noise, (2.50) is not in JKO canonical form (2.13). If the initial PDF ρ_0 is Dirac delta at x_0, then the FPK PDE for (2.50) admits closed form solution:

$$\rho(x,t) = \begin{cases} c\exp(-(u+v))\left(\dfrac{v}{u}\right)^{q/2} I_q\left(2\sqrt{uv}\right), & x > 0, \\ 0, & \text{otherwise}, \end{cases} \tag{2.51}$$

where $I_q(\cdot)$ is the modified Bessel function of order q, and

$$q := \frac{2a\theta}{b^2} - 1, \qquad c := \frac{2a}{b^2\left(1 - \exp(-at)\right)}, \tag{2.52a}$$

$$u := cx_0\exp(-at), \qquad v := cx. \tag{2.52b}$$

The transient solutions (2.51) are non-central chi-squared PDFs. The stationary solution is a Gamma PDF $\rho_\infty(x) \propto x^q \exp\left(-2ax/b^2\right)$, $x > 0$. We will benchmark our proximal algorithm against (2.51).

In order to transcribe (2.50) in the JKO canonical form, we employ the Lamperti transform [153, 162], where the idea is to find a change of variable $y = \varsigma(x)$ such that the SDE for new variable y has unity diffusion coefficient. From Itô's lemma, it follows that the requisite transformation for (2.50) is $y = \varsigma(x) := \frac{2}{b}\sqrt{x}$, and the

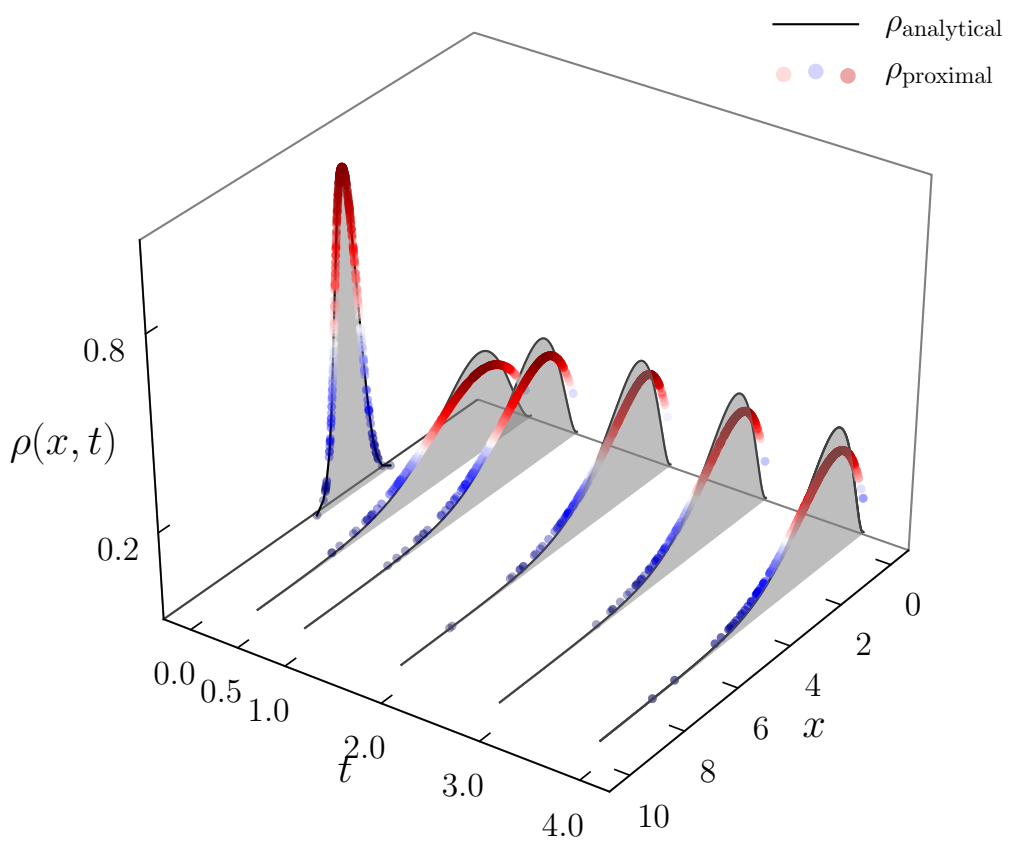

Figure 2.11: Comparison of the analytical and proximal transient PDFs of the FPK PDE for (2.50) with time step $h = 10^{-3}$, and with parameters $a = 3$, $b = 2$, $\theta = 2$, $x_0 = 5$, $\epsilon = 5 \times 10^{-2}$. To approximate the analytical PDFs resulting from $\rho_0(x) = \delta(x - 5)$, the proximal recursions were performed with the initial PDF $\mathcal{N}(5, 10^{-4})$.

resulting SDE in y becomes

$$\mathrm{d}y = \left\{ \left(\frac{2a\theta}{b^2} - \frac{1}{2} \right) \frac{1}{y} - \frac{a}{2} y \right\} \mathrm{d}t \, + \, \mathrm{d}w. \tag{2.53}$$

Clearly, (2.53) is in the JKO canonical form (2.13) with $\beta = 2$, and

$$\psi(y) = ay^2/4 - (q + 1/2) \log y. \tag{2.54}$$

So the proposed proximal algorithm (Algorithm 1) can be applied to (2.53), and at each time t, the resulting PDF $\rho_Y(y, t)$ can be transformed back to the PDF

$\rho_X(x,t)$ via the change-of-measure formula for the push-forward map $x = \varsigma^{-1}(y) = by^2/4$. Fig. 2.11 shows the comparison of the analytical and proximal transient PDFs for (2.50) resulting from $\rho_0 = \delta(x-5)$ with $a=3$, $b=2$, $\theta=2$, $h=10^{-3}$, $\epsilon = 5\times 10^{-2}$, $N=400$, $\delta = 10^{-3}$, and $L=100$.

2.5.2 Mixed Conservative-Dissipative Drift

In many engineering applications, one encounters Itô SDEs where the drift vector fields have both dissipative (gradient) and conservative (Hamiltonian) components. For example, stochastic systems arising from Newton's law in mechanics often have mixed conservative-dissipative structure. Our intent here is to illustrate that such SDEs are amenable to the proposed proximal recursion framework. As an example, we will work out the details for a perturbed two-body problem in celestial mechanics similar to the one treated in [205].

We consider the relative motion of a satellite in geocentric orbit, given by the second order Langevin equation

$$\ddot{\boldsymbol{q}} = -\frac{\mu \boldsymbol{q}}{\| \boldsymbol{q} \|_2^3} \;-\; \gamma \dot{\boldsymbol{q}} \;+\; \boldsymbol{f}_{\text{pert}}(\boldsymbol{q}) \\ +\; \sqrt{2\beta^{-1}\gamma} \;\times\; \text{stochastic forcing}, \tag{2.55}$$

where $\boldsymbol{q} := (x,y,z)^\top \in \mathbb{R}^3$ is the relative position vector for the satellite, μ is a constant (product of the Gavitational constant and the mass of Earth), $-\gamma\dot{\boldsymbol{q}}$ models linear drag[†], $\boldsymbol{f}_{\text{pert}}(\boldsymbol{q})$ models (deterministic) perturbative force due to the oblateness of Earth, and the stochastic forcing is due to solar radiation pressure,

[†]More generally, for nonlinear drag of the form $-\gamma\nabla_{\dot{\boldsymbol{q}}}\widetilde{V}(\dot{\boldsymbol{q}})$, the term $-\gamma\boldsymbol{p}$ in (2.59) will become $-\gamma\nabla\widetilde{V}(\boldsymbol{p})$. This will entail modifying the functional $\widehat{F}$ in (2.61) as $\mathbb{E}_\rho\left[\widetilde{V}(\boldsymbol{p}) \;+\; \beta^{-1}\log\rho\right]$. The linear drag illustrated here corresponds to the special case $\widetilde{V}(\boldsymbol{p}) \equiv \| \boldsymbol{p} \|_2^2/2$.

free-molecular aerodynamic forcing etc. Using the shorthands for sines and cosines as $c\tau := \cos\tau$, $s\tau := \sin\tau$, and recalling the relations among spherical coordinates (r, θ, ϕ) and cartesian coordinates (x, y, z), given by $r = \sqrt{x^2 + y^2 + z^2}$, $c\phi = x/r$, $s\phi = \sqrt{1 - (c\phi)^2}$, $c\theta = z/r$, $s\theta = \sqrt{1 - (c\theta)^2}$, we can write $\boldsymbol{f}_{\text{pert}}(\boldsymbol{q})$ in spherical coordinates as [205, eqn. (12)-(13)]

$$\begin{pmatrix} f_r \\ f_\theta \\ f_\phi \end{pmatrix}_{\text{pert}} = \begin{pmatrix} \frac{k}{2r^4}\left(3(s\theta)^2 - 1\right) \\ -\frac{k}{r^5} s\theta\, c\theta \\ 0 \end{pmatrix}, \; k := 3J_2 R_{\text{E}}^2 \mu = \text{ constant}, \tag{2.56}$$

and the same in cartesian coordinates as

$$\boldsymbol{f}_{\text{pert}}(\boldsymbol{q}) = \begin{pmatrix} f_x \\ f_y \\ f_z \end{pmatrix}_{\text{pert}} = \begin{pmatrix} s\theta\, c\phi & c\theta\, c\phi & -s\phi \\ s\theta\, s\phi & c\theta\, s\phi & c\phi \\ c\theta & -s\theta & 0 \end{pmatrix} \begin{pmatrix} f_r \\ f_\theta \\ f_\phi \end{pmatrix}_{\text{pert}} . \tag{2.57}$$

In (2.56), the Earth oblateness coefficient $J_2 = 1.082 \times 10^{-3}$, the radius of Earth $R_{\text{E}} = 6.3781 \times 10^6$ m, and the Earth standard Gravitational parameter $\mu = 3.9859 \times 10^{14}$ m^3/s^2. Modeling the stochastic forcing in (2.55) as standard Gaussian

white noise (as in [205]), (2.55) can then be expressed as an Itô SDE in $\mathbb{R}^6$:

$$\begin{pmatrix} \mathrm{d}x \\ \mathrm{d}y \\ \mathrm{d}z \\ \mathrm{d}v_x \\ \mathrm{d}v_y \\ \mathrm{d}v_z \end{pmatrix} = \begin{pmatrix} v_x \\ v_y \\ v_z \\ -\dfrac{\mu x}{r^3} + (f_x)_{\mathrm{pert}} - \gamma v_x \\ -\dfrac{\mu y}{r^3} + (f_y)_{\mathrm{pert}} - \gamma v_y \\ -\dfrac{\mu z}{r^3} + (f_z)_{\mathrm{pert}} - \gamma v_z \end{pmatrix} \mathrm{d}t + \sqrt{2\beta^{-1}\gamma} \begin{pmatrix} 0 \\ 0 \\ 0 \\ \mathrm{d}w_1 \\ \mathrm{d}w_2 \\ \mathrm{d}w_3 \end{pmatrix}, \tag{2.58}$$

or more succinctly,

$$\begin{pmatrix} \mathrm{d}\boldsymbol{q} \\ \mathrm{d}\boldsymbol{p} \end{pmatrix} = \begin{pmatrix} \boldsymbol{p} \\ -\nabla V(\boldsymbol{q}) - \gamma \boldsymbol{p} \end{pmatrix} \mathrm{d}t \; + \; \sqrt{2\beta^{-1}\gamma} \begin{pmatrix} \mathbf{0}_{3\times 1} \\ \mathrm{d}\boldsymbol{w}_{3\times 1} \end{pmatrix}, \tag{2.59}$$

where $\boldsymbol{p} := (v_x, v_y, v_z)^\top \in \mathbb{R}^3$ is the velocity vector, and

$$V(\boldsymbol{q}) := V_{\mathrm{gravitational}}(\boldsymbol{q}) + V_{\mathrm{pert}}(\boldsymbol{q}), \tag{2.60a}$$

$$\nabla_{\boldsymbol{q}} V_{\mathrm{gravitational}}(\boldsymbol{q}) = \frac{\mu \boldsymbol{q}}{\| \boldsymbol{q} \|_2^3}, \; -\nabla_{\boldsymbol{q}} V_{\mathrm{pert}}(\boldsymbol{q}) = \begin{pmatrix} f_x \\ f_y \\ f_z \end{pmatrix}_{\mathrm{pert}}. \tag{2.60b}$$

Introducing a "Hamiltonian-like function" $H(\boldsymbol{q}, \boldsymbol{p}) := \| \boldsymbol{p} \|_2^2 / 2 + V(\boldsymbol{q})$, one can verify that the stationary PDF for (2.59) is $\rho_\infty \propto \exp(-\beta H(\boldsymbol{q}, \boldsymbol{p}))$, and that the "total free energy" $F(\rho) := \mathbb{E}_\rho [H + \beta^{-1} \log \rho]$ serves as Lyapunov functional

for the associated FPK PDE, i.e., $\frac{\mathrm{d}F}{\mathrm{d}t} < 0$. However, the proximal recursion (2.22b) does not apply as is, instead needs to be modified to account the joint conservative-dissipative effect as

$$\varrho_k = \underset{\varrho \in D_2}{arg\,inf}\ \frac{1}{2}\widehat{W}_h^2(\varrho_{k-1}, \varrho) + h\gamma\, \widehat{F}(\varrho), \quad k \in \mathbb{N}. \tag{2.61}$$

We implement a recursion from [81, Scheme 2b] where

$$\widehat{F}(\varrho) := \mathbb{E}_\rho \left[\frac{1}{2} \parallel \boldsymbol{p} \parallel_2^2 + \beta^{-1} \log \rho \right], \tag{2.62}$$

and $\widehat{W}_h^2(\varrho_{k-1}, \varrho)$ is the optimal mass transport cost (as in squared 2-Wasserstein metric (2.10)) with modified cost function (modified integrand in (2.10)), i.e.,

$$\widehat{W}_h^2(\rho_1, \rho_2) := \underset{d\pi \in \Pi(\pi_1, \pi_2)}{inf} \int_{\mathcal{X}\times\mathcal{Y}} \widehat{s}_h\left(\boldsymbol{q}, \boldsymbol{p}; \widetilde{\boldsymbol{q}}, \widetilde{\boldsymbol{p}}\right)\, \mathrm{d}\pi\left(\boldsymbol{q}, \boldsymbol{p}, \widetilde{\boldsymbol{q}}, \widetilde{\boldsymbol{p}}\right), \tag{2.63}$$

where $(\boldsymbol{q}, \boldsymbol{p})^\top$ and $(\widetilde{\boldsymbol{q}}, \widetilde{\boldsymbol{p}})^\top$ are two realizations of the state vector governed by (2.59)-(2.60); the integrand in (2.63) is

$$\begin{aligned} \widehat{s}_h\left(\boldsymbol{q}, \boldsymbol{p}; \widetilde{\boldsymbol{q}}, \widetilde{\boldsymbol{p}}\right) \ :=\ & \parallel \widetilde{\boldsymbol{p}} - \boldsymbol{p} + h\nabla V(\boldsymbol{q}) \parallel_2^2 \\ & + 12 \left\| \frac{\widetilde{\boldsymbol{q}} - \boldsymbol{q}}{h} - \frac{\widetilde{\boldsymbol{p}} + \boldsymbol{p}}{2} \right\|_2^2. \end{aligned} \tag{2.64}$$

That the proximal recursion (2.61) with the above choices of functionals $\widehat{F}$ and $\widehat{W}_h$ guarantees $\varrho_k(\boldsymbol{q}, \boldsymbol{p}, h) \to \rho(\boldsymbol{q}, \boldsymbol{p}, t = kh)$ for $k \in \mathbb{N}$ as $h \downarrow 0$, where $\rho(\boldsymbol{q}, \boldsymbol{p}, t)$ is the joint PDF generated by the FPK flow for (2.59) at time t, was proved in [81]. Our proximal algorithm in Section 2.3.2 applies by simply modifying (2.24)-(2.25)

as

$$\boldsymbol{\psi}_{k-1}(i) := \frac{1}{2}(\boldsymbol{p}^i_{k-1})^\top \boldsymbol{p}^i_{k-1}, \quad i = 1, \ldots, N, \tag{2.65}$$

$$\boldsymbol{C}_k(i,j) := \widehat{s}_h(\boldsymbol{q}^i_k, \boldsymbol{p}^i_k; \boldsymbol{q}^j_{k-1}, \boldsymbol{p}^j_{k-1}), \quad i,j = 1, \ldots, N, \tag{2.66}$$

respectively. Notice that $\widehat{s}_h$ in (2.64) (and consequently, $\widehat{W}_h$ in (2.63)) is not a metric (in particular, non-symmetric). Thus, the matrices $\boldsymbol{C}_k$ in (2.66) for $k \in \mathbb{N}$ are not symmetric.

To apply the proposed framework, we first non-dimensionalize the variables $\boldsymbol{q}$ (m), $\boldsymbol{p}$ (m/s), t (s), $\boldsymbol{w}$ ($\sqrt{\text{s}}$) in (2.59) as

$$\boldsymbol{q}' = \boldsymbol{q}/R, \quad \boldsymbol{p}' = \boldsymbol{p}/(R/T), \quad t' = t/T, \quad \boldsymbol{w}' = \boldsymbol{w}/\sqrt{T}, \tag{2.67}$$

where $R := 4.2164 \times 10^7$ m is the radius of the nominal geostationary orbit, and $T = 86164$ s is its period. In (2.67), the primed variables are non-dimensionalized. Using (2.67) and Itô's lemma on (2.59), the non-dimensional SDE in $(\boldsymbol{q}', \boldsymbol{p}')^\top \in \mathbb{R}^6$ becomes

$$\mathrm{d}\boldsymbol{q}' = \boldsymbol{p}' \, \mathrm{d}t', \tag{2.68a}$$

$$\mathrm{d}\boldsymbol{p}' = \Big\{ -\frac{T\mu}{R^3} \frac{\boldsymbol{q}'}{\| \boldsymbol{q}' \|_2^3} + \frac{T^2}{R} \boldsymbol{f}_{\text{pert}}(R\boldsymbol{q}') - \gamma T \boldsymbol{p}' \Big\} \mathrm{d}t' + \frac{T^{3/2}}{R} \sqrt{2\beta^{-1}\gamma} \, \mathrm{d}\boldsymbol{w}'. \tag{2.68b}$$

To avoid numerical conditioning issues, we will apply our algorithm for recursion (2.61) associated with the non-dimensional SDE (2.68), and then transform the PDF in $(\boldsymbol{q}', \boldsymbol{p}')^\top$ to the same in original variables $(\boldsymbol{q}, \boldsymbol{p})^\top$ via change-of-measure formula. Such a computational pipeline, i.e., density propagation in non-dimensional

SDE and then transforming the density back in dimensional variables, is standard in celestial mechanics due to different orders of magnitude in different state variables (see e.g., [205, Sec. III.B]).

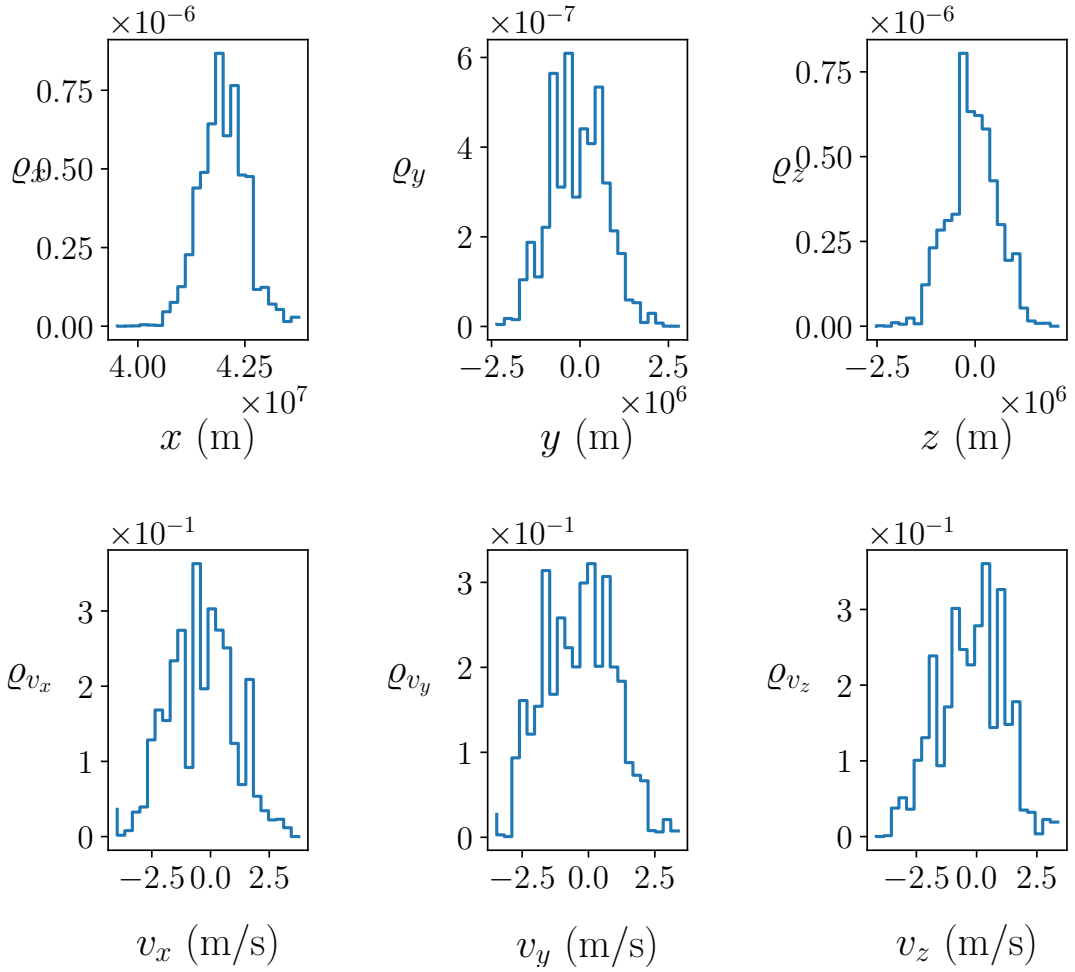

Figure 2.12: Univariate marginal PDFs at $t = 0.005$ s for (2.59)-(2.60) computed from the joint PDF at that time obtained via the proposed proximal algorithm for recursion (2.61) with time step $h = 10^{-5}$, and with parameters $\beta = 1$ m^2/s^2, $\gamma = 1$ s^{-1}, $\epsilon = 5 \times 10^{-2}$, $\delta = 10^{-3}$, $L = 100$, and $N = 400$.

Figs. 2.12 and 2.13 show the univariate marginal PDFs at $t = 0.005$s and $t = 0.010$s respectively, associated with the joint PDFs supported on $\mathbb{R}^6$ at those times, computed through the proposed proximal algorithm. For this simulation, the initial joint PDF $\rho_0(\boldsymbol{q}', \boldsymbol{p}') = \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$, where $\boldsymbol{\mu}_0 = (1, 0, 0, 0, 0, 0)^\top$, and

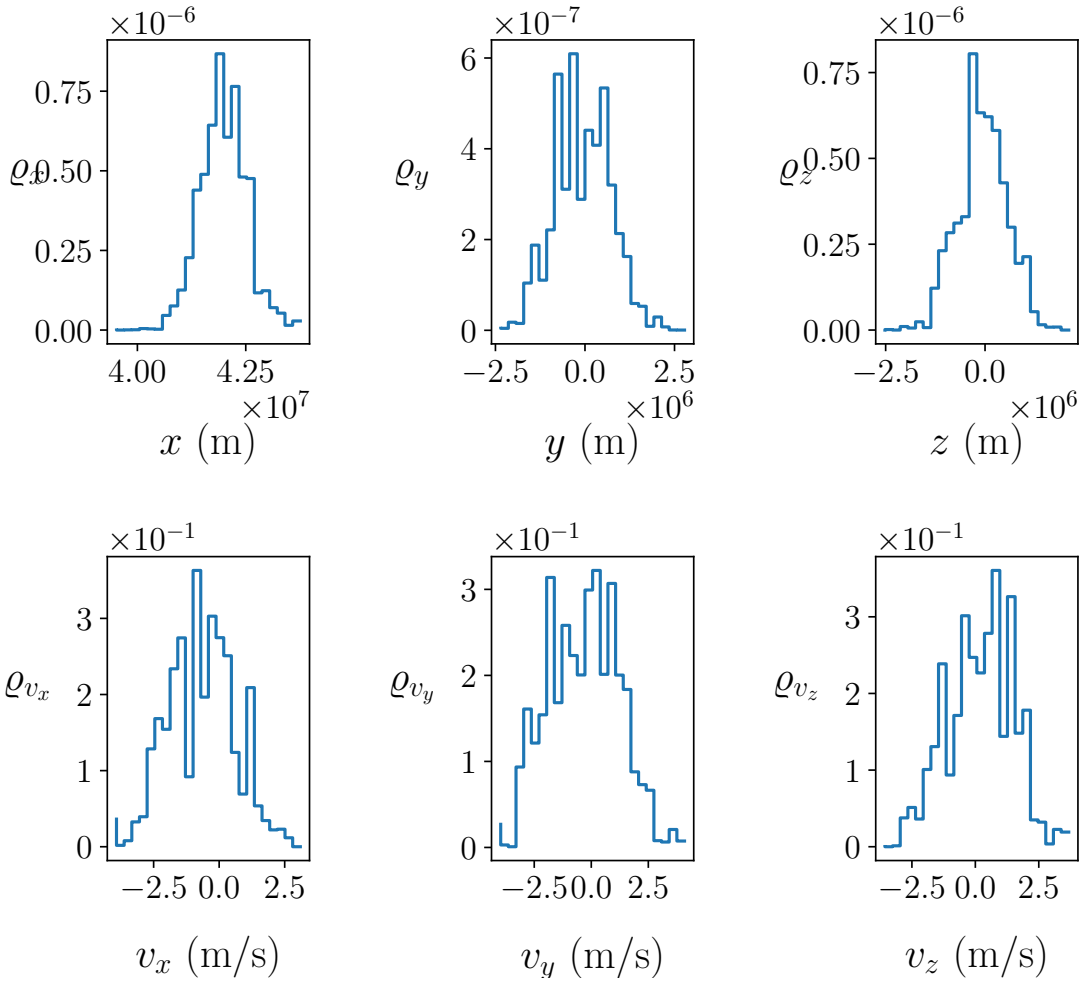

Figure 2.13: Univariate marginal PDFs at $t = 0.01$ s for (2.59)-(2.60) computed from the joint PDF at that time obtained via the proposed proximal algorithm for recursion (2.61) with time step $h = 10^{-5}$, and with parameters $\beta = 1$ m^2/s^2, $\gamma = 1$ s^{-1}, $\epsilon = 5 \times 10^{-2}$, $\delta = 10^{-3}$, $L = 100$, and $N = 400$.

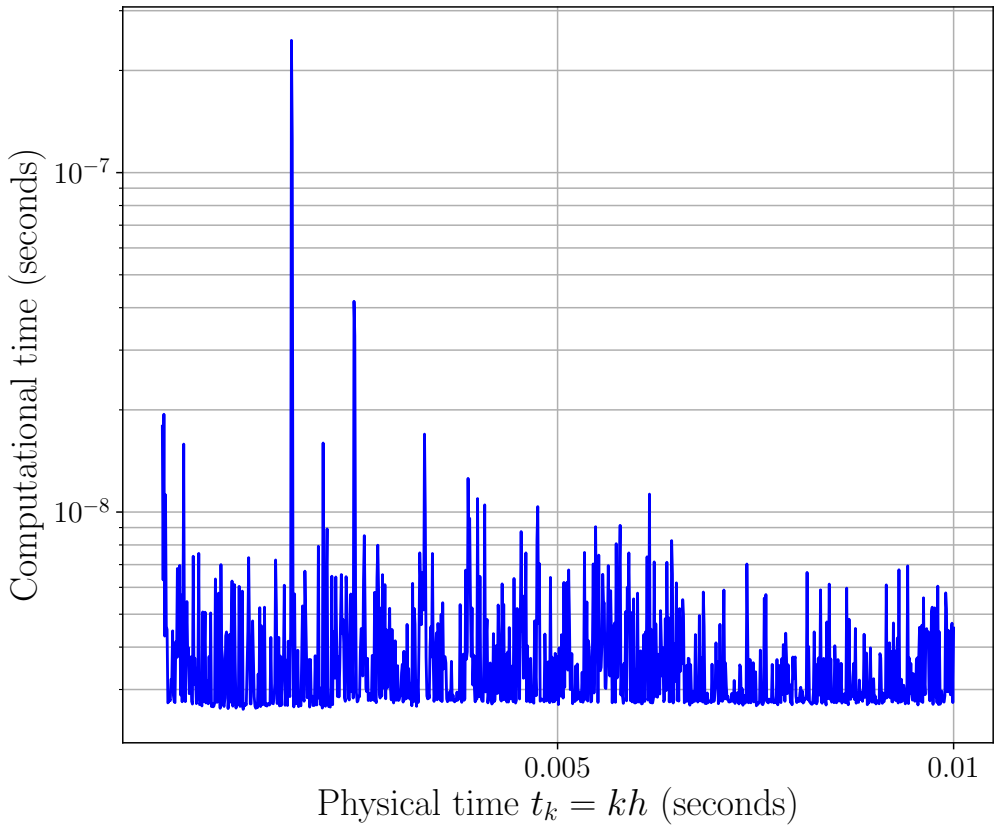

Figure 2.14: The computational times needed for proximal updates in the 6-state numerical example reported in Section 2.5.2. Here, the physical time-step $h = 10^{-5}$ s, and $k \in \mathbb{N}$.

$\mathbf{\Sigma}_0 = 10^{-4} \times \mathrm{diag}\,(3.335, 6.133, 3.933, 6.562, 9.246, 5.761)$. The parameter values used in the simulation are: $h = 10^{-5}$, $\beta = 1$ m^2/s^2, $\gamma = 1$ s^{-1}, $\epsilon = 5 \times 10^{-2}$, $\delta = 10^{-3}$, $L = 100$, and $N = 400$. The computational times for this 6-state example reported in Fig. 2.14 reveal that the proximal recursions run remarkably fast (indeed, faster than the runtime for the 2-state example in Section 2.4.2, cf. Fig. 2.9).

2.6 Conclusions

In this paper, novel uncertainty propagation algorithms are presented for computing the flow of the joint PDFs associated with continuous-time stochastic nonlinear systems. By interpreting the PDF flow as gradient descent on the manifold of joint PDFs w.r.t. a suitable metric, the proposed computational framework implements proximal algorithms which are proved to be convergent due to certain contraction properties established herein. Numerical examples are provided to demonstrate the practical use of the proposed algorithm and its efficiency in terms of computational runtime. In contrast to the conventional function approximation algorithms for this problem, the proposed non-parametric framework does not make any spatial discretization, instead performs finite sample probability-weighted scattered data computation in the form of temporal recursion. The location of the atomic measures is delegated to a Euler-Maruyama scheme, thus avoiding spatial discretization. The results of this paper provide computational teeth to the emerging systems-theoretic viewpoint [100, 101] that the PDF flows in uncertainty propagation can be seen as gradient flow.

Chapter 3

Wasserstein Proximal Algorithms for the Schrödinger Bridge Problem: Density Control with Nonlinear Drift

We study the Schrödinger bridge problem (SBP) with nonlinear prior dynamics. In control-theoretic language, this is a problem of minimum effort steering of a given joint state probability density function (PDF) to another over a finite time horizon, subject to a controlled stochastic differential evolution of the state vector. As such, it can be seen as a stochastic optimal control problem in continuous time with endpoint density constraints – a topic that originated in the physics literature in 1930s, and in the recent years, has garnered burgeoning interest in the systems-control community.

For generic nonlinear drift, we reduce the SBP to solving a system of forward

and backward Kolmogorov partial differential equations (PDEs) that are coupled through the boundary conditions, with unknowns being the “Schrödinger factors" – so named since their product at any time yields the optimal controlled joint state PDF at that time. We show that if the drift is a gradient vector field, or is of mixed conservative-dissipative nature, then it is possible to transform these PDEs into a pair of initial value problems (IVPs) involving the same forward Kolmogorov operator. Combined with a recently proposed fixed point recursion that is contractive in the Hilbert metric, this opens up the possibility to numerically solve the SBPs in these cases by computing the Schrödinger factors via a single IVP solver for the corresponding (uncontrolled) forward Kolmogorov PDE. The flows generated by such forward Kolmogorov PDEs, for the two aforementioned types of drift, in turn, enjoy gradient descent structures on the manifold of joint PDFs with respect to suitable distance functionals. We employ a proximal algorithm developed in our prior work, that exploits this geometric viewpoint, to solve these IVPs and compute the Schrödinger factors via weighted scattered point cloud evolution in the state space. We provide the algorithmic details and illustrate the proposed framework of solving the SBPs with nonlinear prior dynamics by numerical examples.

3.1 Introduction

The Schrödinger bridge problem (SBP) is a non-standard *finite horizon* stochastic optimal control problem in continuous time. The “non-standard" aspect stems from the fact that the SBP concerns with steering the flow of the joint state probability density function (PDF), and not the state trajectory *per se*, in a controlled stochastic dynamical system while minimizing the expected control

effort to do so. In other words, the SBP is a two-point density control problem subject to controlled trajectory-level dynamics. Fig. 3.1 shows a schematic of the SBP.

Solving the SBP amounts to *feedback synthesis for ensemble shaping*. The problem is of broad contemporary interest due to its potential applications in controlling a *physical population* such as robotic swarm [15], ensemble of neurons [163], and density of highway traffic [63]. From a probabilistic perspective, one can alternatively interpret controlling the joint state PDF as that of dynamically reshaping uncertainties in a feedback loop – a viewpoint that promotes "control *of* uncertainties" in lieu of the usual "control *with* uncertainties". That the role of feedback could be strategically synthesizing, instead of simply mitigating the uncertainties, is a recent line of thought [33, 34, 157].

During 1931-32, the SBP was introduced by Erwin Schrödinger in the two articles [191, 192]. Schrödinger's original motivation was to come up with a classical reformulation of quantum mechanics via diffusion processes; see e.g., [221, 222]. Early mathematical treatment [25,86,124] of the subject considered the case of no prior dynamics, which is what we refer to as the "classical SBP". The stochastic optimal control interpretation for the same emerged in [75,159]. We point the readers to [145,216] for survey of the classical SBP. More recent works [38,39,54,55,59] have considered SBP with prior dynamics. Except the case of linear prior dynamics with Gaussian endpoint PDFs, the SBP with prior dynamics–while most relevant in control applications–remains computationally challenging in general. The computational difficulty is particularly evident in the case of *nonlinear* prior dynamics, since then, solving the SBP reduces to solving a system of coupled nonlinear partial differential equations (PDEs). Numerical algorithms for solving the SBP are only beginning to appear now [147, 182]. The purpose of this paper

is to show that if the nonlinear prior dynamics is either a gradient vector field, or has both dissipative (gradient) and conservative (Hamiltonian) components, then we can design scalable point cloud algorithms based on the recursive evaluation of certain proximal operators with respect to (w.r.t.) the Wasserstein metric, that solves the associated SBPs. The types of nonlinearities considered in this paper subsume well-known stochastic models in practical applications such as the Nyquist-Johnson resistors [35,148], rotational diffusions for describing the orientation in liquid crystals [64, Ch. 16.5], and statistical mechanics of polymers [64, Ch. 17].

The Wasserstein proximal recursions we employ, solve the forward-backward infinite-dimensional flows in the space of the *Schrödinger factors* (see Section 3.2.1), defined as functions of the state vector and time, whose product at any given time equals the optimal controlled joint state PDF at that time. These factors satisfy a system of boundary-coupled PDEs, which we solve via the proximal recursions. An intriguing aspect of our proposed framework is that the Wasserstein proximal recursions originate in the theory of optimal mass transport (OMT) [210], and two connections between the OMT and the SBP are well-known: *first*, the dynamic version [21] of the OMT can be recovered as the zero-noise limit [143, 160] of the classical SBP; *second*, the classical SBP can be recovered as the solution of the dynamic OMT problem with Fisher information regularization [59, Sec. 5], [70, Sec. 4]. The use of Wasserstein gradient flows to solve the SBP with prior nonlinear drift, as proposed here, offers a *third* connection between the OMT and the SBP. This can be of independent interest.

The remaining of this paper is structured as follows. In Section 3.2, we recap the basics of the SBP. Section 3.3 considers the SBP with generic nonlinear prior dynamics, deduces the necessary conditions of optimality for the same, and derives

an associated Schrödinger system via the Hopf-Cole transform. In Section 3.4, we specialize the development in Section 3.3 for two important types of prior nonlinear dynamics: gradient drift, and mixed conservative-dissipative drift. We show how to reformulate the corresponding Schrödinger systems to make the same algorithmically amenable. In Section 3.5, we show that the reformulations derived in Section 3.4 can be solved via Wasserstein proximal recursions. Section 3.6 provides numerical examples to illustrate the proposed algorithmic framework. Section 3.7 concludes the paper.

3.2 Background

Notation

We use bold-faced capital letters for matrices, and bold-faced lower-case letters for column vectors. The Euclidean gradient, divergence, Laplacian and Hessian operators are denoted by ∇, $\nabla\cdot$, Δ, and $\mathbf{Hess}(\cdot)$, respectively. We will sometimes put subscript to these differential operators to clarify that the operator is w.r.t. the subscripted variable (e.g, $\nabla_{\boldsymbol{x}}$ to mean that the gradient is w.r.t. vector $\boldsymbol{x}$); we will drop the subscript when there is no scope for confusion. We use $\langle\cdot,\cdot\rangle$ to denote the standard Euclidean inner product. For vectors $\boldsymbol{x},\boldsymbol{y}\in\mathbb{R}^n$, we have $\langle\boldsymbol{x},\boldsymbol{y}\rangle := \boldsymbol{x}^\top\boldsymbol{y}$, and that the squared Euclidean 2-norm $\|\ \boldsymbol{x}\ \|_2^2 := \langle\boldsymbol{x},\boldsymbol{x}\rangle$. For matrices $\boldsymbol{A},\boldsymbol{B}$ of appropriate dimensions, $\langle\boldsymbol{A},\boldsymbol{B}\rangle := \mathrm{trace}(\boldsymbol{A}^\top\boldsymbol{B})$ denotes the Frobenius inner product. The symbol $\mathbf{0}$ denotes either a column vector or a matrix (depending on context), with all entries equal to zeros; the symbol $\mathbf{1}$ denotes the column vector

containing all ones; $\delta(\boldsymbol{x}-\boldsymbol{y})$ denotes the Dirac delta located at $\boldsymbol{y}$.

$$\mathcal{P}_2\left(\mathbb{R}^n\right) := \left\{\rho : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0} \middle| \int_{\mathbb{R}^n} \rho \,\mathrm{d}\boldsymbol{x} = 1, \int_{\mathbb{R}^n} \| \boldsymbol{x} \|_2^2 \, \rho \,\mathrm{d}\boldsymbol{x} < \infty \right\}.$$

We use the shorthand $\boldsymbol{x} \sim \rho$ to mean that the random vector $\boldsymbol{x}$ has the joint PDF ρ. The notation $\mathcal{N}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}\right)$ stands for a joint Gaussian PDF with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. The symbol $\boldsymbol{I}$ denotes an identity matrix of appropriate dimension. We use the the notations $\odot$ and $\oslash$ for element-wise (i.e., Hadamard) multiplication and division, respectively. The operands $\log(\cdot)$ and $\geq$ are to be understood element-wise. The set of natural numbers is denoted by $\mathbb{N}$.

In the following, we briefly review the classical SBP, i.e., the case of no prior dynamics. We then point out a recent extension: the SBP with linear prior dynamics. Along the way, we introduce the Schrödinger system, and the Schrödinger factors, which will be important in our development.

3.2.1 Classical SBP and the Schrödinger System

Let the state space be $\mathbb{R}^n$. The *classical SBP* is a stochastic optimal control problem of the form

$$\inf_{\boldsymbol{u} \in U} \quad \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2 \,\mathrm{d}t\right\}, \tag{3.1a}$$

$$\text{subject to} \quad \mathrm{d}\boldsymbol{x}(t) = \boldsymbol{u}(\boldsymbol{x},t)\,\mathrm{d}t + \sqrt{2\epsilon}\,\mathrm{d}\boldsymbol{w}(t), \tag{3.1b}$$

$$\boldsymbol{x}(t=0) \sim \rho_0(\boldsymbol{x}), \quad \boldsymbol{x}(t=1) \sim \rho_1(\boldsymbol{x}), \tag{3.1c}$$

where the set of feasible controls $\mathcal{U}$ comprises of finite energy inputs, i.e., $\mathcal{U} := \{\boldsymbol{u} : \mathbb{R}^n \times [0,1] \mapsto \mathbb{R}^n \mid \langle \boldsymbol{u}, \boldsymbol{u} \rangle < \infty\}$. Here, $\boldsymbol{w}(t)$ denotes the standard Wiener

process in $\mathbb{R}^n$, the diffusion coefficient $\epsilon > 0$ (not necessarily small), the prescribed initial and terminal joint state PDFs are ρ_0 and ρ_1, respectively. The expectation operator in (3.1a) is taken w.r.t. the controlled joint state PDF $\rho(\boldsymbol{x}, t)$.

The SBP (3.1) can be transcribed into the following equivalent variational problem:

$$\begin{aligned} \underset{\rho, \boldsymbol{u}}{inf} \quad & \int_0^1 \int_{\mathbb{R}^n} \frac{1}{2} \|\boldsymbol{u}(\boldsymbol{x}, t)\|_2^2 \, \rho(\boldsymbol{x}, t) \, \mathrm{d}\boldsymbol{x} \, \mathrm{d}t, && \text{(3.2a)} \\ \text{subject to} \quad & \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{u}) = \epsilon \Delta \rho, && \text{(3.2b)} \\ & \rho(\boldsymbol{x}, 0) = \rho_0(\boldsymbol{x}), \quad \rho(\boldsymbol{x}, 1) = \rho_1(\boldsymbol{x}), && \text{(3.2c)} \end{aligned}$$

where (3.2b) is the Fokker-Planck or Kolmogorov's forward PDE, hereafter abbreviated as the FPK PDE, associated with the controlled stochastic differential equation (SDE) (3.1b). Let $\mathcal{P}(\mathbb{R}^n)$ denote the collection of all joint PDFs supported on $\mathbb{R}^n$. Notice that the decision variable in problem (4.3) is the pair $(\rho, \boldsymbol{u}) \in \mathcal{P}(\mathbb{R}^n) \times \mathcal{U}$.

From the first order conditions of optimality for (4.3), it is easy to verify that the optimal pair $(\rho^{\rho^{\text{opt}}}, \boldsymbol{u}^{\rho^{\text{opt}}})$ satisfies the following system of PDEs:

$$\begin{aligned} &\frac{\partial \psi}{\partial t} + \frac{1}{2} \|\nabla \psi\|_2^2 = -\epsilon \Delta \psi, \\ &\frac{\partial}{\partial t} \rho^{\rho^{\text{opt}}} + \nabla \cdot (\rho^{\rho^{\text{opt}}} \nabla \psi) = \epsilon \Delta \rho^{\rho^{\text{opt}}}, \end{aligned} \tag{3.3}$$

and the optimal control $\boldsymbol{u}^{\rho^{\text{opt}}}(\boldsymbol{x}, t) = \nabla \psi(\boldsymbol{x}, t)$.

Furthermore, (3.3) can be transformed into a system of *linear* PDEs via the

mapping $\left(\rho^{\rho^{\text{opt}}}, \psi\right) \mapsto (\varphi, \hat{\varphi})$ given by

$$\varphi(\boldsymbol{x}, t) = \exp\left(\frac{\psi(\boldsymbol{x}, t)}{2\epsilon}\right), \tag{3.4a}$$

$$\hat{\varphi}(\boldsymbol{x}, t) = \rho^{\rho^{\text{opt}}}(\boldsymbol{x}, t) \exp\left(-\frac{\psi(\boldsymbol{x}, t)}{2\epsilon}\right). \tag{3.4b}$$

Specifically, the transformed variables $(\varphi, \hat{\varphi})$ satisfy the following pair of forward-backward heat equations

$$\frac{\partial \varphi}{\partial t} = -\epsilon \Delta \varphi, \tag{3.5a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = \epsilon \Delta \hat{\varphi}, \tag{3.5b}$$

with coupled boundary conditions

$$\varphi(\boldsymbol{x}, t = 0)\hat{\varphi}(\boldsymbol{x}, t = 0) = \rho_0(\boldsymbol{x}), \tag{3.6a}$$

$$\varphi(\boldsymbol{x}, t = 1)\hat{\varphi}(\boldsymbol{x}, t = 1) = \rho_1(\boldsymbol{x}). \tag{3.6b}$$

To ease notation, let

$$\varphi_1(\boldsymbol{x}) := \varphi(\boldsymbol{x}, t = 1), \quad \hat{\varphi}_0(\boldsymbol{x}) := \hat{\varphi}(\boldsymbol{x}, t = 0). \tag{3.7}$$

Then, the solution of (3.5) can be formally written as

$$\varphi(\boldsymbol{x}, t) = \int_{\mathbb{R}^n} K(t, \boldsymbol{x}, 1, \boldsymbol{y})\varphi_1(\boldsymbol{y}) \, \mathrm{d}\boldsymbol{y}, \quad t \leq 1, \tag{3.8a}$$

$$\hat{\varphi}(\boldsymbol{x}, t) = \int_{\mathbb{R}^n} K(0, \boldsymbol{y}, t, \boldsymbol{x})\hat{\varphi}_0(\boldsymbol{y}) \, \mathrm{d}\boldsymbol{y}, \quad t \geq 0, \tag{3.8b}$$

where

$$K(t, \boldsymbol{x}, s, \boldsymbol{y}) := (4\pi\epsilon(t - s))^{-n/2} \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{y}\|_2^2}{4\epsilon(t - s)}\right) \tag{3.9}$$

is the heat kernel or the Markov kernel associated with the pure diffusion SDE $\mathrm{d}\boldsymbol{x}(t) = \sqrt{2\epsilon}\,\mathrm{d}\boldsymbol{w}(t)$.

Combining (3.6) and (3.8), it becomes clear that finding the minimizer for the classical SBP amounts to solving for the pair $(\varphi_1, \hat{\varphi}_0)$ which satisfies the following system of nonlinear integral equations, referred to as the *Schrödinger system*, given by

$$\rho_0(\boldsymbol{x}) = \hat{\varphi}_0(\boldsymbol{x}) \int_{\mathbb{R}^n} K(0, \boldsymbol{x}, 1, \boldsymbol{y}) \varphi_1(\boldsymbol{y})\, \mathrm{d}\boldsymbol{y}, \tag{3.10a}$$

$$\rho_1(\boldsymbol{x}) = \varphi_1(\boldsymbol{x}) \int_{\mathbb{R}^n} K(0, \boldsymbol{y}, 1, \boldsymbol{x}) \hat{\varphi}_0(\boldsymbol{y})\, \mathrm{d}\boldsymbol{y}. \tag{3.10b}$$

The existence and uniqueness of solutions to the Schrödinger system (3.10) were established in [25,86,124]. To compute the pair $(\varphi_1, \hat{\varphi}_0)$ from (3.10), a fixed point recursion was proposed in [53]. Such a recursion was also proved [53, Sec. III] to be contractive in Hilbert's projective metric [36, 113]. Once the pair $(\varphi_1, \hat{\varphi}_0)$ is obtained from (3.10), then using (3.8) one can compute the pair $(\varphi, \hat{\varphi})$. Finally, from (5.50), the original decision variables $(\rho^{\rho^{\mathrm{opt}}}, \boldsymbol{u}^{\rho^{\mathrm{opt}}})$ can be recovered via the mapping $(\varphi, \hat{\varphi}) \mapsto \left(\rho^{\rho^{\mathrm{opt}}}, \boldsymbol{u}^{\rho^{\mathrm{opt}}}\right)$ given by

$$\rho^{\mathrm{opt}}(\boldsymbol{x}, t) = \varphi(\boldsymbol{x}, t)\, \hat{\varphi}(\boldsymbol{x}, t), \tag{3.11a}$$

$$\boldsymbol{u}^{\mathrm{opt}}(\boldsymbol{x}, t) = 2\epsilon \nabla \log \varphi(\boldsymbol{x}, t). \tag{3.11b}$$

From (3.11a), the optimal controlled joint state PDF at any time is a product of the factors φ and $\hat{\varphi}$ at that time, and hence we refer to $(\varphi, \hat{\varphi})$ as the *Schrödinger factors*. Notice that the factors solve the boundary-coupled PDE system (3.5)-(3.6).

The computational steps outlined in the preceding paragraph were used to solve

the classical SBP in Fig. 3.1 with $\epsilon = 0.5$. Using the optimal control computed from (3.11b), 100 sample paths (shown in the (x, t) plane in Fig. 3.1) of the optimal closed-loop SDE were simulated via the Euler-Maruyama scheme with time step-size 10^{-3}.

3.2.2 SBP with Linear Prior Dynamics

Recently, the classical SBP has been extended [60] to the case when the prior dynamics is a linear time-varying (LTV) system, i.e., (3.1b) is replaced with the more general controlled SDE

$$\mathrm{d}\boldsymbol{x}(t) = \boldsymbol{A}(t)\boldsymbol{x}(t)\,\mathrm{d}t + \boldsymbol{B}(t)\boldsymbol{u}(\boldsymbol{x}, t)\,\mathrm{d}t + \sqrt{2\epsilon}\boldsymbol{B}(t)\,\mathrm{d}\boldsymbol{w}(t), \tag{3.12}$$

where the system matrices $\boldsymbol{A}(t) \in \mathbb{R}^{n\times n}, \boldsymbol{B}(t) \in \mathbb{R}^{n\times m}$, $m \leq n$, and the pair $(\boldsymbol{A}(t), \boldsymbol{B}(t))$ is assumed to be controllable for all t. When $\boldsymbol{A}(t)$ is identically zero, and $\boldsymbol{B}(t)$ is identity, then the setup reduces to the classical SBP. A *Schrödinger system* for the LTV case can be established wherein the heat kernel (3.9) is to be replaced by the Markov kernel associated with the uncontrolled SDE

$$\mathrm{d}\boldsymbol{x}(t) = \boldsymbol{A}(t)\boldsymbol{x}(t)\,\mathrm{d}t + \sqrt{2\epsilon}\boldsymbol{B}(t)\,\mathrm{d}\boldsymbol{w}(t). \tag{3.13}$$

We refer the readers to [60, Sec. 4] for the details.

Thanks to the availability of the Markov kernel associated with (3.13), the fixed point recursion idea and the contraction results mentioned in Section 3.2.1 carry through in this case, making the computation tractable as in the classical SBP. We note that in (3.12), the noise and the control act through the same channels, i.e., the same matrix $\boldsymbol{B}(t)$ appears as the coefficient of both. This is the case we

will focus in this paper. The more general case of different coefficient matrices was treated in [57].

3.3 Problem Formulation

In this Section, we consider a much less investigated version of the SBP: a generalization of problem (3.1) with nonlinear prior dynamics given by a (deterministic) vector field $\boldsymbol{f} : \mathbb{R}^n \times \mathbb{R}_{>0} \mapsto \mathbb{R}^n$. In particular, we address the following minimum energy stochastic optimal control problem:

$$\inf_{\boldsymbol{u}\in U} \quad \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2 \,\mathrm{d}t\right\}, \tag{3.14a}$$

$$\text{subject to} \quad \mathrm{d}\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x},t)\,\mathrm{d}t + \boldsymbol{B}(t)\boldsymbol{u}(\boldsymbol{x},t)\,\mathrm{d}t + \sqrt{2\epsilon}\boldsymbol{B}(t)\,\mathrm{d}\boldsymbol{w}(t), \tag{3.14b}$$

$$\boldsymbol{x}(t=0) \sim \rho_0(\boldsymbol{x}), \quad \boldsymbol{x}(t=1) \sim \rho_1(\boldsymbol{x}), \tag{3.14c}$$

where $\boldsymbol{x} \in \mathbb{R}^n$, and the set $\mathcal{U}$ comprises of all finite energy inputs, as before. Clearly, (3.12) corresponds to the special case $\boldsymbol{f}(\boldsymbol{x},t) \equiv \boldsymbol{A}(t)\boldsymbol{x}(t)$ in (3.14b). Given $\boldsymbol{f}(\boldsymbol{x},t)$, $\boldsymbol{B}(t)$ and ϵ, our objective is to steer the joint state PDF $\rho(\boldsymbol{x},t)$ from a prescribed initial PDF ρ_0 at $t = 0$ to another prescribed terminal PDF ρ_1 at $t = 1$ while minimizing the expected control effort. A schematic for the problem is shown in Fig. 3.2.

Define the diffusion tensor $\boldsymbol{D}(t) := \boldsymbol{B}(t)\boldsymbol{B}(t)^\top$. Problem (3.14) can be formally

recast into a "fluid dynamics" version [21] given by:

$$\inf_{(\rho,\boldsymbol{u})} \quad \frac{1}{2}\int_0^1 \int_{\mathbb{R}^n} \|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2\, \rho(\boldsymbol{x},t)\, \mathrm{d}\boldsymbol{x}\, \mathrm{d}t \tag{3.15a}$$

$$\text{subject to} \quad \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho\,(\boldsymbol{f} + \boldsymbol{B}(t)\boldsymbol{u})) = \epsilon \mathbf{1}^\top (\boldsymbol{D}(t) \odot \mathbf{Hess}\,(\rho))\,\mathbf{1}, \tag{3.15b}$$

$$\rho(\boldsymbol{x},0) = \rho_0(\boldsymbol{x}), \quad \rho(\boldsymbol{x},1) = \rho_1(\boldsymbol{x}), \tag{3.15c}$$

where the infimum is taken over all pairs $(\rho, \boldsymbol{u}) \in \mathcal{P}(\mathbb{R}^n) \times \mathcal{U}$ satisfying (3.15b)-(3.15c). We note that (3.15b) is the controlled FPK PDE which governs the flow of the joint PDF associated with the SDE (3.14b). Conceptually, problem (3.15) is to problem (3.14) what problem (4.3) is to problem (3.1).

3.3.1 Existence and Uniqueness

To establish the existence and uniqueness of minimizer for (3.15), we consider the convexity of the cost functional (3.15a). For this purpose, we perform a change of variable $(\rho, \boldsymbol{u}) \mapsto (\rho, \boldsymbol{m})$ by setting $\boldsymbol{m} := \rho\boldsymbol{u}$. By direct substitution, we obtain the following reformulation of (3.15):

$$\underset{(\rho,\boldsymbol{m})}{\text{minimize}} \quad \frac{1}{2}\int_0^1 \int_{\mathbb{R}^n} J(\rho,\boldsymbol{m})\, \mathrm{d}\boldsymbol{x}\, \mathrm{d}t \tag{3.16a}$$

$$\text{subject to} \quad \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho\boldsymbol{f} + \boldsymbol{B}(t)\boldsymbol{m}) - \epsilon \mathbf{1}^\top (\boldsymbol{D}(t) \odot \mathbf{Hess}\,(\rho))\,\mathbf{1} = 0, \tag{3.16b}$$

$$\rho(\boldsymbol{x},0) = \rho_0(\boldsymbol{x}), \quad \rho(\boldsymbol{x},1) = \rho_1(\boldsymbol{x}), \tag{3.16c}$$

where

$$J(\rho, \boldsymbol{m}) := \begin{cases} \| \boldsymbol{m} \|_2^2 / \rho & \text{if} \quad \rho > 0, \\ 0 & \text{if} \quad (\boldsymbol{m}, \rho) = (\mathbf{0}, 0), \\ +\infty & \text{otherwise.} \end{cases}$$

Since $J(\rho, \boldsymbol{m})$ is the perspective function of the strictly convex map $\boldsymbol{m} \mapsto \|\boldsymbol{m}\|_2^2$, hence J is jointly strictly convex in $(\rho, \boldsymbol{m})$. The constraints (3.16b)-(3.16c) are linear in $(\rho, \boldsymbol{m})$. Thus, the convex optimization problem (3.16), and therefore problem (3.15), admits a unique minimizing pair $(\rho^{\text{opt}}, \boldsymbol{u}^{\text{opt}})$.

3.3.2 Conditions for Optimality

In this Section, we show that the first order optimality conditions for (3.15) correspond to a coupled system of nonlinear PDEs. Consider the Lagrangian associated with (3.15):

$$\begin{aligned} \mathcal{L}(\rho, \boldsymbol{u}, \psi) := \int_0^1 \int_{\mathbb{R}^n} \Big\{ \frac{1}{2} \|\boldsymbol{u}(\boldsymbol{x}, t)\|_2^2 \, \rho(\boldsymbol{x}, t) + \psi(\boldsymbol{x}, t) \times \Big(\frac{\partial \rho}{\partial t} \\ + \nabla \cdot ((\boldsymbol{f} + \boldsymbol{B}(t)\boldsymbol{u}) \, \rho(\boldsymbol{x}, t)) - \epsilon \mathbf{1}^\top (\boldsymbol{D}(t) \odot \mathbf{Hess}\,(\rho)) \, \mathbf{1} \Big) \Big\} \, \mathrm{d}\boldsymbol{x} \, \mathrm{d}t, \end{aligned} \tag{3.17}$$

where $\psi(\boldsymbol{x}, t)$ is a $C^1\left(\mathbb{R}^n; \mathbb{R}_{>0}\right)$ Lagrange multiplier. Let

$$\begin{aligned} \mathcal{P}_{01}\left(\mathbb{R}^n\right) := \Big\{ \rho(\boldsymbol{x}, t) \mid \rho \geq 0, \int_{\mathbb{R}^n} \rho \, \mathrm{d}\boldsymbol{x} = 1, \\ \rho(\boldsymbol{x}, 0) = \rho_0, \rho(\boldsymbol{x}, 1) = \rho_1 \Big\}. \end{aligned} \tag{3.18}$$

Performing the unconstrained minimization of the Lagrangian $\mathcal{L}$ over $\mathcal{P}_{01}\left(\mathbb{R}^n\right) \times \mathcal{U}$ yields the following result (proof in Appendix A.3).

Proposition 5. *The pair $(\rho^{\text{opt}}(\boldsymbol{x}, t), \boldsymbol{u}^{\text{opt}}(\boldsymbol{x}, t))$ that solves (3.15), must satisfy*

the system of coupled PDEs

$$\frac{\partial \psi}{\partial t} + \frac{1}{2}\|\boldsymbol{B}(t)^\top \nabla \psi\|_2^2 + \langle \nabla \psi, \boldsymbol{f} \rangle = -\epsilon \langle \boldsymbol{D}(t), \mathbf{Hess}\,(\psi) \rangle, \tag{3.19a}$$

$$\frac{\partial}{\partial t}\rho^{\mathrm{opt}} + \nabla \cdot (\rho^{\mathrm{opt}}(\boldsymbol{f} + \boldsymbol{B}(t)^\top \nabla \psi)) = \epsilon \mathbf{1}^\top \Big(\boldsymbol{D}(t) \odot \mathbf{Hess}\left(\rho^{\mathrm{opt}}\right) \Big) \mathbf{1}, \tag{3.19b}$$

with boundary conditions

$$\rho^{\mathrm{opt}}(\boldsymbol{x}, 0) = \rho_0(\boldsymbol{x}), \quad \rho^{\mathrm{opt}}(\boldsymbol{x}, 1) = \rho_1(\boldsymbol{x}), \tag{3.20}$$

and

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

Remark 3. *Notice that the system of coupled PDEs (3.19) is the same as that appearing in classical mean field games, but instead of the usual boundary conditions (see e.g., [139, equation (2)])*

$$\rho^{\mathrm{opt}}(\boldsymbol{x}, 0) = \rho_0(\boldsymbol{x}) \text{ (given)}, \quad \psi(\boldsymbol{x}, 1) = \psi_1(\boldsymbol{x}) \text{ (given)}, \tag{3.21}$$

our boundary conditions are given by (3.20).

The PDE (3.19a) is the Hamilton-Jacobi-Bellman (HJB) equation while the PDE (3.19b) is the controlled FPK equation. Computing the optimal pair $(\rho^{\rho^{\mathrm{opt}}}, \boldsymbol{u}^{\rho^{\mathrm{opt}}})$, or equivalently $(\rho^{\rho^{\mathrm{opt}}}, \psi)$ form (3.19)-(3.20) calls for solving a system of coupled nonlinear PDEs with atypical boundary conditions, and is challenging in general.

To tackle the issues of coupling and nonlinearity, we resort to the *Hopf-Cole transform* [68,114] given by (5.50). The geometric interpretation of this transform for optimal control problems were investigated in [140, 141]. Similar ideas have

also appeared in the stochastic control literature [76, 85]. In our context, this transform allows to convert the HJB-FPK system of coupled nonlinear PDEs (3.19)-(3.20) into a system of *boundary-coupled linear PDEs*. We summarize this in the following Theorem (proof in Appendix A.4).

Theorem 6. ***(Hopf-Cole transform)*** *Given* $\boldsymbol{f}$, ϵ, ρ_0, ρ_1, *consider the Hopf-Cole transform* $(\rho^{\mathrm{opt}}, \psi) \mapsto (\varphi, \hat{\varphi})$ *defined via (5.50), applied to the system of nonlinear PDEs (3.19)-(3.20). Then the pair* $(\varphi, \hat{\varphi})$ *satisfies the system of linear PDEs*

$$\frac{\partial \varphi}{\partial t} = -\langle \nabla \varphi, \boldsymbol{f} \rangle - \epsilon \langle \boldsymbol{D}(t), \mathbf{Hess}\,(\varphi) \rangle, \tag{3.22a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = -\nabla \cdot (\hat{\varphi} \boldsymbol{f}) + \epsilon \mathbf{1}^{\top} (\boldsymbol{D}(t) \odot \mathbf{Hess}\,(\hat{\varphi}))\, \mathbf{1}, \tag{3.22b}$$

with boundary conditions (borrowing notations (3.7))

$$\varphi_0(\boldsymbol{x}) \hat{\varphi}_0(\boldsymbol{x}) = \rho_0(\boldsymbol{x}), \quad \varphi_1(\boldsymbol{x}) \hat{\varphi}_1(\boldsymbol{x}) = \rho_1(\boldsymbol{x}). \tag{3.23}$$

Moreover, the optimal controlled state PDF $\rho^{\mathrm{opt}}(\boldsymbol{x}, t)$ *solving (3.15) is given by (3.11a). The optimal control for the same is given by*

$$\boldsymbol{u}^{\mathrm{opt}}(\boldsymbol{x}, t) = 2\epsilon \boldsymbol{B}(t)^{\top} \nabla \log \varphi.$$

We note that (3.22a) is the *backward Kolmogorov equation* in variable φ, and (4.4a) is the *forward Kolmogorov* or the FPK equation in variable $\hat{\varphi}$, associated with the uncontrolled nonlinear SDE

$$\mathrm{d}\boldsymbol{x}(t) = \boldsymbol{f}(\boldsymbol{x}, t)\, \mathrm{d}t + \sqrt{2\epsilon}\, \boldsymbol{B}(t)\, \mathrm{d}\boldsymbol{w}(t). \tag{3.24}$$

Indeed, (3.22) generalizes the forward-backward heat equations (3.5). As expected, setting $\boldsymbol{f} \equiv \mathbf{0}$ and $\boldsymbol{B}(t) \equiv \boldsymbol{I}$ in (3.22) recovers (3.5). Following the nomenclature in Section 3.2.1, the pair $(\varphi, \hat{\varphi})$ solving (3.22)-(3.23) defines the *Schrödinger factors* for problem (3.15).

The essence of Theorem 6 is that instead of solving the system of coupled nonlinear PDEs (3.19), we can solve the system of linear PDEs (3.22) provided that we compute the pair $(\varphi_1, \hat{\varphi}_0)$ which serves as the endpoint data for

$$\frac{\partial \varphi}{\partial t} = -\langle \nabla \varphi, \boldsymbol{f} \rangle - \epsilon \langle \boldsymbol{D}(t), \mathbf{Hess}\,(\varphi) \rangle,\ \varphi(\boldsymbol{x}, 1) = \varphi_1(\boldsymbol{x}), \tag{3.25a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = -\nabla \cdot (\hat{\varphi} \boldsymbol{f}) + \epsilon \mathbf{1}^{\top} (\boldsymbol{D}(t) \odot \mathbf{Hess}\,(\hat{\varphi}))\, \mathbf{1},\ \hat{\varphi}(\boldsymbol{x}, 0) = \hat{\varphi}_0(\boldsymbol{x}). \tag{3.25b}$$

Denoting the forward and backward Kolmogorov operators in (3.25) as L_{FK} and L_{BK} respectively, we can write (3.25) succinctly as an infinite dimensional two point boundary value problem

$$\frac{\partial}{\partial t} \begin{pmatrix} \varphi \\ \hat{\varphi} \end{pmatrix} = \begin{pmatrix} L_{\mathrm{BK}} & 0 \\ 0 & L_{\mathrm{FK}} \end{pmatrix} \begin{pmatrix} \varphi \\ \hat{\varphi} \end{pmatrix},\ \begin{pmatrix} \varphi(\boldsymbol{x}, 1) \\ \hat{\varphi}(\boldsymbol{x}, 0) \end{pmatrix} = \begin{pmatrix} \varphi_1(\boldsymbol{x}) \\ \hat{\varphi}_0(\boldsymbol{x}) \end{pmatrix}. \tag{3.26}$$

Solving (3.25) or (3.26) will yield $(\varphi(\boldsymbol{x}, t), \hat{\varphi}(\boldsymbol{x}, t))$ for all $t \in [0, 1]$, which in turn can be used to determine $(\rho^{\rho^{\mathrm{opt}}}(\boldsymbol{x}, t), \boldsymbol{u}^{\rho^{\mathrm{opt}}}(\boldsymbol{x}, t))$ via (3.11). Notice that this computational pipeline hinges on the ability to first compute the pair $(\varphi_1, \hat{\varphi}_0)$ using (3.23) *and* (3.25), and then utilize $(\varphi_1, \hat{\varphi}_0)$ to solve (3.25) or (3.26). However, unlike the situation in Section 3.2, the difficulty now is that the closed-form expression of the Markov kernel associated with (3.24) is not available in general. This prevents us from setting up a *Schrödinger system* like (3.10) to solve for the pair $(\varphi_1, \hat{\varphi}_0)$.

In the following Section 3.4, we will reformulate the Schrödinger system for two cases: when $\boldsymbol{f}$ is gradient of a potential, and when the prior drift has mixed conservative-dissipative structure, i.e., a *degenerate* diffusion of the form (3.14b) with $\boldsymbol{u}, \boldsymbol{w} \in \mathbb{R}^m$, the state space dimension $n \equiv 2m$, and a constant input matrix $\boldsymbol{B} \in \mathbb{R}^{n \times m}$. In Section 3.5, we will then show how these reformulations can harness the recently-introduced proximal recursions [41, 42] for computing the pair $(\varphi_1, \hat{\varphi}_0)$, and consequently the pair $(\varphi(\boldsymbol{x}, t), \hat{\varphi}(\boldsymbol{x}, t))$. This is appealing since these proximal algorithms do not require spatial discretization or function approximation, and instead evolve weighted scattered point cloud data. Therefore, such algorithms hold the promise for solving the SBP with nonlinear prior dynamics in high dimensions when no analytical handle on the Markov kernel is available.

3.4 Reformulation of the Schrödinger Systems

Next, we provide algorithmically tractable reformulations of the Schrödinger systems for two important types of prior nonlinear dynamics: gradient drift in Section 3.4.1, and mixed conservative-dissipative drift (i.e., degenerate diffusion, see e.g., [32, Sec. 7, 8]) in Section 3.4.2. For clarity of exposition, in Section 3.4.1, we consider $\boldsymbol{B}(t) \equiv \boldsymbol{I}$. In Section 3.4.2, the matrix $\boldsymbol{B}$ will be non-identity.

3.4.1 The Case of Gradient Drift

We now consider $\boldsymbol{f}$ to be a gradient vector field, i.e., $\boldsymbol{f} = -\nabla V(\boldsymbol{x})$ for some $C^2(\mathbb{R}^n)$ function $V : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0}$, and $\boldsymbol{B}(t) = \boldsymbol{I}$. This reduces (3.15) to the

following:

$$\inf_{(\rho,\boldsymbol{u})} \quad \frac{1}{2}\int_0^1 \int_{\mathbb{R}^n} \|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2\, \rho(\boldsymbol{x},t)\, \mathrm{d}\boldsymbol{x}\, \mathrm{d}t, \tag{3.27a}$$

$$\text{subject to} \quad \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho(\boldsymbol{u} - \nabla V)) = \epsilon \Delta \rho, \tag{3.27b}$$

$$\rho(\boldsymbol{x},0) = \rho_0(\boldsymbol{x}), \quad \rho(\boldsymbol{x},1) = \rho_1(\boldsymbol{x}). \tag{3.27c}$$

Applying Proposition 5 and Theorem 6 to (3.27), we arrive at the system of linear PDEs for $(\varphi, \hat{\varphi})$ given by

$$\frac{\partial \varphi}{\partial t} = \langle \nabla \varphi, \nabla V \rangle - \epsilon \Delta \varphi, \tag{3.28a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = \nabla \cdot (\hat{\varphi} \nabla V) + \epsilon \Delta \hat{\varphi}, \tag{3.28b}$$

with boundary conditions

$$\varphi_0(\boldsymbol{x})\hat{\varphi}_0(\boldsymbol{x}) = \rho_0(\boldsymbol{x}), \quad \varphi_1(\boldsymbol{x})\hat{\varphi}_1(\boldsymbol{x}) = \rho_1(\boldsymbol{x}). \tag{3.29}$$

As expected, (3.28a) is a backward Kolmogorov PDE, and (3.28b) is a forward Kolmogorov or FPK* PDE.

We would like to exploit the structure of (3.28) to develop an algorithm that computes the pair $(\varphi_1, \hat{\varphi}_0)$ which can then serve as the terminal and initial data for the following system solving for the pair $(\varphi(\boldsymbol{x},t), \hat{\varphi}(\boldsymbol{x},t))$:

$$\frac{\partial \varphi}{\partial t} = \langle \nabla \varphi, \nabla V \rangle - \epsilon \Delta \varphi, \quad \varphi(\boldsymbol{x},1) = \varphi_1(\boldsymbol{x}), \tag{3.30a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = \nabla \cdot (\hat{\varphi} \nabla V) + \epsilon \Delta \hat{\varphi}, \quad \hat{\varphi}(\boldsymbol{x},0) = \hat{\varphi}_0(\boldsymbol{x}). \tag{3.30b}$$

*This particular instance of the FPK PDE (3.28b) is also known as the "Smoluchowski equation" [51, Ch. II.4.(vi)].

After having $(\varphi(\boldsymbol{x},t),\hat{\varphi}(\boldsymbol{x},t))$ from (3.30), we can obtain the pair $(\rho^{\rho^{\text{opt}}},\boldsymbol{u}^{\rho^{\text{opt}}})$ via (3.11). To this end, the following result (proof in Appendix A.5) will be an important step.

Theorem 7. *Given $V(\boldsymbol{x}),\epsilon,\varphi_1(\boldsymbol{x})$, consider the terminal value problem (TVP) (3.30a) in unknown $\varphi(\boldsymbol{x},t)$. Let $s:=1-t$, and $q(\boldsymbol{x},s):=\varphi(\boldsymbol{x},t)=\varphi(\boldsymbol{x},1-s)$. Then q satisfies the initial value problem (IVP)*

$$\frac{\partial q}{\partial s}=-\langle\nabla q,\nabla V\rangle+\epsilon\Delta q,\quad q(\boldsymbol{x},0)=\varphi_1(\boldsymbol{x}). \tag{3.31}$$

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

$$\frac{\partial p}{\partial s}=\nabla\cdot(p\nabla V)+\epsilon\Delta p,\quad p(\boldsymbol{x},0)=q(\boldsymbol{x},0)\exp\left(-V(\boldsymbol{x})/\epsilon\right), \tag{3.32}$$

where q is a smooth solution of (3.31), and V is such that

$$\int_{\mathbb{R}^n}q(\boldsymbol{x},0)\exp\left(-V(\boldsymbol{x})/\epsilon\right)\,\mathrm{d}\boldsymbol{x}<\infty,\quad \textit{for all}\ \epsilon>0.$$

Thanks to Theorem 7, solving (3.30) amounts to solving the system

$$\frac{\partial\hat{\varphi}}{\partial t}=\nabla\cdot(\hat{\varphi}\nabla V)+\epsilon\Delta\hat{\varphi},\quad \hat{\varphi}(\boldsymbol{x},0)=\hat{\varphi}_0(\boldsymbol{x}), \tag{3.33a}$$

$$\frac{\partial p}{\partial s}=\nabla\cdot(p\nabla V)+\epsilon\Delta p,\quad p(\boldsymbol{x},0)=\varphi_1(\boldsymbol{x})\exp\left(-V(\boldsymbol{x})/\epsilon\right). \tag{3.33b}$$

Notice that (3.33a) and (3.33b) involve the exact same PDE with different initial conditions, to be integrated in different time coordinates t and s, where $t=1-s$. This implies that availability of a single FPK IVP solver is enough to set up a fixed point recursion for the pair $(\varphi_1,\hat{\varphi}_0)$ via (3.33). Once p is computed from

(3.33b), we can recover φ by the relation

$$\varphi(\boldsymbol{x}, t) = \varphi(\boldsymbol{x}, 1-s) = p(\boldsymbol{x}, s)/\exp\left(-V(\boldsymbol{x})/\epsilon\right) \tag{3.34}$$

3.4.2 The Case of Mixed Conservative-Dissipative Drift

We now consider a *degenerate* diffusion of the form (3.14b) with $\boldsymbol{u}, \boldsymbol{w} \in \mathbb{R}^m$. The state $\boldsymbol{x}$ consists of two sub-vectors $\boldsymbol{\xi}, \boldsymbol{\eta} \in \mathbb{R}^m$, i.e., $\boldsymbol{x} := (\boldsymbol{\xi}, \boldsymbol{\eta})^\top \in \mathbb{R}^n$ with $n \equiv 2m$. The controlled SDE is given by

$$\begin{aligned}\begin{pmatrix} \mathrm{d}\boldsymbol{\xi} \\ \mathrm{d}\boldsymbol{\eta} \end{pmatrix} = \Big\{ \underbrace{\begin{pmatrix} \boldsymbol{\eta} \\ -\nabla_{\boldsymbol{\xi}} V\left(\boldsymbol{\xi}\right) - \kappa\boldsymbol{\eta} \end{pmatrix}}_{\boldsymbol{f}(\boldsymbol{x})} + \underbrace{\begin{pmatrix} \boldsymbol{0}_{m\times m} \\ \boldsymbol{I}_{m\times m} \end{pmatrix}}_{\boldsymbol{B}} \boldsymbol{u} \Big\} \mathrm{d}t \\ + \sqrt{2\epsilon\kappa} \underbrace{\begin{pmatrix} \boldsymbol{0}_{m\times m} \\ \boldsymbol{I}_{m\times m} \end{pmatrix}}_{\boldsymbol{B}} \mathrm{d}\boldsymbol{w}, \quad \kappa > 0,\end{aligned} \tag{3.35}$$

i.e., we consider the SBP (3.14) with the constraint (3.14b) replaced by (3.35). Here, we assume that $V \in C^2(\mathbb{R}^m)$, $\inf V > -\infty$, and that $\mathbf{Hess}\,(V)$ is uniformly lower bounded. To glean physical motivation, one may consider $m = 3$ and think of $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$ as three-dimensional position and velocity vectors, respectively [41, Sec. V.B]; see also [148, Examples 2,3]. So, (3.35) is a controlled Langevin equation. We are thus led to solve an instance of (3.15) wherein for the constraint (3.15b), we use the vector field $\boldsymbol{f}$ and the matrix $\boldsymbol{B}$ as shown above, and set the diffusion tensor as $\boldsymbol{D} := \kappa\boldsymbol{B}\boldsymbol{B}^\top \in \mathbb{R}^{n\times n}$.

Using Proposition 5, the first order conditions of optimality for this variant of

SBP yields the coupled HJB-FPK system:

$$\frac{\partial \psi}{\partial t} = -\langle \boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}} \psi \rangle + \langle \nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa \boldsymbol{\eta}, \nabla_{\boldsymbol{\eta}} \psi \rangle - \epsilon \kappa \Delta_{\boldsymbol{\eta}} \psi - \frac{1}{2} \| \nabla_{\boldsymbol{\eta}} \psi \|_2^2, \tag{3.36a}$$

$$\frac{\partial \rho^{\text{opt}}}{\partial t} = -\langle \boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}} \rho^{\text{opt}} \rangle + \nabla_{\boldsymbol{\eta}} \cdot \left(\rho^{\text{opt}} \left(\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa \boldsymbol{\eta} - \nabla_{\boldsymbol{\eta}} \psi \right) \right) + \epsilon \kappa \Delta_{\boldsymbol{\eta}} \rho^{\text{opt}}, \tag{3.36b}$$

with boundary conditions $\rho^{\text{opt}}(\boldsymbol{x}, 0) = \rho_0(\boldsymbol{x})$ and $\rho^{\text{opt}}(\boldsymbol{x}, 1) = \rho_1(\boldsymbol{x})$, and the optimal control $\boldsymbol{u}^{\text{opt}}(\boldsymbol{x}, t) = \nabla_{\boldsymbol{\eta}} \psi(\boldsymbol{x}, t)$.

Following Theorem 6, we now apply the Hopf-Cole transform $(\psi, \rho^{\text{opt}}) \mapsto (\varphi, \hat{\varphi})$ given by (5.50), to the system (3.36). This results in the following system of linear PDEs for $(\varphi, \hat{\varphi})$:

$$\frac{\partial \varphi}{\partial t} = -\langle \boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}} \varphi \rangle + \langle \nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa \boldsymbol{\eta}, \nabla_{\boldsymbol{\eta}} \varphi \rangle - \epsilon \kappa \Delta_{\boldsymbol{\eta}} \varphi, \tag{3.37a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = -\langle \boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}} \hat{\varphi} \rangle + \nabla_{\boldsymbol{\eta}} \cdot \left(\hat{\varphi} \left(\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa \boldsymbol{\eta} \right) \right) + \epsilon \kappa \Delta_{\boldsymbol{\eta}} \hat{\varphi}, \tag{3.37b}$$

with boundary conditions (3.29). In particular, (3.37a) is a backward Kolmogorov PDE, and (3.37b) is a forward Kolmogorov or FPK[†] PDE.

Next, we establish a result (proof in Appendix A.6) for (3.37) that is similar in flavor to Theorem 7, and will be useful for designing proximal algorithm in Section 3.5.

Theorem 8. *Given $\epsilon, \varphi_1(\boldsymbol{\xi}, \boldsymbol{\eta})$, consider the TVP comprising of the PDE (3.37a) in variable φ together with the boundary condition $\varphi(\boldsymbol{\xi}, \boldsymbol{\eta}, 1) = \varphi_1(\boldsymbol{\xi}, \boldsymbol{\eta})$. Let*

[T] This particular instance of the FPK PDE (3.37b) is also known as the "kinetic Fokker-Planck equation" [212].

$s := 1 - t$, and $q(\boldsymbol{\xi}, \boldsymbol{\eta}, s) := \varphi(\boldsymbol{\xi}, \boldsymbol{\eta}, t) = \varphi(\boldsymbol{\xi}, \boldsymbol{\eta}, 1 - s)$. Then q satisfies the IVP

$$\frac{\partial q}{\partial s} = \langle \boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}} q \rangle - \langle \nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa \boldsymbol{\eta}, \nabla_{\boldsymbol{\eta}} q \rangle + \epsilon \kappa \, \Delta_{\boldsymbol{\eta}} q, \tag{3.38a}$$

$$q(\boldsymbol{\xi}, \boldsymbol{\eta}, 0) = \varphi_1(\boldsymbol{\xi}, \boldsymbol{\eta}). \tag{3.38b}$$

Further, let $\boldsymbol{\vartheta} := -\boldsymbol{\eta}$, and

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

Then, $p(\boldsymbol{\xi}, \boldsymbol{\vartheta}, s) := \widetilde{p}(\boldsymbol{\xi}, -\boldsymbol{\eta}, s)$ satisfies the IVP

$$\frac{\partial p}{\partial s} = -\langle \boldsymbol{\vartheta}, \nabla_{\boldsymbol{\xi}} p \rangle + \nabla_{\boldsymbol{\vartheta}} \cdot \left(p \left(\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa \boldsymbol{\vartheta}\right)\right) + \epsilon \kappa \, \Delta_{\boldsymbol{\vartheta}} p, \tag{3.39a}$$

$$p(\boldsymbol{\xi}, \boldsymbol{\vartheta}, 0) = q(\boldsymbol{\xi}, \boldsymbol{\eta}, 0) \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right), \tag{3.39b}$$

where q is a smooth solution of (3.38), and V is such that

$$\int_{\mathbb{R}^m} \int_{\mathbb{R}^m} q(\boldsymbol{\xi}, \boldsymbol{\eta}, 0) \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right) \, \mathrm{d}\boldsymbol{\xi} \, \mathrm{d}\boldsymbol{\eta} < \infty,$$

for all $\epsilon > 0$.

Thanks to Theorem 8, solving the system (3.37) with boundary conditions

$$\varphi(\boldsymbol{\xi}, \boldsymbol{\eta}, 1) = \varphi_1(\boldsymbol{\xi}, \boldsymbol{\eta}), \quad \hat{\varphi}(\boldsymbol{\xi}, \boldsymbol{\eta}, 0) = \hat{\varphi}_0(\boldsymbol{\xi}, \boldsymbol{\eta}),$$

reduces to solving the system

$$\frac{\partial \hat{\varphi}}{\partial t} = -\langle \boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}} \hat{\varphi} \rangle + \nabla_{\boldsymbol{\eta}} \cdot \left(\hat{\varphi} \left(\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa \boldsymbol{\eta}\right)\right) + \epsilon \kappa \Delta_{\boldsymbol{\eta}} \hat{\varphi}, \tag{3.40a}$$

$$\frac{\partial p}{\partial s} = -\langle \boldsymbol{\vartheta}, \nabla_{\boldsymbol{\xi}} p \rangle + \nabla_{\boldsymbol{\vartheta}} \cdot \left(p \left(\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa \boldsymbol{\vartheta}\right)\right) + \epsilon \kappa \, \Delta_{\boldsymbol{\vartheta}} p, \tag{3.40b}$$

with initial conditions

$$\hat{\varphi}(\boldsymbol{\xi},\boldsymbol{\eta},0) = \hat{\varphi}_0(\boldsymbol{\xi},\boldsymbol{\eta}), \tag{3.41a}$$

$$\begin{aligned} p(\boldsymbol{\xi},\boldsymbol{\vartheta},0) &= q(\boldsymbol{\xi},\boldsymbol{\eta},0)\exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right) \\ &= \varphi_1(\boldsymbol{\xi},-\boldsymbol{\vartheta})\exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\vartheta}\|_2^2 + V(\boldsymbol{\xi})\right)\right). \end{aligned} \tag{3.41b}$$

Similar to Section 3.4.1, notice that (3.40a) and (3.40b) involve the exact same kinetic Fokker-Planck PDE with different initial conditions, to be integrated in different time coordinates t and s, where $t = 1 - s$. This implies that availability of a single kinetic FPK IVP solver is enough to set up a fixed point recursion for the pair $(\varphi_1, \hat{\varphi}_0)$ via (3.40). Once $p(\boldsymbol{\xi},\boldsymbol{\vartheta},s)$ is computed from (3.40b), we can recover $\varphi(\boldsymbol{\xi},\boldsymbol{\eta},t)$ using the relation

$$\varphi(\boldsymbol{\xi},\boldsymbol{\eta},t) = \varphi(\boldsymbol{\xi},\boldsymbol{\eta},1-s) = p(\boldsymbol{\xi},\boldsymbol{\vartheta},s)/\exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right). \tag{3.42}$$

Remark 4. *While both Theorem 7 and Theorem 8 reduce the respective Schrödinger systems–which are two point forward-backward Kolmogorov systems–to forward-forward Kolmogorov IVPs, it is important to recognize the subtle difference between the transformations employed in these two theorems. Unlike (3.28), the Schrödinger system (3.37) is not reversible (see e.g., [180, Ch. 6.1]). To account this, Theorem 8 used a spatial transform $(\boldsymbol{\xi},\boldsymbol{\eta}) \mapsto (\boldsymbol{\xi},\boldsymbol{\vartheta})$, in addition to the time transform $t \mapsto s$. This resulted in a sequence of change-of-variables $\varphi \mapsto q \mapsto \widetilde{p} \mapsto p$ in Theorem 8. To put the matter in perspective, Theorem 7 involved the change-of-variables $\varphi \mapsto q \mapsto p$.*

3.5 Wasserstein Proximal Algorithms

In this Section, we will design variational recursions to solve the SBP with gradient or mixed conservative-dissipative drift. For the case of the gradient drift, the variational recursions will iteratively solve the system of IVPs (3.33). For the case of the mixed conservative-dissipative drift, the variational recursions will iteratively solve the system of IVPs (3.40)-(3.41). The main idea here is to design certain Lyapunov functionals (usually referred to as the 'free energy functionals') such that the FPK PDEs in (3.33) and (3.40) can be recast as the gradient flow of the respective Lyapunov functionals w.r.t. suitable metric on the manifold $\mathcal{P}_2(\mathbb{R}^n)$, which is the space of all PDFs supported on $\mathbb{R}^n$ with finite second moment. For theoretical developments along this line, we refer the readers to [6]. Our recent works [41, 42] showed that such infinite dimensional gradient flow structure can be exploited for designing grid-less recursive algorithms to solve for the transient solutions of the FPK PDE. We collect the key conceptual ingredients next.

3.5.1 Proximal Operators for Infinite Dimensional Gradient Flows

First, we introduce the 2-Wasserstein metric W, that will play an important role in the development that follows.

Definition 4. *(**2-Wasserstein metric**) The 2-Wasserstein metric between two probability measures μ_0, μ_1, each supported on $\mathbb{R}^n$, is*

$$W(\mu_0, \mu_1) := \left(\underset{m \in M(\mu_0, \mu_1)}{inf} \int_{\mathbb{R}^n \times \mathbb{R}^n} \|\boldsymbol{x} - \boldsymbol{y}\|_2^2 \, \mathrm{d}m(\boldsymbol{x}, \boldsymbol{y}) \right)^{\frac{1}{2}}, \tag{3.43}$$

where $\mathcal{M}(\mu_0, \mu_1)$ is the set of all joint probability measures (i.e., couplings) supported on the product space $\mathbb{R}^n \times \mathbb{R}^n$ whose marginals are μ_0 and μ_1, respectively. Whenever μ_0 and μ_1 are absolutely continuous, their respective PDFs ρ_0 and ρ_1 exist (i.e., $\mathrm{d}\mu_0 = \rho_0(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$, $\mathrm{d}\mu_1 = \rho_1(\boldsymbol{y})\mathrm{d}\boldsymbol{y}$), and we use the equivalent notation $W(\rho_0, \rho_1)$.

Although Definition 4 introduces the 2-Wasserstein metric W as a static variational problem, Brenier and Benamou [21] pointed out that W^2 admits the following equivalent dynamic reformulation:

$$W^2(\rho_0, \rho_1) = \inf_{\rho, \boldsymbol{u}} \quad \int_0^1 \int_{\mathbb{R}^n} \frac{1}{2} \|\boldsymbol{u}(\boldsymbol{x}, t)\|_2^2 \, \rho(\boldsymbol{x}, t) \, \mathrm{d}\boldsymbol{x}\mathrm{d}t, \tag{3.44a}$$

$$\text{subject to } \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{u}) = 0, \tag{3.44b}$$

$$\rho(\boldsymbol{x}, 0) = \rho_0(\boldsymbol{x}), \quad \rho(\boldsymbol{x}, 1) = \rho_1(\boldsymbol{x}), \tag{3.44c}$$

which can be interpreted as the minimum energy control problem to steer the PDF ρ_0 to ρ_1 over the time interval $[0, 1]$ subject to the deterministic dynamics $\dot{\boldsymbol{x}} = \boldsymbol{u}$. Formally, the classical SBP (4.3) can be viewed as a dynamic stochastic regularization of (3.44), and this regularization is captured by the diffusion term $\epsilon\Delta\rho$ in (3.2b). These ideas can be made rigorous by proving weak convergence of measures [143, 144, 160, 161] in the $\epsilon \downarrow 0$ limit, i.e., the solution of (4.3) converges to that of (3.44) in the said limit.

Now the idea is to compute the flows generated by the IVPs (3.33) or (3.40), via variational recursions over discrete time index pair $(t_{k-1}, s_{k-1}) := ((k-1)\tau, (k -$

$1)\sigma)$, $k \in \mathbb{N}$, given by the following map

$$\begin{pmatrix} \hat{\phi}_{t_{k-1}} \\ \varpi_{s_{k-1}} \end{pmatrix} \mapsto \begin{pmatrix} \hat{\phi}_{t_k} \\ \varpi_{s_k} \end{pmatrix} := \begin{pmatrix} \underset{\hat{\phi} \in P_2(R^n)}{arg\,inf}\ \frac{1}{2} d^2\left(\hat{\phi}_{t_{k-1}}, \hat{\phi}\right) + \tau F(\hat{\phi}) \\ \underset{\varpi \in P_2(R^n)}{arg\,inf}\ \frac{1}{2} d^2\left(\varpi_{s_{k-1}}, \varpi\right) + \sigma F(\varpi) \end{pmatrix}, \tag{3.45}$$

where $\tau, \sigma > 0$ are time-step sizes, and the functionals (d, F) are chosen such that

$$\hat{\phi}_{t_{k-1}}(\boldsymbol{x}) \to \hat{\varphi}(\boldsymbol{x}, t = (k-1)\tau) \qquad \text{in } L^1(\mathbb{R}^n) \text{ as } \quad \tau \downarrow 0, \tag{3.46a}$$

$$\varpi_{s_{k-1}}(\boldsymbol{x}) \to p(\boldsymbol{x}, s = (k-1)\sigma) \qquad \text{in } L^1(\mathbb{R}^n) \text{ as } \quad \sigma \downarrow 0. \tag{3.46b}$$

In words, the variational recursions (3.45) are designed to approximate the flows generated by (3.33) or (3.40) in the small time-step limit.

The variational maps appearing in (3.45) resemble *proximal operators* [19, 176] well-known in the Euclidean and general Hilbert space optimization literature. In our setting too, d will be a distance metric on $\mathcal{P}_2(\mathbb{R}^n)$, and F will be a Lyapunov functional, i.e., non-negative and decreasing along the flows generated by the IVPs (3.33) or (3.40), and thus motivate defining the proximal operators

$$\text{prox}^d_{\tau F}\left(\hat{\phi}_{t_{k-1}}\right) := \underset{\hat{\phi} \in P_2(R^n)}{arg\,inf}\ \frac{1}{2} d^2\left(\hat{\phi}_{t_{k-1}}, \hat{\phi}\right) + \tau F(\hat{\phi}), \tag{3.47a}$$

$$\text{prox}^d_{\sigma F}\left(\varpi_{s_{k-1}}\right) := \underset{\varpi \in P_2(R^n)}{arg\,inf}\ \frac{1}{2} d^2\left(\varpi_{s_{k-1}}, \varpi\right) + \sigma F(\varpi), \tag{3.47b}$$

which respectively read as the proximal operator of τF and σF, w.r.t. the metric d. This metric viewpoint allows clear geometric interpretation of the proximal recursions (3.45): they define gradient descent of the functionals τF and σF, measured w.r.t. the distance metric d; see e.g., [41, 100–102]. For (3.33), the

distance d will turn out to be the 2-Wasserstein metric W. For (3.40), the distance metric d will be a variant of W.

3.5.2 Proximal Recursions

We now outline how the proximal recursions (3.47) can be constructed for the two nonlinear drifts of our interest: gradient and mixed conservative-dissipative.

Gradient drift

The seminal paper [125] showed that the flows generated by the FPK PDEs of the form (3.33) can be seen as the gradient descent of the Lyapunov functional

$$F_{\text{gradient}}(\cdot) := \int_{\mathbb{R}^n} V(\boldsymbol{x})(\cdot)\, \mathrm{d}\boldsymbol{x} + \epsilon \int_{\mathbb{R}^n} (\cdot) \log (\cdot)\, \mathrm{d}\boldsymbol{x}, \tag{3.48}$$

w.r.t. the distance metric W in $\mathcal{P}_2(\mathbb{R}^n)$. Here, $(\cdot)$ is a placeholder for either $\hat{\varphi}$ or p in (3.33). From (3.45) and (3.47), the solutions of the IVPs (3.33), denoted by $(\hat{\varphi}(\boldsymbol{x}, t), p(\boldsymbol{x}, s))$, can then be approximated (as in (3.46)) by the following proximal recursions:

$$\begin{pmatrix} \hat{\phi}_{t_k} \\ \varpi_{s_k} \end{pmatrix} = \begin{pmatrix} \text{prox}^W_{\tau F_{\text{gradient}}}\left(\hat{\phi}_{t_{k-1}}\right) \\ \text{prox}^W_{\sigma F_{\text{gradient}}}\left(\varpi_{s_{k-1}}\right) \end{pmatrix}, \quad k \in \mathbb{N}, \tag{3.49}$$

where F_{gradient} is given by (3.48), and the initial conditions:

$$\begin{aligned} \hat{\phi}_{t_0} = \hat{\varphi}(\boldsymbol{x}, 0) &\overset{(3.33\text{a})}{=} \hat{\varphi}_0(\boldsymbol{x}), \\ \varpi_{s_0} = p(\boldsymbol{x}, 0) &\overset{(3.33\text{b})}{=} \varphi_1(\boldsymbol{x}) \exp(-V(\boldsymbol{x})/\epsilon). \end{aligned}$$

In other words, for the IVPs (3.33), the pair $(d, F) \equiv (W, F_{\text{gradient}})$.

To numerically implement the proximal recursions (3.49), we will use the weighted scattered point cloud algorithm developed in our previous work [41, Sec. III]. Since the metric W (see Definition 4) is itself defined as a variational problem, evaluating proximal operators w.r.t. W, requires solving nested functional minimization problems. The algorithm in [41, Sec. III.B] uses an entropic or Sinkhorn regularization followed by block coordinate ascent in the dual space, and provides convergence guarantees via contraction mapping [41, Sec. III.C].

Mixed conservative-dissipative drift

For FPK PDEs of the form (3.40), recall that $\boldsymbol{x} := (\boldsymbol{\xi}, \boldsymbol{\eta})^\top \in \mathbb{R}^n$, $\boldsymbol{y} := (\boldsymbol{\xi}, \boldsymbol{\vartheta})^\top \in \mathbb{R}^n$, and let the "Hamiltonian-like" function $H(\boldsymbol{x}) := \frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi}) = \frac{1}{2}\|\boldsymbol{\vartheta}\|_2^2 + V(\boldsymbol{\xi}) =: H(\boldsymbol{y})$. It can be verified that

$$F^{\text{a}}_{\text{mixed}}(\hat{\varphi}) := \int_{\mathbb{R}^n} H(\boldsymbol{x})\hat{\varphi}(\boldsymbol{x}, t)\ \mathrm{d}\boldsymbol{x} + \epsilon \int_{\mathbb{R}^n} \hat{\varphi}(\boldsymbol{x}, t) \log \hat{\varphi}(\boldsymbol{x}, t)\ \mathrm{d}\boldsymbol{x},$$

and

$$F^{\text{b}}_{\text{mixed}}(p) := \int_{\mathbb{R}^n} H(\boldsymbol{y})p(\boldsymbol{y}, s)\ \mathrm{d}\boldsymbol{y} + \epsilon \int_{\mathbb{R}^n} p(\boldsymbol{y}, s) \log p(\boldsymbol{y}, s)\ \mathrm{d}\boldsymbol{y},$$

are Lyapunov functionals along the flows of (3.40a) and (3.40b), respectively. However, since the degenerate diffusions do not allow statistical reversibility, the flows for (3.40) cannot be recast as the gradient descent of the above functionals w.r.t. W. To circumvent this, consider instead the functionals

$$\widetilde{F}^{\text{a}}_{\text{mixed}}(\hat{\varphi}) := \int_{\mathbb{R}^n} \frac{1}{2}\|\boldsymbol{\eta}\|_2^2 \hat{\varphi}(\boldsymbol{x}, t)\ \mathrm{d}\boldsymbol{x} + \epsilon \int_{\mathbb{R}^n} \hat{\varphi}(\boldsymbol{x}, t) \log \hat{\varphi}(\boldsymbol{x}, t)\ \mathrm{d}\boldsymbol{x}, \tag{3.50a}$$

$$\widetilde{F}^{\text{b}}_{\text{mixed}}(p) := \int_{\mathbb{R}^n} \frac{1}{2}\|\boldsymbol{\vartheta}\|_2^2 p(\boldsymbol{y}, s)\ \mathrm{d}\boldsymbol{y} + \epsilon \int_{\mathbb{R}^n} p(\boldsymbol{y}, s) \log p(\boldsymbol{y}, s)\ \mathrm{d}\boldsymbol{y}. \tag{3.50b}$$

Also, consider the following distance functionals which are modified versions of (3.43):

$$\widetilde{W}_\tau(\mu_0,\mu_1) := \left(\underset{m\in M(\mu_0,\mu_1)}{inf} \int_{\mathbb{R}^n\times\mathbb{R}^n} s_\tau\left(\boldsymbol{x},\overline{\boldsymbol{x}}\right)\, \mathrm{d}m(\boldsymbol{x},\overline{\boldsymbol{x}}) \right)^{\frac{1}{2}}, \tag{3.51a}$$

$$\widetilde{W}_\sigma(\mu_0,\mu_1) := \left(\underset{m\in M(\mu_0,\mu_1)}{inf} \int_{\mathbb{R}^n\times\mathbb{R}^n} s_\sigma\left(\boldsymbol{y},\overline{\boldsymbol{y}}\right)\, \mathrm{d}m(\boldsymbol{y},\overline{\boldsymbol{y}}) \right)^{\frac{1}{2}}, \tag{3.51b}$$

where $\boldsymbol{x} = (\boldsymbol{\xi},\boldsymbol{\eta})^\top$ and $\overline{\boldsymbol{x}} = (\overline{\boldsymbol{\xi}},\overline{\boldsymbol{\eta}})^\top$ are two realizations of the state vector in (3.35); similarly, $\boldsymbol{y} = (\boldsymbol{\xi},\boldsymbol{\vartheta})^\top$, $\overline{\boldsymbol{y}} = (\overline{\boldsymbol{\xi}},\overline{\boldsymbol{\vartheta}})^\top$, and

$$s_\tau\left(\boldsymbol{x},\overline{\boldsymbol{x}}\right) := \|\overline{\boldsymbol{\eta}} - \boldsymbol{\eta} + \tau\nabla V(\boldsymbol{\xi})\|_2^2 + 12\left\|\frac{\overline{\boldsymbol{\xi}}-\boldsymbol{\xi}}{\tau} - \frac{\overline{\boldsymbol{\eta}}+\boldsymbol{\eta}}{2}\right\|_2^2, \tag{3.52a}$$

$$s_\sigma\left(\boldsymbol{y},\overline{\boldsymbol{y}}\right) := \|\overline{\boldsymbol{\vartheta}} - \boldsymbol{\vartheta} + \sigma\nabla V(\boldsymbol{\xi})\|_2^2 + 12\left\|\frac{\overline{\boldsymbol{\xi}}-\boldsymbol{\xi}}{\sigma} - \frac{\overline{\boldsymbol{\vartheta}}+\boldsymbol{\vartheta}}{2}\right\|_2^2. \tag{3.52b}$$

Following [81, Scheme 2b], we then set up the proximal recursions

$$\begin{pmatrix} \hat{\phi}_{t_k} \\ \varpi_{s_k} \end{pmatrix} = \begin{pmatrix} \mathrm{prox}^{\widetilde{W}_\tau}_{\kappa\tau\widetilde{F}^{\mathrm{a}}_{\mathrm{mixed}}}\left(\hat{\phi}_{t_{k-1}}\right) \\ \mathrm{prox}^{\widetilde{W}_\sigma}_{\kappa\sigma\widetilde{F}^{\mathrm{b}}_{\mathrm{mixed}}}\left(\varpi_{s_{k-1}}\right) \end{pmatrix}, \quad k\in\mathbb{N}, \tag{3.53}$$

with initial conditions (3.41). For the above recursions, the results from [81, Theorem 2.4] provide the consistency guarantees (3.46) for the flows generated by (3.40).

The weighted scattered point cloud algorithm from [41] that we mentioned in Section 3.5.2, can be applied to this case too (see [41, Section V.B]), and we will use the same to solve (3.40) by numerically performing the proximal recursion (3.53).

3.5.3 Sinkhorn Proximal Algorithms for Solving (3.49) and (3.53)

Our development so far have reduced solving the SBPs with gradient and mixed conservative-dissipative drifts to that of solving the proximal recursions (3.49) and (3.53), respectively. For numerical computation, we perform these proximal recursions over weighted scattered point clouds with sample size N. Specifically, let $\boldsymbol{x}^i(t)$ be the state vector for the i-th sample at time t, and for $i = 1, \dots, N$, $k \in \mathbb{N}$, define the $N \times 1$ vectors

$$\hat{\boldsymbol{\phi}}^i_{k-1} := \hat{\phi}_{t_{k-1}}\left(\boldsymbol{x}^i(t_{k-1})\right), \tag{3.54a}$$

$$\boldsymbol{\varpi}^i_{k-1} := \varpi_{s_{k-1}}\left(\boldsymbol{x}^i(s_{k-1})\right), \tag{3.54b}$$

wherein the superscript i in the left-hand-side denotes the "i-th component of". The recursions (3.49) and (3.53) are performed over the point clouds

$$\left\{\boldsymbol{x}^i(t_{k-1}), \hat{\boldsymbol{\phi}}^i_{k-1}\right\}_{i=1}^N, \left\{\boldsymbol{x}^i(s_{k-1}), \boldsymbol{\varpi}^i_{k-1}\right\}_{i=1}^N \tag{3.55}$$

respectively, following the algorithm given in [41,42]. To do so, the state vectors $\boldsymbol{x}^i$ are updated by applying the Euler-Maruyama[‡] scheme to the appropriate version of the *uncontrolled* SDE (3.24). For instance, in the gradient drift case, the uncontrolled SDE associated with (3.33) is

$$\mathrm{d}\boldsymbol{x} = -\nabla V(\boldsymbol{x})\,\mathrm{d}t + \sqrt{2\epsilon}\,\mathrm{d}\boldsymbol{w}. \tag{3.56}$$

For $k \in \mathbb{N}$, we suppose that $\{\boldsymbol{x}^i(t_{k-1})\}_{i=1}^N \mapsto \{\boldsymbol{x}^i(t_k)\}_{i=1}^N$ and $\{\boldsymbol{x}^i(s_{k-1})\}_{i=1}^N \mapsto$

[‡]Here, we use the Euler-Maruyama scheme for ease of implementation; it can be replaced by any SDE integrator; see [41, Remark 1 in Sec. III.B].

$\{\boldsymbol{x}^i(s_k)\}_{i=1}^N$, are the Euler-Maruyama updates associated with (3.56) with time step-sizes τ and σ, respectively. Following [41], we write the Sinkhorn regularized proximal recursions for (3.49) in vector form:

$$\hat{\boldsymbol{\phi}}_k = \underset{\hat{\boldsymbol{\phi}}}{arg\,min}\Big\{\min_{\boldsymbol{M}\in\Pi(\hat{\boldsymbol{\phi}}_{k-1},\hat{\boldsymbol{\phi}})}\frac{1}{2}\Big\langle \boldsymbol{C}\left(\{\boldsymbol{x}^i(t_{k-1}),\boldsymbol{x}^i(t_k)\}_{i=1}^N\right),\boldsymbol{M}\Big\rangle$$
$$+\gamma\langle\boldsymbol{M},\log\boldsymbol{M}\rangle+\tau\langle V\left(\{\boldsymbol{x}^i(t_{k-1})\}_{i=1}^N\right)+\epsilon\log\hat{\boldsymbol{\phi}},\hat{\boldsymbol{\phi}}\rangle\Big\}, \tag{3.57a}$$
$$\boldsymbol{\varpi}_k = \underset{\boldsymbol{\varpi}}{arg\,min}\Big\{\min_{\boldsymbol{M}\in\Pi(\boldsymbol{\varpi}_{k-1},\boldsymbol{\varpi})}\frac{1}{2}\Big\langle \boldsymbol{C}\left(\{\boldsymbol{x}^i(s_{k-1}),\boldsymbol{x}^i(s_k)\}_{i=1}^N\right),\boldsymbol{M}\Big\rangle$$
$$+\gamma\langle\boldsymbol{M},\log\boldsymbol{M}\rangle+\sigma\langle V\left(\{\boldsymbol{x}^i(s_{k-1})\}_{i=1}^N\right)+\epsilon\log\boldsymbol{\varpi},\boldsymbol{\varpi}\rangle\Big\}, \tag{3.57b}$$

wherein for $k\in\mathbb{N}$, the $(i,j)^{\text{th}}$ component of the matrix $\boldsymbol{C}\in\mathbb{R}^{N\times N}$ in (3.57a) equals $\|\boldsymbol{x}^i(t_{k-1})-\boldsymbol{x}^j(t_k)\|_2^2$; likewise, the $(i,j)^{\text{th}}$ component of the matrix $\boldsymbol{C}$ in (3.57b) equals $\|\boldsymbol{x}^i(s_{k-1})-\boldsymbol{x}^j(s_k)\|_2^2$. In (3.57), the notation $\Pi(\boldsymbol{a},\boldsymbol{b})$ stands for the set of all matrices $\boldsymbol{M}\in\mathbb{R}^{N\times N}$ such that $\boldsymbol{M}\geq 0$, $\boldsymbol{M}\mathbf{1}=\boldsymbol{a}$, and $\boldsymbol{M}^\top\mathbf{1}=\boldsymbol{b}$, for given admissible[§] vectors $\boldsymbol{a},\boldsymbol{b}\in\mathbb{R}^N$. Furthermore, $V\left(\{\boldsymbol{x}^i(t_{k-1})\}_{i=1}^N\right)$ returns the $N\times 1$ vector whose i^{th} element equals $V\left(\boldsymbol{x}^i(t_{k-1})\right)$; likewise for $V\left(\{\boldsymbol{x}^i(s_{k-1})\}_{i=1}^N\right)$. The term $\langle\boldsymbol{M},\log\boldsymbol{M}\rangle$ is the Sinkhorn/entropic regularization, and $\gamma>0$ is a small regularization parameter. Dualizing (3.57) lead to certain generalized Sinkhorn-type fixed point iterations [73, 130] whose solutions yield the proximal updates $\left(\hat{\boldsymbol{\phi}}_k,\boldsymbol{\varpi}_k\right)$ for $k\in\mathbb{N}$. For algorithmic details, we refer the readers to [41, Sec. III.B]. The convergence guarantees for such iterations follow from the contractive properties of cone preserving nonlinear maps that arise as first order conditions of optimality for (3.57), and are detailed in [41, Sec. III.C].

For the mixed conservative-dissipative drift case, the *uncontrolled* SDE associated

[§]Here, "admissible" means that the vectors $\boldsymbol{a},\boldsymbol{b}$ are component-wise nonnegative, and have equal element-wise sum. This is due to the fact that the generators in (3.33) and (3.40) are both integral preserving and nonnegativity preserving.

with (3.40) is

$$\begin{pmatrix} \mathrm{d}\boldsymbol{\xi} \\ \mathrm{d}\boldsymbol{\eta} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\eta} \\ -\nabla_{\boldsymbol{\xi}} V\left(\boldsymbol{\xi}\right) - \kappa\boldsymbol{\eta} \end{pmatrix} \mathrm{d}t + \sqrt{2\epsilon\kappa} \begin{pmatrix} \boldsymbol{0}_{m\times m} \\ \boldsymbol{I}_{m\times m} \end{pmatrix} \mathrm{d}\boldsymbol{w}. \tag{3.58}$$

As before, for $k \in \mathbb{N}$, we denote the Euler-Maruyama updates associated with (3.58) with time step-sizes τ and σ, respectively, as $\{\boldsymbol{\xi}^i(t_{k-1}), \boldsymbol{\eta}^i(t_{k-1})\}_{i=1}^N \mapsto \{\boldsymbol{\xi}^i(t_k), \boldsymbol{\eta}^i(t_k)\}_{i=1}^N$ and $\{\boldsymbol{\xi}^i(s_{k-1}), \boldsymbol{\vartheta}^i(s_{k-1})\}_{i=1}^N \mapsto \{\boldsymbol{\xi}^i(s_k), \boldsymbol{\vartheta}^i(s_k)\}_{i=1}^N$. To reduce notational overload, let us recall the shorthands we used in (3.51)-(3.52): $\boldsymbol{x} = (\boldsymbol{\xi}, \boldsymbol{\eta})^\top$, $\boldsymbol{y} = (\boldsymbol{\xi}, \boldsymbol{\vartheta})^\top$, and write the Sinkhorn regularized proximal recursions for (3.53) in vector form (see [41, Sec. V.B]):

$$\begin{aligned} \hat{\boldsymbol{\phi}}_k = \ & \underset{\hat{\boldsymbol{\phi}}}{arg\,min} \Bigg\{ \min_{\boldsymbol{M}\in\Pi(\hat{\boldsymbol{\phi}}_{k-1}, \hat{\boldsymbol{\phi}})} \frac{1}{2} \Big\langle \boldsymbol{S}_\tau \left(\{\boldsymbol{x}^i(t_{k-1}), \boldsymbol{x}^i(t_k)\}_{i=1}^N \right), \boldsymbol{M} \Big\rangle \\ & + \gamma \langle \boldsymbol{M}, \log \boldsymbol{M} \rangle + \tau \Big\langle \Big\{ \frac{1}{2} \|\boldsymbol{\eta}^i(s_{k-1})\|_2^2 \Big\}_{i=1}^N + \epsilon \log \boldsymbol{\varpi}, \boldsymbol{\varpi} \Big\rangle \Bigg\}, \end{aligned} \tag{3.59a}$$

$$\begin{aligned} \boldsymbol{\varpi}_k = \ & \underset{\boldsymbol{\varpi}}{arg\,min} \Bigg\{ \min_{\boldsymbol{M}\in\Pi(\boldsymbol{\varpi}_{k-1}, \boldsymbol{\varpi})} \frac{1}{2} \Big\langle \boldsymbol{S}_\sigma \left(\{\boldsymbol{y}^i(s_{k-1}), \boldsymbol{y}^i(s_k)\}_{i=1}^N \right), \boldsymbol{M} \Big\rangle \\ & + \gamma \langle \boldsymbol{M}, \log \boldsymbol{M} \rangle + \sigma \Big\langle \Big\{ \frac{1}{2} \|\boldsymbol{\vartheta}^i(s_{k-1})\|_2^2 \Big\}_{i=1}^N + \epsilon \log \boldsymbol{\varpi}, \boldsymbol{\varpi} \Big\rangle \Bigg\}, \end{aligned} \tag{3.59b}$$

wherein for $k \in \mathbb{N}$, the $(i,j)^{\text{th}}$ component of the matrix $\boldsymbol{S}_\tau \in \mathbb{R}^{N\times N}$ in (3.59a) equals $s_\tau\left(\boldsymbol{x}^i(t_{k-1}), \boldsymbol{x}^j(t_k)\right)$, and $s_\tau(\cdot,\cdot)$ is given by (3.52a). Similarly, the $(i,j)^{\text{th}}$ component of the matrix $\boldsymbol{S}_\sigma \in \mathbb{R}^{N\times N}$ in (3.59b) equals $s_\sigma\left(\boldsymbol{y}^i(s_{k-1}), \boldsymbol{y}^j(s_k)\right)$, and $s_\sigma(\cdot,\cdot)$ is given by (3.52b).

3.5.4 Overall Algorithm

We now bring together the ideas from Sections 3.4 and 3.5.3, and outline the overall algorithm to solve the SBPs with gradient and mixed-conservative dissipative drifts via scattered point cloud-based computation. We perform a recursion over

the pair $(\hat{\varphi}_0, \varphi_1)$, and the computational steps for the same are:

Step 1. Guess $\hat{\varphi}_1(\boldsymbol{x})$ (everywhere nonnegative).

Step 2. Compute $\varphi_1(\boldsymbol{x}) = \rho_1(\boldsymbol{x})/\hat{\varphi}_1(\boldsymbol{x})$.

Step 3. For the case of gradient drift, compute

$$p(\boldsymbol{x}, s=0) = \varphi_1(\boldsymbol{x}) \exp\left(-V(\boldsymbol{x})/\epsilon\right).$$

For the case of mixed conservative-dissipative drift, compute

$$p(\boldsymbol{\xi}, \boldsymbol{\vartheta}, s=0) = \varphi_1(\boldsymbol{\xi}, -\boldsymbol{\vartheta}) \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\vartheta}\|_2^2 + V(\boldsymbol{\xi})\right)\right).$$

Step 4. Solve IVP (3.33b) or (3.40b) till $s = 1$ to obtain $p(\boldsymbol{x}, s = 1)$.

Step 5. For the case of gradient drift, compute

$$\varphi_0(\boldsymbol{x}) = p(\boldsymbol{x}, s=1) \exp\left(V(\boldsymbol{x})/\epsilon\right).$$

For the case of mixed conservative-dissipative drift, compute

$$\varphi_0(\boldsymbol{\xi}, \boldsymbol{\eta}) = p(\boldsymbol{\xi}, \boldsymbol{\vartheta}, s=1) \exp\left(\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right).$$

Step 6. Compute $\hat{\varphi}_0(\boldsymbol{x}) = \rho_0(\boldsymbol{x})/\varphi_0(\boldsymbol{x})$.

Step 7. Solve IVP (3.33a) or (3.40a) till $t = 1$ to obtain $\hat{\varphi}_1(\boldsymbol{x})$.

Step 8. Repeat until the pair $(\hat{\varphi}_0, \varphi_1)$ has converged[¶].

[¶]We check whether the Wasserstein distances between the current and previous iterates for the pair $(\hat{\varphi}_0, p(\boldsymbol{x}, s=0))$, which is a proxy for the pair $(\hat{\varphi}_0, \varphi_1)$, are below some user-specified numerical tolerance. These Wasserstein distances are computed by solving linear programs, i.e., discrete versions of (41).

That such a recursion will converge, follows from the fact [53, Proposition 1 in Sec. III] that the recursion is in fact contractive in Hilbert's projective metric [36, 113] provided (i) the endpoint PDFs have compact supports, and (ii) the transition probability densities for (3.33) and (3.40) are continuous and positive. Under the stated assumptions on the function V, the transition densities indeed satisfy these properties; see Appendix A.7. Since our computational framework involves finite set of scattered points, the compactness condition also holds. Once the *endpoint* Schrödinger factors are found as outlined in Steps 1–8 above, we then compute the Schrödinger factors at any time t, i.e., the pair $(\hat{\varphi}(\boldsymbol{x}, t), \varphi(\boldsymbol{x}, t))$ using the IVPs (3.33) or (3.40).

For scattered point cloud-based computation with sample size N, the aforesaid steps lead to a recursion over the pair of vectors $(\hat{\boldsymbol{\phi}}_0, \boldsymbol{\phi}_1)$, each of these vectors being of size $N \times 1$. We refer to this proposed procedure as ComputeFactorsSBP which takes the $N \times 1$ endpoint PDF vectors $\boldsymbol{\rho}_0, \boldsymbol{\rho}_1$ (i.e., the endpoint joint PDF values evaluated at the scattered sample state vectors) as input, and returns the converged pair $(\hat{\boldsymbol{\phi}}_0, \boldsymbol{\phi}_1)$. As detailed in Algorithm 1, ComputeFactorsSBP performs the discrete versions of the Steps 1–8 above, followed by computation of the transient Schrödinger factors. For conceptual clarity, in Algorithm 1, we list the steps for the case of gradient drift; adapting Algorithm 1 to the case of mixed conservative-dissipative drift is straightforward following the modifications in Steps 3, 4, 5 and 7 mentioned before.

The procedure ComputeFactorsSBP involves a main iteration where we guess an initial vector $\hat{\boldsymbol{\phi}}_1$, and repeatedly update the vectors $\hat{\boldsymbol{\phi}}_0, \hat{\boldsymbol{\phi}}_1, \boldsymbol{\phi}_0, \boldsymbol{\phi}_1$ using a secondary function called ProxRecur that performs the proximal recursions (3.57)

or (3.59) for solving the IVPs (3.33) or (3.40) in Steps 4 and 7 above. The main iteration requires the parameters $\gamma, \epsilon, \tau, \sigma, N$, the function $V(\cdot)$ described before, and numSteps which is the number of time-steps for the proximal update. Furthermore, we use the two pairs of parameters $(\mathrm{tol}_{\mathrm{SB}}, \mathrm{maxIter}_{\mathrm{SB}}), (\mathrm{tol}_{\mathrm{PR}}, \mathrm{maxIter}_{\mathrm{PR}})$: the first pair specifies the numerical tolerance and the maximum number of iterations for the main iteration in ComputeFactorsSBP; the second pair specifies the same for the Sinkhorn algorithm used in ProxRecur. Notice that the procedure ProxRecur subsumes the Euler-Maruyama scheme to update the states $\{\boldsymbol{x}^i\}_{i=1}^N$; see [41, Sec. III.B].

Once the converged vector pair $(\hat{\boldsymbol{\phi}}_0, \boldsymbol{\phi}_1)$ is obtained, we again invoke ProxRecur to solve the IVP pair (3.33) or (3.40) via (3.57) or (3.59), and compute the discrete transient solutions $(\hat{\boldsymbol{\phi}}_k, \boldsymbol{\varpi}_k)$. We remind the readers that since the function p is integrated in the time coordinate s, its discrete version $\boldsymbol{\varpi}_k$ is really evaluated at $t = 1 - k\tau$, i.e., $\boldsymbol{\varpi}_k \approx p(\boldsymbol{x}, t = 1 - k\tau)$ for $k \in \mathbb{N}$. We next employ the discrete versions of the mappings $p \mapsto \varphi$, given by

$$\begin{aligned} \varphi(\boldsymbol{x}, t) &= p(\boldsymbol{x}, s) \exp\left(V(\boldsymbol{x})/\epsilon\right), \\ \varphi(\boldsymbol{\xi}, \boldsymbol{\eta}, t) &= p(\boldsymbol{\xi}, \boldsymbol{\vartheta}, s) \exp\left(\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right), \end{aligned}$$

to compute $\boldsymbol{\varpi}_k \mapsto \boldsymbol{\phi}_k$, i.e., $\left(\hat{\boldsymbol{\phi}}_k, \boldsymbol{\varpi}_k\right) \mapsto \left(\hat{\boldsymbol{\phi}}_k, \boldsymbol{\phi}_k\right)$. This completes the computation of the transient Schrödinger factors $\left(\hat{\boldsymbol{\phi}}_k, \boldsymbol{\phi}_k\right)$ as a pair of weighted scattered point clouds, i.e., as function values $(\hat{\varphi}(\boldsymbol{x}, t), \varphi(\boldsymbol{x}, t))$ evaluated at the state coordinates that are updated by the Euler-Maruyama scheme.

To compute the discrete optimal pair $(\boldsymbol{\rho}_k^{\mathrm{opt}}, \boldsymbol{u}_k^{\rho^{\mathrm{opt}}})$ from $\left(\hat{\boldsymbol{\phi}}_k, \boldsymbol{\phi}_k\right)$, we need to eval-

uate the Hadamard product

$$\boldsymbol{\rho}_k^{\text{opt}} = \boldsymbol{\phi}_{\text{numSteps}+1-k} \odot \hat{\boldsymbol{\phi}}_k, \tag{3.60}$$

which in our case, is not well-defined as is, since $\boldsymbol{\phi}^{\text{trans}}_{\text{numSteps}+1-k}$ and $\hat{\boldsymbol{\phi}}_k^{\text{trans}}$ are evaluated on different state coordinates (i.e., are supported on different finite sets). On the other hand, the optimal feedback control

$$\boldsymbol{u}_k^{\text{opt}} = 2\epsilon \boldsymbol{B}^\top \nabla \log \boldsymbol{\phi}_k, \tag{3.61}$$

requires computing a gradient w.r.t. the state, which again is not well-defined for scattered data. To circumvent these issues, we perform scattered data interpolation for the transient Schrödinger factors $\left(\hat{\boldsymbol{\phi}}_k, \boldsymbol{\phi}_k\right)$ via *multiquadric* scheme [107,108] that uses radial basis functions to fit a surface to the weighted scattered point cloud. The interpolation requires the vectors $(\boldsymbol{\phi}_k, \hat{\boldsymbol{\phi}}_k)$, their corresponding state space coordinates, a set of user-provided query points, and returns the interpolated values evaluated at the query points. Using the same set of query points to interpolate the two point clouds for the pair $(\boldsymbol{\phi}_k, \hat{\boldsymbol{\phi}}_k)$, we perform the element-wise multiplication (3.60). As for (3.61), we use standard finite difference techniques to compute the gradient of the logarithm of interpolated values of $\boldsymbol{\phi}_k$.

3.6 Numerical Examples

3.6.1 SBP Example with Gradient Drift

We consider an instance of the SBP given in Section 3.4.1 with $\boldsymbol{x} \in \mathbb{R}^2$, and $V(x_1, x_2) = \frac{1}{4}(1 + x_1^4) + \frac{1}{2}(x_2^2 - x_1^2)$. The controlled prior dynamics is

$$\begin{pmatrix} \mathrm{d}x_1 \\ \mathrm{d}x_2 \end{pmatrix} = -\nabla V(x_1, x_2)\,\mathrm{d}t + \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}\,\mathrm{d}t + \sqrt{2\epsilon}\begin{pmatrix} \mathrm{d}w_1 \\ \mathrm{d}w_2 \end{pmatrix}. \tag{3.62}$$

The control objective is to steer the prescribed joint PDF of the initial condition $\boldsymbol{x}(t=0) \sim \rho_0 = \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ to the prescribed joint PDF of the terminal condition $\boldsymbol{x}(t=1) \sim \rho_1 = c_1\mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) + c_2\mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ over $t \in [0, 1]$, subject to (3.62), while minimizing the control effort (3.27a). Here, we fix

$$\boldsymbol{\mu}_0 = (-2, 0)^\top,\ \boldsymbol{\Sigma}_0 = \mathrm{diag}(0.8, 0.7),\ c_1 = c_2 = 0.5,\ \boldsymbol{\mu}_1 = (1.5, 2)^\top,$$
$$\boldsymbol{\mu}_2 = (1.5, -2)^\top,\ \boldsymbol{\Sigma}_1 = \mathrm{diag}(0.5, 0.8),\ \boldsymbol{\Sigma}_2 = \mathrm{diag}(0.7, 0.8).$$

Notice that in the absence of control ($\boldsymbol{u} \equiv 0$), the transient joint PDFs of (3.62) tend to the stationary solution $\rho_\infty \propto \exp\left(-V(x_1, x_2)/\epsilon\right)$ which has two modes along the horizontal axis; see [41, Fig. 8]. In contrast, the prescribed terminal bimodal PDF ρ_1, specified as a mixture of Gaussians, has two modes along the vertical axis.

Fig. 3.3 shows the evolution of the optimal controlled joint PDF $\rho^{\mathrm{opt}}(\boldsymbol{x}, t)$ obtained by solving the SBP using Algorithm 1 given in Section 3.5.4. For numerical simulation, we set $\epsilon = 6$, $\gamma = 0.5$, $\tau = \sigma = 10^{-3}$, $N = 500$, numSteps $= 1000$, $\mathrm{tol}_{\mathrm{SB}} = 0.1$, $\mathrm{maxIter}_{\mathrm{SB}} = \mathrm{maxIter}_{\mathrm{PR}} = 500$, and $\mathrm{tol}_{\mathrm{PR}} = 10^{-3}$. In Fig. 3.3,

each subplot corresponds to a different snapshot in time; each subplot is plotted over the domain $[-4,4]\times[-6,6]$. For the purpose of comparison, Fig. 3.4 shows the uncontrolled joint state PDF evolution starting from the same initial PDF ρ_0. The components of the optimal feedback control $\boldsymbol{u}^{\mathrm{opt}}(\boldsymbol{x},t)$ obtained from the SBP solution, are shown in Figs. 3.5 and 3.6. Fig. 3.7 depicts the magnitude of $\boldsymbol{u}^{\rho^{\mathrm{opt}}}$.

3.6.2 SBP Example with Mixed Conservative-Dissipative Drift

We next consider an instance of the SBP given in Section 3.4.2 with $n=2m=2$, i.e., $\boldsymbol{x}=(\xi,\eta)^\top\in\mathbb{R}^2$, and $V(\xi)=5\xi^4$. In other words, the controlled prior dynamics is

$$\begin{pmatrix}\mathrm{d}\xi\\ \mathrm{d}\eta\end{pmatrix}=\begin{pmatrix}\eta\\ -\frac{\partial}{\partial\xi}5\xi^4-\kappa\eta+u\end{pmatrix}\mathrm{d}t+\sqrt{2\epsilon\kappa}\begin{pmatrix}0\\ \mathrm{d}w\end{pmatrix}. \tag{3.63}$$

Notice that the function V satisfies the conditions mentioned in Section 3.4.2. We use the same endpoint PDFs ρ_0,ρ_1 as in Section 3.6.1, and solve the SBP using Algorithm 1 with the Steps 3, 4, 5 and 7 modified for the mixed conservative-dissipative case, as mentioned before. Notice that in the absence of control ($u\equiv 0$), the transient joint PDFs of (3.63) tend to the stationary solution $\rho_\infty\propto\exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\eta^2+V(\xi)\right)\right)$.

Fig. 3.8 shows the evolution of the optimal controlled joint state PDF $\rho^{\mathrm{opt}}(\boldsymbol{x},t)$ obtained from the SBP solution, where each subplot is plotted over the domain $[-4,4]\times[-10,10]$. The parameters used in our numerical simulation are: $\epsilon=5$, $\kappa=0.5$, $\gamma=0.5$, $\tau=\sigma=10^{-3}$, $N=100$, numSteps $=1000$, $\mathrm{tol}_{\mathrm{SB}}=0.1$,

$\text{maxIter}_{\text{SB}} = \text{maxIter}_{\text{PR}} = 500$, and $\text{tol}_{\text{PR}} = 10^{-3}$. The optimal feedback control $\boldsymbol{u}^{\text{opt}}(\boldsymbol{x}, t)$ obtained from our SBP solution is shown in Fig. 3.9.

3.7 Conclusions

We address the problem of minimum energy finite horizon steering of state density between two prescribed endpoint densities via feedback control, subject to trajectory-level dynamics with nonlinear drift. This is a generalization of the classical SBP formulated by Schrödinger in the 1930s, and has many potential applications such as shaping of biological and robotic swarms through feedback synthesis. We derive optimality conditions for the case of a generic nonlinear drift, and show that two specific cases, viz. gradient and mixed-conservative dissipative drifts, are particularly amenable for computational purpose, based on certain infinite-dimensional proximal recursions that exploit the structure of these nonlinearities. These ideas make contact with the theory of Wasserstein gradient flow, and optimal mass transport. Building on our previous work [41], we design proximal algorithms for solving the density steering problems in these two cases of practical interest. Numerical examples are provided to illustrate the proposed framework. Our main contribution is to make algorithmic advances in solving the variants of SBPs that are of interest to the systems-control community. From a probabilistic perspective, while there exists a substantial body of work [12, 104, 115, 170] in steering the second order state statistics (also known as "covariance control"), steering the joint state PDF subject to controlled nonlinear dynamics is a relatively new direction of research. The results of this article is expected to further the theoretical and algorithmic development for the same.

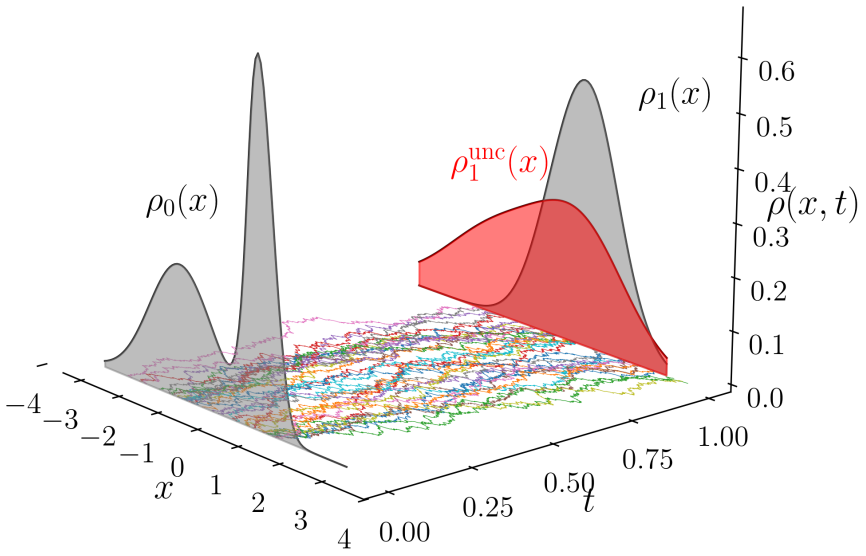

Figure 3.1: The classical SBP with 1-dimensional state space shown above, with state variable $x \in \mathbb{R}$, concerns determining the optimal control $u^{\mathrm{opt}}(x,t)$ that steers the prescribed initial state PDF $\rho_0(x)$ at time $t = 0$ to the prescribed terminal state PDF $\rho_1(x)$ at time $t = 1$, while minimizing $\mathbb{E}\left\{\int_0^1 \frac{1}{2}|u(x,t)|^2 \,\mathrm{d}t\right\}$ subject to a controlled diffusion (3.1b), i.e., the SBP solves (3.1) wherein the expectation operator in the objective is w.r.t. the controlled state PDF $\rho(x,t)$. The optimal controlled sample paths $x^{\mathrm{opt}}(t)$ for 100 randomly chosen initial states are shown in the (x,t) plane. In the absence of control, starting from $\rho_0(x)$, the (uncontrolled) state PDF at $t = 1$ becomes $\rho_1^{\mathrm{unc}}(x)$, as depicted. For numerically solving the classical SBP with $\epsilon = 0.5$ and ρ_0, ρ_1 as shown above, we used the fixed point recursion proposed in [53]; see Section 3.2.1 for details.

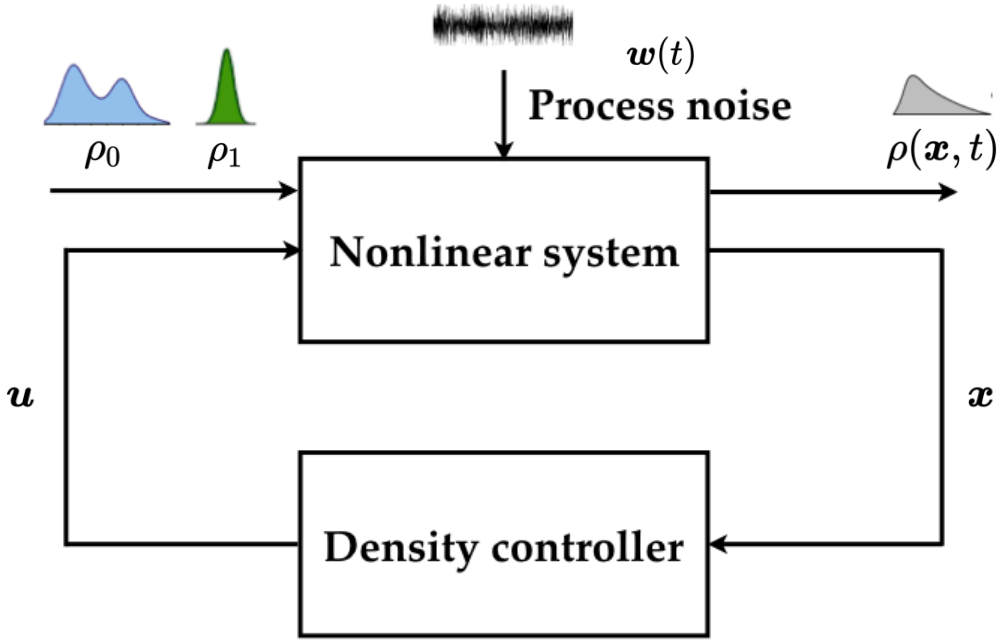

Figure 3.2: Schematic of the SBP with nonlinear prior dynamics (see Section 3.3 for the problem formulation).

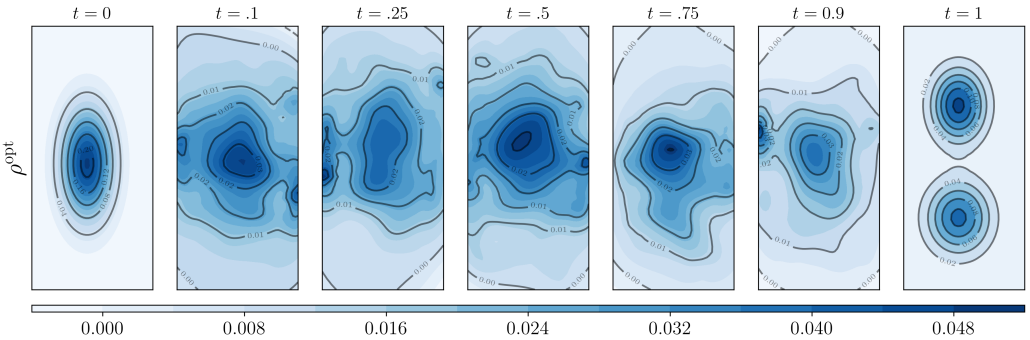

Figure 3.3: For the SBP in Section 3.6.1, shown here are the contour plots of the optimal controlled transient joint state PDFs $\rho^{\text{opt}}(\boldsymbol{x}, t)$, $t \in [0, 1]$, along with the endpoint joint PDFs $\rho_0(\boldsymbol{x}), \rho_1(\boldsymbol{x})$. Each subplot corresponds to a different snapshot in time; all subplots are plotted on the domain $[-4, 4] \times [-6, 6]$. The color denotes the joint PDF value; see colorbar (dark hue = high, light hue = low).

Algorithm 2 ComputeFactorsSBP using ProxRecur to compute the Schrödinger factors $(\hat{\boldsymbol{\phi}}_k^{\text{transient}}, \boldsymbol{\phi}_k^{\text{transient}})$ for gradient drift

procedure ComputeFactorsSBP($\boldsymbol{\rho}_0$, $\boldsymbol{\rho}_1$, γ, ϵ, τ, σ, N, $V(\cdot)$, numSteps, tol_{SB}, $\text{maxIter}_{\text{SB}}$, tol_{PR}, $\text{maxIter}_{\text{PR}}$)
  $\hat{\boldsymbol{\phi}}_1 \leftarrow \left[\text{rand}_{N\times 1}, \mathbf{0}_{N\times(\text{maxIter}_{\text{SB}}-1)}\right]$ ▷ Step 1
  $\boldsymbol{\phi}_0 \leftarrow \left[\mathbf{0}_{N\times\text{maxIter}_{\text{SB}}}\right]$ ▷ initialize
  $\hat{\boldsymbol{\phi}}_0 \leftarrow \left[\mathbf{0}_{N\times\text{maxIter}_{\text{SB}}}\right]$
  $\boldsymbol{\phi}_1 \leftarrow \left[\mathbf{0}_{N\times\text{maxIter}_{\text{SB}}}\right]$
  $\boldsymbol{p}_0 \leftarrow \left[\mathbf{0}_{N\times\text{maxIter}_{\text{SB}}}\right]$
  $\boldsymbol{p}_1 \leftarrow \left[\mathbf{0}_{N\times\text{maxIter}_{\text{SB}}}\right]$
  $\boldsymbol{p}^{\text{temp}} \leftarrow \left[\mathbf{0}_{N\times\text{numSteps}+1}\right]$
  $\hat{\boldsymbol{\phi}}^{\text{temp}} \leftarrow \left[\mathbf{0}_{N\times\text{numSteps}+1}\right]$
  $\ell = 1$ ▷ iteration index for ComputeFactorsSBP
  while $\ell \leq \text{maxIter}_{\text{SB}}$ **do**
    $\boldsymbol{\phi}_1(:,\ell+1) = \boldsymbol{\rho}_1 \oslash \hat{\boldsymbol{\phi}}_1(:,\ell)$ ▷ Step 2
    $\boldsymbol{p}_0(:,\ell+1) = \boldsymbol{\phi}_1(:,\ell+1) \odot \exp(-V(\{\boldsymbol{x}^i\}_{i=1}^N)/\epsilon)$ ▷ Step 3
    $\boldsymbol{p}^{\text{temp}}(:,1) = \boldsymbol{p}_0(:,\ell+1)$
    for $i \leftarrow 1$ to numSteps **do**
      $\boldsymbol{p}^{\text{temp}}(:,i+1) \leftarrow$ ProxRecur($\boldsymbol{p}^{\text{temp}}(:,i), \gamma, \epsilon, \sigma, N, V(\cdot), \text{tol}_{\text{PR}}, \text{maxIter}_{\text{PR}}$)
    end for
    $\boldsymbol{p}_1(:,\ell+1) \leftarrow \boldsymbol{p}^{\text{temp}}(:,\text{numSteps}+1)$ ▷ Step 4
    $\boldsymbol{\phi}_0(:,\ell+1) \leftarrow \boldsymbol{p}_1(:,\ell+1) \odot \exp(V\left(\{\boldsymbol{x}^i\}_{i=1}^N\right)/\epsilon)$ ▷ Step 5
    $\hat{\boldsymbol{\phi}}_0(:,\ell+1) \leftarrow \boldsymbol{\rho}_0 \oslash \boldsymbol{\phi}_0(:,\ell+1)$ ▷ Step 6
    $\hat{\boldsymbol{\phi}}_{\text{temp}}(:,1) \leftarrow \hat{\boldsymbol{\phi}}_0(:,\ell+1)$
    for $j \leftarrow 1$ to numSteps **do**
      $\hat{\boldsymbol{\phi}}^{\text{temp}}(:,j+1) \leftarrow$ ProxRecur($\hat{\boldsymbol{\phi}}^{\text{temp}}(:,j), \gamma, \epsilon, \tau, N, V(\cdot), \text{tol}_{\text{PR}}, \text{maxIter}_{\text{PR}}$)
    end for
    $\hat{\boldsymbol{\phi}}_1(:,\ell+1) \leftarrow \hat{\boldsymbol{\phi}}^{\text{temp}}(:,\text{numSteps}+1)$ ▷ Step 7
    if $W^2(\boldsymbol{p}_0(:,\ell+1), \boldsymbol{p}_0(:,\ell)) < \text{tol}_{\text{SB}}$ & $W^2(\hat{\boldsymbol{\phi}}_0(:,\ell+1), \hat{\boldsymbol{\phi}}_0(:,\ell)) < \text{tol}_{\text{SB}}$ **then**
      break
    else
      $\ell \leftarrow \ell + 1$ ▷ Step 8
    end if
  end while
  return $\hat{\boldsymbol{\phi}}_0(:,\ell), \boldsymbol{\phi}_1(:,\ell)$ ▷ converged endpoint pair
  $\hat{\boldsymbol{\phi}}_{k=1}^{\text{transient}} \leftarrow \hat{\boldsymbol{\phi}}_0(:,\ell)$ ▷ initialize
  $\boldsymbol{p}_{k=1}^{\text{transient}} \leftarrow \boldsymbol{\phi}_1(:,\ell) \odot \exp(-V\left(\{\boldsymbol{x}^i\}_{i=1}^N\right)_1/\epsilon)$ ▷ initialize
  for $k \leftarrow 1$ to numSteps **do**
    $\hat{\boldsymbol{\phi}}_{k+1}^{\text{transient}} \leftarrow$ ProxRecur($\hat{\boldsymbol{\phi}}_k^{\text{transient}}, \gamma, \epsilon, \tau, N, V(\cdot), \text{tol}_{\text{PR}}, \text{maxIter}_{\text{PR}}$)
    $\boldsymbol{p}_{k+1}^{\text{transient}} \leftarrow$ ProxRecur($\boldsymbol{p}_k^{\text{transient}}, \gamma, \epsilon, \sigma, N, V(\cdot), \text{tol}_{\text{PR}}, \text{maxIter}_{\text{PR}}$)
    $\boldsymbol{\phi}_{k+1}^{\text{transient}} \leftarrow \boldsymbol{p}_{k+1}^{\text{transient}} \odot \exp(V\left(\{\boldsymbol{x}^i\}_{i=1}^N\right)/\epsilon)$
  end for
  return $(\hat{\boldsymbol{\phi}}_k^{\text{transient}}, \boldsymbol{\phi}_k^{\text{transient}})$ ▷ transient Schrödinger factors
end procedure

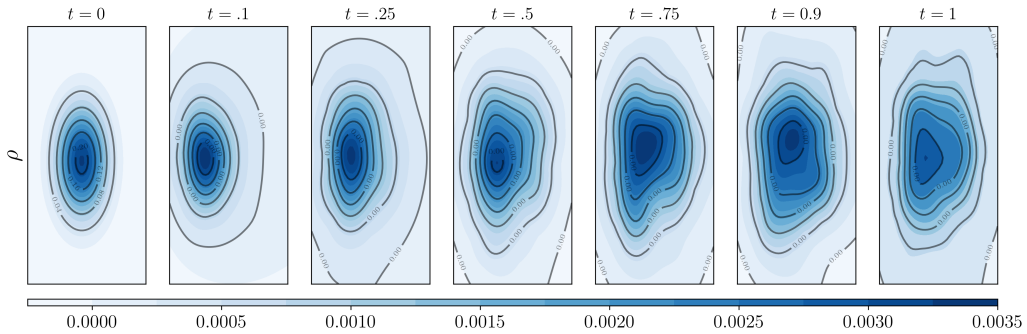

Figure 3.4: The contour plots of the uncontrolled ($\boldsymbol{u} \equiv 0$) transient joint state PDFs $\rho(\boldsymbol{x}, t)$, $t \in [0, 1]$, for (3.62) starting from the initial joint state PDF ρ_0 given in Section 3.6.1. Each subplot corresponds to a different snapshot in time. Each subplot corresponds to a different snapshot in time; all subplots are plotted on the domain $[-4, 4] \times [-6, 6]$. The color denotes the joint PDF value; see colorbar (dark hue = high, light hue = low).

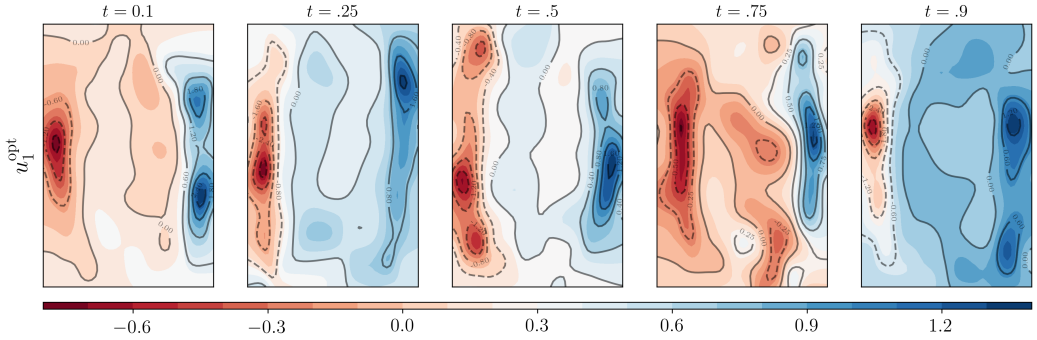

Figure 3.5: For the SBP in Section 3.6.1, shown here are the contour plots of $u_1^{\rho^{\text{opt}}}(\boldsymbol{x}, t)$, the first component of the optimal feedback control. Each subplot is plotted on the domain $[-4, 4] \times [-6, 6]$. The color (blue = high, red = low) denotes the value of $u_1^{\rho^{\text{opt}}}$ at each snapshot in time.

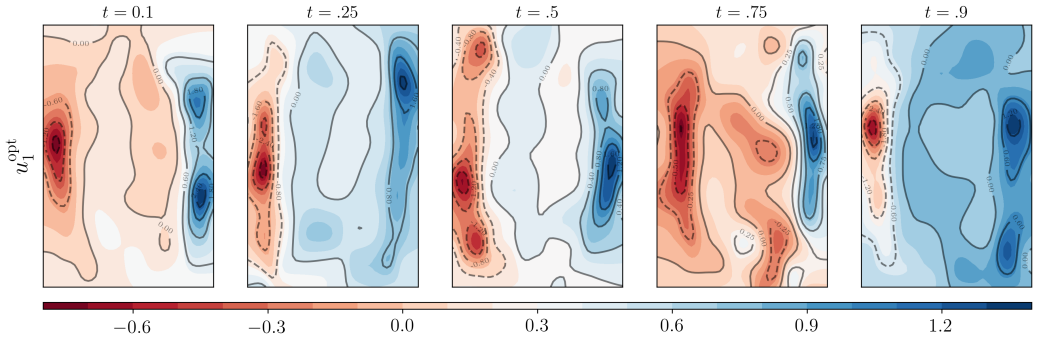

Figure 3.6: For the SBP in Section 3.6.1, shown here are the contour plots of $u_2^{\rho^{\text{opt}}}(\boldsymbol{x}, t)$, the second component of the optimal feedback control. Each subplot is plotted on the domain $[-4, 4] \times [-6, 6]$. The color (blue = high, red = low) denotes the value of $u_1^{\rho^{\text{opt}}}$ at each snapshot in time.

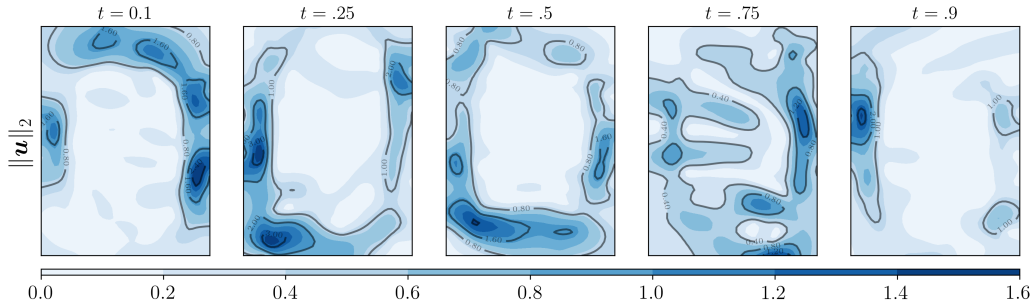

Figure 3.7: For the SBP in Section 3.6.1, shown here are the contour plots for the magnitude (dark hue = high, light hue = low; see colorbar) of the optimal feedback control, i.e., $\|\boldsymbol{u}\|_2$.

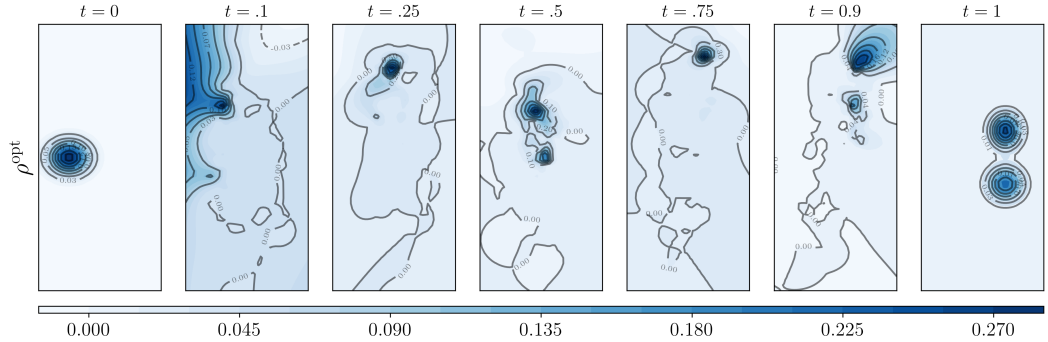

Figure 3.8: For the SBP in Section 3.6.2, shown here are the contour plots of the optimal controlled transient joint state PDFs $\rho^{\text{opt}}(\boldsymbol{x}, t)$, $t \in [0, 1]$, along with the endpoint joint PDFs $\rho_0(\boldsymbol{x}), \rho_1(\boldsymbol{x})$. Each subplot corresponds to a different snapshot in time; all subplots are plotted on the domain $[-4, 4] \times [-10, 10]$. The color denotes the joint PDF value; see colorbar (dark hue = high, light hue = low).

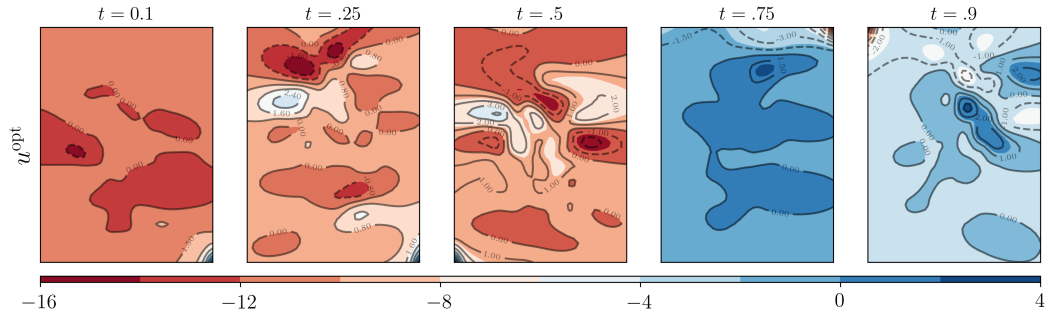

Figure 3.9: For the SBP in Section 3.6.2, shown here are the contour plots of the optimal feedback control $u^{\rho^{\text{opt}}}(\boldsymbol{x}, t)$. Each subplot is plotted on the domain $[-4, 4] \times [-10, 10]$. The color (blue = high, red = low) denotes the value of $u^{\rho^{\text{opt}}}$ at each snapshot in time; see colorbar.

Chapter 4

Reflected Schrödinger Bridge: Density Control with Path Constraints

How to steer a given joint state probability density function to another over finite horizon subject to a controlled stochastic dynamics with hard state (sample path) constraints? In applications, state constraints may encode safety requirements such as obstacle avoidance. In this paper, we perform the feedback synthesis for minimum control effort density steering (a.k.a. Schrödinger bridge) problem subject to state constraints. We extend the theory of Schrödinger bridges to account the reflecting boundary conditions for the sample paths, and provide a computational framework building on our previous work on proximal recursions, to solve the same.

4.1 Introduction

We consider finite horizon feedback steering of an ensemble of trajectories subject to a controlled stochastic differential equation (SDE) with endpoint joint state probability density function (PDF) constraints – a topic of growing interest in the systems-control literature. Motivating applications include belief space motion planning for vehicular autonomy, and the steering of robotic or biological swarms via decentralized feedback. While early contributions focused on the covariance control [92,115,195], more recent papers [56,58,60] addressed the optimal feedback synthesis for steering an arbitrary prescribed initial joint state PDF to another prescribed terminal joint state PDF subject to controlled linear dynamics, and revealed the connections between the associated stochastic optimal control problem, the theory of optimal mass transport [210], and the Schrödinger bridge [191, 192]. Follow up works have accounted terminal cost [104], input constraints [12, 171], output feedback [11], and some nonlinear dynamics [38, 40, 44]. As for the state or path constraints, prior work [170] incorporated the same in soft probabilistic sense. The contribution of the present paper is to account *hard deterministic path constraints* in the problem of minimum effort finite horizon PDF steering via feedback synthesis. This can be intuitively phrased as the "hard safety with soft endpoint" problem.

The main idea underlying the ensuing development is to modify the unconstrained Itô SDEs to the "reflected Itô SDEs" [110,121,151,217], i.e., the controlled sample paths in the state space (in addition to the control-affine deterministic drift) are driven by two stochastic processes: a Wiener process, and a local time stochastic process. The latter enforces the sample paths in the state space to satisfy the

deterministic non-strict* path containment constraints at all times. These considerations engender a Schrödinger bridge-like formulation–referred hereafter as the *Reflected Schrödinger Bridge Problem (RSBP)*–which unlike its classical counterpart, has extra boundary conditions involving the gradients of the so-called Schrödinger factors. We show how recent developments in contraction mapping w.r.t. the Hilbert metric, and the proximal recursion over the Schrödinger factors can be harnessed to solve the RSBP.

4.2 Reflected Schrödinger Bridge Problem

4.2.1 Formulation

Consider a connected, smooth† and bounded domain $\mathcal{X} \subset \mathbb{R}^n$. Let $\overline{\mathcal{X}} := \mathcal{X} \cup \partial\mathcal{X}$ denote the closure of $\mathcal{X}$. For time $t \in [0,1]$, consider the stochastic control problem

$$\inf_{\boldsymbol{u}\in U} \quad \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|\boldsymbol{u}(t,\boldsymbol{x}_t^{\boldsymbol{u}})\|_2^2 \,\mathrm{d}t\right\} \tag{4.1a}$$

$$\text{subject to} \quad \mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}} = \boldsymbol{f}(t,\boldsymbol{x}_t^{\boldsymbol{u}})\,\mathrm{d}t + \boldsymbol{u}(t,\boldsymbol{x}_t^{\boldsymbol{u}})\,\mathrm{d}t + \sqrt{2\theta}\,\mathrm{d}\boldsymbol{w}_t + \boldsymbol{n}(\boldsymbol{x}_t^{\boldsymbol{u}})\mathrm{d}\gamma_t, \tag{4.1b}$$

$$\boldsymbol{x}_0^{\boldsymbol{u}} := \boldsymbol{x}_t^u\,(t=0) \sim \rho_0, \quad \boldsymbol{x}_1^{\boldsymbol{u}} := \boldsymbol{x}_t^u\,(t=1) \sim \rho_1, \tag{4.1c}$$

*There is no loss of generality in allowing the sample paths to satisfy *non-strict* path containment in given $\mathcal{X} \subset \mathbb{R}^n$ since *strict* containment can be enforced by reflecting them from ϵ-inner boundary layer of $\partial\mathcal{X}$ for ϵ small enough.

†More precisely, there exists $\xi \in C_b^2(\mathbb{R}^n)$ such that $\mathcal{X} \equiv \{\boldsymbol{x} \in \mathbb{R}^n \mid \xi(\boldsymbol{x}) > 0\}$ with boundary $\partial\mathcal{X} \equiv \{\boldsymbol{x} \in \mathbb{R}^n \mid \xi(\boldsymbol{x}) = 0\}$.

where $\boldsymbol{w}_t$ is the standard Wiener process in $\mathbb{R}^n$, the controlled state $\boldsymbol{x}_t^{\boldsymbol{u}} \in \overline{\mathcal{X}}$, and the endpoint joint state PDFs ρ_0, ρ_1 are prescribed[‡] such that their supports are in $\overline{\mathcal{X}}$, both are everywhere nonnegative, have finite second moments, and $\int \rho_0 = \int \rho_1 = 1$. The parameter $\theta > 0$ is referred to as the thermodynamic temperature, and the expectation operator $\mathbb{E}\{\cdot\}$ in (4.1a) is w.r.t. the law of the controlled state $\boldsymbol{x}_t^{\boldsymbol{u}}$. The set $\mathcal{U}$ consists of all admissible feedback policies $\boldsymbol{u}(t, \boldsymbol{x}_t^u)$, given by $\mathcal{U} := \{\boldsymbol{u} : [0,1] \times \overline{\mathcal{X}} \mapsto \mathbb{R}^n \mid \|\boldsymbol{u}\|_2^2 < \infty, \boldsymbol{u}(t,\cdot) \in \text{Lipschitz}\left(\overline{\mathcal{X}}\right) \text{ for all } t \in [0,1]\}$. We assume that the prior drift vector field $\boldsymbol{f}$ is bounded Borel measurable in $(t, \boldsymbol{x}_t^{\boldsymbol{u}}) \in [0,1] \times \overline{\mathcal{X}}$, and Lipschitz continuous w.r.t. $\boldsymbol{x}_t^{\boldsymbol{u}} \in \overline{\mathcal{X}}$. The vector field $\boldsymbol{n}$ is set to be the inward unit normal to the boundary $\partial\mathcal{X}$, and gives the direction of reflection. Furthermore, for $t \in [0,1]$, γ_t is *minimal local time*: a continuous, non-negative and non-decreasing stochastic process [89, 109, 184] that restricts $\boldsymbol{x}_t^u$ to the domain $\mathcal{X}$, with $\gamma_0 \equiv 0$. Specifically, letting $\mathbb{1}_{\{\}}$ denote the indicator function of the subscripted set, we have

$$\gamma_t = \int_0^t \mathbb{1}_{\{\boldsymbol{x}_s^{\boldsymbol{u}} \in \partial\mathcal{X}\}} \, \mathrm{d}\gamma_s, \quad \int_0^1 \mathbb{1}_{\{\boldsymbol{x}_t^{\boldsymbol{u}} \notin \partial\mathcal{X}\}} \, \mathrm{d}\gamma_t = 0, \tag{4.2}$$

which is to say that the process γ_t only increases at times $t \in [0,1]$ when $\boldsymbol{x}_t^{\boldsymbol{u}}$ hits the boundary, i.e., when $\boldsymbol{x}_t^{\boldsymbol{u}} \in \partial\mathcal{X}$. Thus, (4.1b) is a controlled reflected SDE, and the tuple $(\boldsymbol{x}_t^{\boldsymbol{u}}, \gamma_t)$ solves the *Skorokhod problem* [137, 197, 198]. We point the readers to [151] for the proof of existence and uniqueness of solutions to (4.1b) under the stated regularity assumptions.

To formalize the probabilistic setting of the problem at hand, let Ω be the space of continuous functions $\omega : [0,1] \mapsto \overline{\mathcal{X}}$. We view Ω as a complete separable metric space endowed with the topology of uniform convergence on compact time

[‡]The notation $\boldsymbol{x} \sim \rho$ means that the random vector $\boldsymbol{x}$ has joint PDF ρ.

intervals. With Ω, we associate the σ-algebra $\mathscr{F} = \sigma\{\omega(s) \mid 0 \leq s \leq 1\}$. Consider the complete filtered probability space $(\Omega, \mathscr{F}, \mathbb{P})$ with filtration $\mathscr{F}_t = \sigma\{\omega(s) \mid 0 \leq s \leq t \leq 1\}$ wherein "complete" means that $\mathscr{F}_0$ contains all $\mathbb{P}$-null sets, and $\mathscr{F}_t$ is right continuous. The processes $\boldsymbol{w}_t$, $\boldsymbol{x}_t^{\boldsymbol{u}}$ (for a given feedback policy $\boldsymbol{u}$) and γ_t are $\mathscr{F}_t$-adapted (i.e., non-anticipating) for $t \in [0,1]$. In (4.1c), the random vectors $\boldsymbol{x}_0^{\boldsymbol{u}}$ and $\boldsymbol{x}_1^{\boldsymbol{u}}$ are respectively $\mathscr{F}_0$-measurable and $\mathscr{F}_1$-measurable.

Denote the Euclidean gradient operator as ∇, the inner product as $\langle \cdot, \cdot \rangle$, and the Laplacian as Δ. Letting

$$\mathcal{L} := \theta\Delta + \langle \boldsymbol{f} + \boldsymbol{u}, \nabla \rangle,$$

the law of the sample path of (4.1b) can be characterized [203] as follows: for each $\boldsymbol{x} \in \overline{\mathcal{X}}$, there is a *unique* probability measure $\mathbb{P}_{\boldsymbol{x}}^{\mu}$ on Ω such that $\mathbb{P}_{\boldsymbol{x}}^{\mu}(\boldsymbol{x}_0^{\boldsymbol{u}} = \boldsymbol{x}) = 1$ and for any $\phi \in C_c^{1,2}\left([0,1]; \overline{\mathcal{X}}\right)$ whose inner normal derivative on $\partial\mathcal{X}$ is nonnegative,

$$\phi(t, \boldsymbol{x}_t^{\boldsymbol{u}}) - \int_0^t \left(\frac{\partial \phi}{\partial s} + \mathcal{L}\phi\right)(s, \boldsymbol{x}_s^{\boldsymbol{u}})\ \mathrm{d}s$$

is $\mathbb{P}_{\boldsymbol{x}}^{\mu}$-submartingale, and (iii) there is a continuous, nonnegative, nondecreasing stochastic process γ_t satisfying (4.2). As a consequence [203, p. 196] of this characterization it follows that the process $\boldsymbol{x}_t^{\boldsymbol{u}}$ is Feller continuous and strongly Markov. In particular, the measure-valued trajectory $\mathbb{P}_{\boldsymbol{x}_t^{\boldsymbol{u}}}^{\mu(t)}$ comprises of absolutely continuous measures w.r.t. Lebesgue measure.

The objective in problem (4.1) is to perform the minimum control effort steering of the given initial state PDF ρ_0 at $t = 0$ to the given terminal state PDF ρ_1 at $t = 1$ subject to the controlled sample path dynamics (4.1b). In other words, the data of the problem consists of the domain $\overline{\mathcal{X}}$, the prior dynamics data $\boldsymbol{f}, \theta$, and the two endpoint PDFs ρ_0, ρ_1.

Formally, we can transcribe (4.1) into the following variational problem [21]:

$$\inf_{(\rho,\boldsymbol{u})\in P_2(\overline{X})\times U} \quad \int_0^1\int_{\overline{\mathcal{X}}} \frac{1}{2}\|\boldsymbol{u}(t,\boldsymbol{x}_t^u)\|_2^2\,\rho(t,\boldsymbol{x}_t^u)\,\mathrm{d}\boldsymbol{x}_t^u\mathrm{d}t \tag{4.3a}$$

$$\text{subject to} \quad \frac{\partial\rho}{\partial t} + \nabla\cdot(\rho(\boldsymbol{u}+\boldsymbol{f})) = \theta\Delta\rho, \tag{4.3b}$$

$$\langle -(\boldsymbol{u}+\boldsymbol{f})\rho + \theta\nabla\rho, \boldsymbol{\nu}\rangle|_{\partial\mathcal{X}} = 0, \tag{4.3c}$$

$$\rho(0,\boldsymbol{x}_t^u) = \rho_0, \quad \rho(1,\boldsymbol{x}_t^u) = \rho_1, \tag{4.3d}$$

where a PDF-valued curve $\rho(t,\cdot) \in \mathcal{P}_2(\overline{\mathcal{X}})$ if for each $t \in [0,1]$, the PDF ρ is supported on $\overline{\mathcal{X}}$, and has finite second moment. In this paper, we will not focus on the rather technical direction of establishing the existence of minimizer for (4.3), which can be pursued along the lines of [210, p. 243–245]. Instead, we will formally derive the conditions of optimality, convert them to the so-called Schrödinger system, and argue the existence-uniqueness of solutions for the same.

4.2.2 Necessary Conditions of Optimality

The following result summarizes how the optimal pair $(\rho^{\mathrm{opt}}, \boldsymbol{u}^{\mathrm{opt}})$ for problem (4.3) can be obtained.

Theorem 2 (**Optimal control and optimal state PDF**). *A pair* $(\rho^{\mathrm{opt}}, \boldsymbol{u}^{\mathrm{opt}})$ *solving the variational problem (4.3) must satisfy the system of coupled nonlinear PDEs:*

$$\frac{\partial\rho^{\mathrm{opt}}}{\partial t} + \nabla\cdot\left(\rho^{\mathrm{opt}}(\nabla\psi+\boldsymbol{f})\right) = \theta\Delta\rho^{\mathrm{opt}}, \tag{4.4a}$$

$$\frac{\partial\psi}{\partial t} + \frac{1}{2}\|\nabla\psi\|_2^2 + \langle\nabla\psi,\boldsymbol{f}\rangle = -\theta\Delta\psi, \tag{4.4b}$$

where

$$\boldsymbol{u}^{\text{opt}}(t,\cdot) = \nabla\psi(t,\cdot), \tag{4.5}$$

subject to the boundary conditions

$$\langle \nabla\psi, \boldsymbol{\nu}\rangle|_{\partial\mathcal{X}} = 0, \quad \textit{for all} \quad t \in [0,1], \tag{4.6a}$$

$$\rho^{\text{opt}}(0,\cdot) = \rho_0, \quad \rho^{\text{opt}}(1,\cdot) = \rho_1, \tag{4.6b}$$

$$\langle \rho^{\text{opt}}(\nabla\psi + \boldsymbol{f}) - \theta\nabla\rho^{\text{opt}}, \boldsymbol{\nu}\rangle|_{\partial\mathcal{X}} = 0, \quad \textit{for all} \quad t \in [0,1]. \tag{4.6c}$$

The PDE (4.4a) is a controlled Fokker-Planck-Kolmogorov (FPK) equation, and (4.4b) is a Hamilton-Jacobi-Bellman (HJB) equation. Because the equations (4.4a)-(4.4b) have one way coupling, and the boundary conditions (4.6a)-(4.6c) are atypical, solving (4.4) is a challenging task in general. In the following, we show that it is possible to transform the *coupled nonlinear* system (4.4) into a boundary coupled *linear* system of PDEs which we refer to as the *Schrödinger system*. We will see that the resulting system paves way to a computational pipeline for solving the density steering problem with path constraints.

4.2.3 Schrödinger System

We now apply the Hopf-Cole transform [68, 114] to the system of nonlinear PDEs (4.4).

Theorem 3 (**Schrödinger system**). *Given the data* $\overline{\mathcal{X}}, \boldsymbol{f}, \theta, \rho_0, \rho_1$ *for problem*

(4.3), consider the Hopf-Cole transform $(\rho^{\mathrm{opt}}, \psi) \mapsto (\varphi, \hat{\varphi})$ *given by*

$$\varphi(t,\cdot) := \exp\left(\psi(t,\cdot)/2\theta\right), \tag{4.7a}$$

$$\hat{\varphi}(t,\cdot) := \rho^{\mathrm{opt}}(t,\cdot)\exp\left(-\psi(t,\cdot)/2\theta\right), \tag{4.7b}$$

applied to (4.4) where $t \in [0,1]$*. For* $k \in \{0,1\}$*, introduce the notation* $\varphi_k := \varphi(k,\cdot)$*,* $\hat{\varphi}_k := \hat{\varphi}(k,\cdot)$*. Then the pair* $(\varphi, \hat{\varphi})$ *satisfies the system of linear PDEs*

$$\frac{\partial \varphi}{\partial t} = -\langle \nabla\varphi, \boldsymbol{f}\rangle - \theta\Delta\varphi, \tag{4.8a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = -\nabla\cdot(\boldsymbol{f}\hat{\varphi}) + \theta\Delta\hat{\varphi}, \tag{4.8b}$$

subject to the boundary conditions

$$\varphi_0\hat{\varphi}_0 = \rho_0, \quad \varphi_1\hat{\varphi}_1 = \rho_1, \tag{4.9a}$$

$$\langle \nabla\varphi, \boldsymbol{\nu}\rangle|_{\partial\mathcal{X}} = \langle \boldsymbol{f}\hat{\varphi} - \theta\nabla\hat{\varphi}, \boldsymbol{\nu}\rangle|_{\partial\mathcal{X}} = 0. \tag{4.9b}$$

For all $t \in [0,1]$*, the pair* $(\rho^{\mathrm{opt}}, \boldsymbol{u}^{\mathrm{opt}})$ *can be recovered as*

$$\rho^{\mathrm{opt}}(t,\cdot) = \varphi(t,\cdot)\hat{\varphi}(t,\cdot), \ \boldsymbol{u}^{\mathrm{opt}}(t,\cdot) = 2\theta\nabla\log\varphi(t,\cdot). \tag{4.10}$$

Remark 5. *From (4.7), both* $\varphi, \hat{\varphi}$ *are nonnegative by definition, and strictly positive if* ψ *is bounded and* ρ^{opt} *is positive.*

Remark 6. *Under the regularity assumptions on* $\boldsymbol{f}$ *and* $\overline{\mathcal{X}}$ *stated in Section 4.2.1, the process* $\boldsymbol{x}_t$ *satisfying the* uncontrolled *reflected Itô SDE*

$$\mathrm{d}\boldsymbol{x}_t = \boldsymbol{f}(t,\boldsymbol{x}_t)\,\mathrm{d}t + \sqrt{2\theta}\,\mathrm{d}\boldsymbol{w}_t + \boldsymbol{n}(\boldsymbol{x}_t)\,\mathrm{d}\gamma_t,\ t \in [0,1], \tag{4.11}$$

is a Feller continuous strongly Markov process. Therefore, the theory of semi-

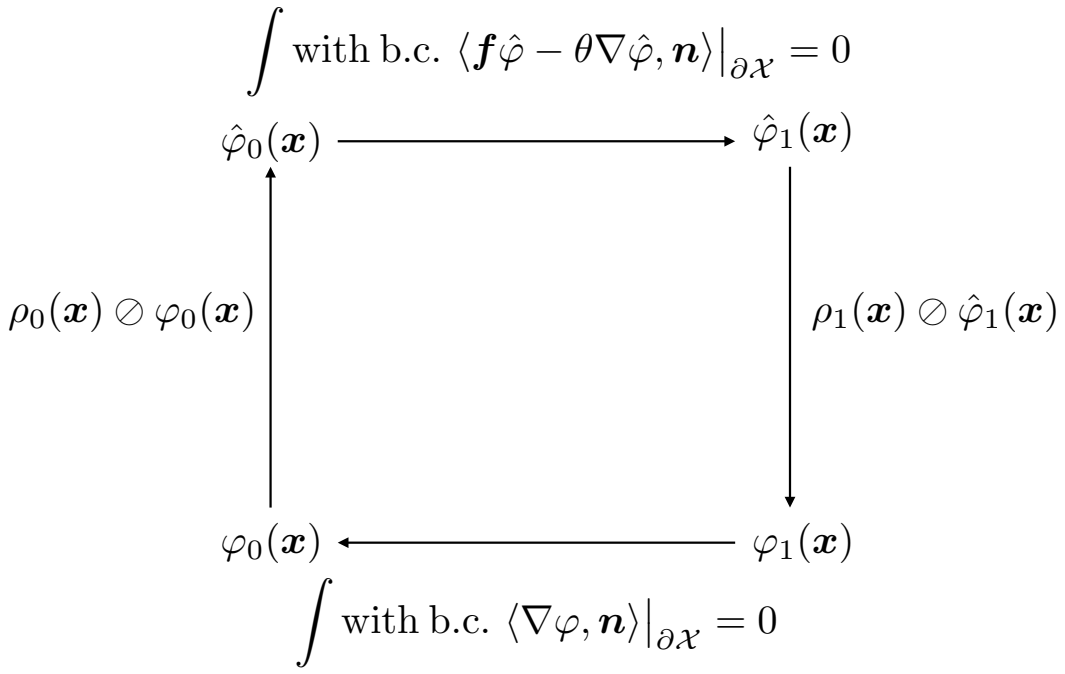

Figure 4.1: Schematic of the fixed point recursion for the Schrödinger system (4.8)-(4.9). The abbreviation "b.c." stands for boundary condition, the symbol $\oslash$ denotes the Hadamard division.

groups applies and the transition density of (4.11) satisfies Kolmogorov's equations. Notice that the transition density or Green's function will depend on the domain $\overline{\mathcal{X}}$. *In particular, we point out that (4.8a) is the* backward Kolmogorov equation *in unkonwn* φ *with the corresponding* Neumann boundary condition $\langle \nabla\varphi, \boldsymbol{\nu}\rangle|_{\partial\mathcal{X}} = 0$ *in (4.9b). On the other hand, (4.8b) is the* forward Kolmogorov equation *in unkonwn* $\hat{\varphi}$ *with the corresponding* Robin boundary condition $\langle \boldsymbol{f}\hat{\varphi} - \theta\nabla\hat{\varphi}, \boldsymbol{\nu}\rangle|_{\partial\mathcal{X}} = 0$ *in (4.9b). These "backward Kolmogorov with Neumann" and "forward Kolmogorov with Robin" system of PDE boundary value problems are coupled via the atypical boundary conditions (4.9a).*

Theorem 3 reduces finding the optimal pair $(\rho^{\text{opt}}, \boldsymbol{u}^{\text{opt}})$ for the RSBP to that of finding the pair[§] $(\varphi(t, \boldsymbol{x}_t), \hat{\varphi}(t, \boldsymbol{x}_t))$ associated with the uncontrolled SDE (4.11). To do so, we need to compute the terminal-initial condition pair $(\varphi_1, \hat{\varphi}_0)$, which

[§]We refer to $\varphi(t, \boldsymbol{x}_t), \hat{\varphi}(t, \boldsymbol{x}_t)$ as the *Schrödinger factors*.

can be obtained by first making an initial guess for $(\varphi_1, \hat{\varphi}_0)$ and then performing time update by integrating the system (4.8)-(4.9b). Using (4.9a), this then sets up a fixed point recursion over the pair $(\varphi_1, \hat{\varphi}_0)$ (see Fig. 4.1). If this recursion converges to a unique pair, then the converged pair $(\varphi_1, \hat{\varphi}_0)$ can be used to compute the transient factors $(\varphi(t, \boldsymbol{x}_t), \hat{\varphi}(t, \boldsymbol{x}_t))$, and we can recover $(\rho^{\text{opt}}, \boldsymbol{u}^{\text{opt}})$ via (4.10). This computational pipeline will be pursued in this paper.

Since the PDEs in (4.8) are linear, and the boundary couplings in (4.9a) are in product form, the nonnegative function pair $(\varphi_1, \hat{\varphi}_0)$ can only be unique in the projective sense, i.e., if $(\varphi_1, \hat{\varphi}_0)$ is a solution then so is $(\alpha\varphi_1, \hat{\varphi}_0/\alpha)$ for any $\alpha > 0$. In [53], it was shown that the aforesaid fixed point recursion is in fact *contractive* on a suitable cone in Hilbert's projective metric, and hence guaranteed to converge to a unique pair $(\varphi_1, \hat{\varphi}_0)$, provided that the transition density for (4.11) is positive and continuous[¶] on $\overline{\mathcal{X}} \times \overline{\mathcal{X}}$ for all $t \in [0, 1]$, and ρ_0, ρ_1 are supported on compact subsets of $\overline{\mathcal{X}}$.

4.3 Case Study: RSBP in 1D without Prior Drift

To illustrate the ideas presented thus far, we now consider a simple instance of problem (4.3) over the state space $\overline{\mathcal{X}} = [a, b] \subset \mathbb{R}$, and with the prior drift $f \equiv 0$. That is to say, we consider the finite horizon density steering subject to the controlled two-sided reflected Brownian motion. Using some properties of the associated Markov kernel, we will show that the Schrödinger system (4.8)-(4.9) corresponding to this particular RSBP has a unique solution which can be obtained by the kind of fixed point recursion mentioned toward the end of Section

[¶]Under the regularity assumptions on $\boldsymbol{f}$ and $\overline{\mathcal{X}}$ stated in Section 4.2.1, the transition density for (4.11) indeed satisfies these conditions.

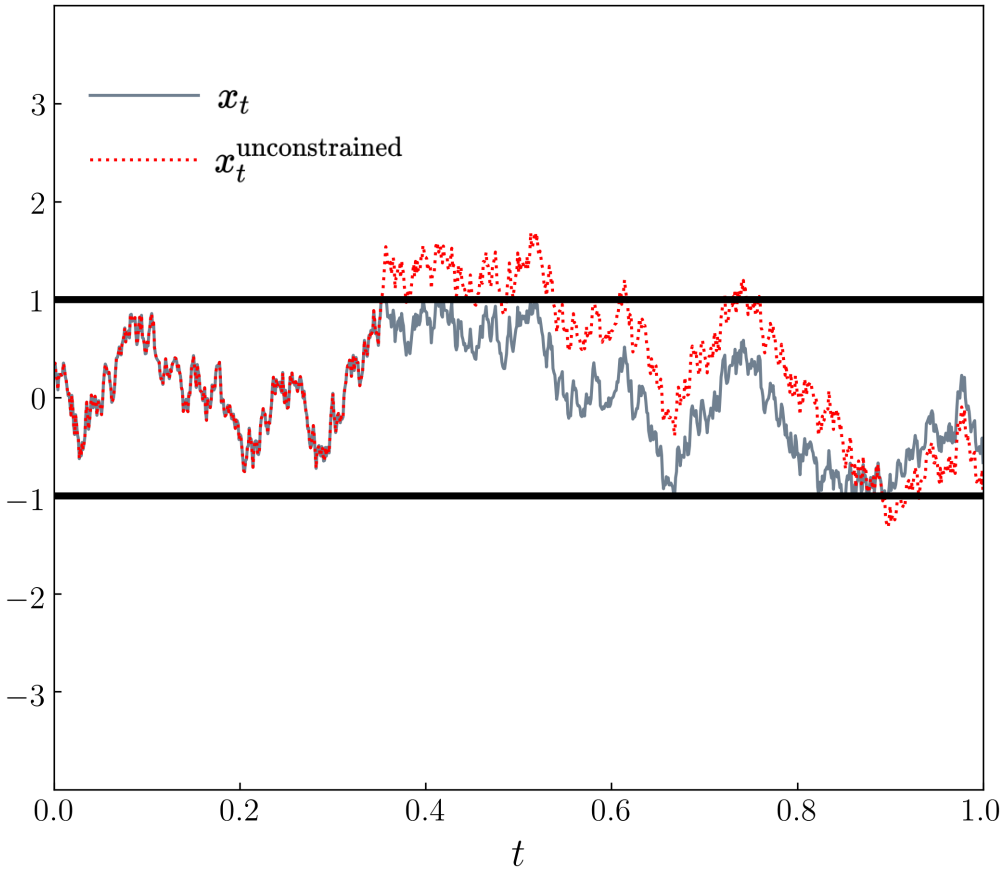

Figure 4.2: For $t \in [0,1]$, the *solid line* shows a sample path x_t for (4.15) with $[a,b] \equiv [-1,1]$, $\theta = 0.5$. The *dotted line* shows the corresponding unconstrained sample path $x_t^{\text{unconstrained}}$, computed using the two-sided Skorokhod map [137].

4.2.3.

In this case, the Schrödinger system (4.8)-(4.9) reduces to

$$\frac{\partial \varphi}{\partial t} = -\theta \frac{\partial^2 \varphi}{\partial x^2}, \tag{4.12a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = \theta \frac{\partial^2 \hat{\varphi}}{\partial x^2}, \tag{4.12b}$$

$$\varphi_0 \hat{\varphi}_0 = \rho_0, \quad \varphi_1 \hat{\varphi}_1 = \rho_1, \tag{4.12c}$$

$$\frac{\partial \varphi}{\partial x}\Big|_{x=a,b} = \frac{\partial \hat{\varphi}}{\partial x}\Big|_{x=a,b} = 0. \tag{4.12d}$$

Notice that (4.12a)-(4.12b) are the backward and forward heat PDEs, respectively,

which subject to (4.12d), have solutions

$$\varphi(x,t) = \int_{[a,b]} K_\theta(x,y,1-t)\varphi_1(y)\,\mathrm{d}y, \quad t \le 1, \tag{4.13a}$$

$$\hat{\varphi}(x,t) = \int_{[a,b]} K_\theta(y,x,t)\hat{\varphi}_0(y)\,\mathrm{d}y, \qquad t \ge 0, \tag{4.13b}$$

where

$$\begin{aligned} K_\theta(x,y,t) := \frac{1}{b-a} + \frac{2}{b-a}\sum_{m=1}^{\infty} \exp\left(-\frac{\theta\pi^2 m^2}{(b-a)^2}t\right) \\ \times \cos\left(\frac{m\pi(x-a)}{b-a}\right)\cos\left(\frac{m\pi(y-a)}{b-a}\right) \end{aligned} \tag{4.14}$$

is the Markov kernel or transition density [150, Sec. 4.1], [26, p. 410-411] associated with the uncontrolled reflected SDE

$$\mathrm{d}x_t = \sqrt{2\theta}\,\mathrm{d}w_t + \mathrm{d}L_t - \mathrm{d}U_t, \quad t \in [0,1]. \tag{4.15}$$

In (4.15), L_t, U_t are the two local time stochastic processes [89, 109] at the lower and upper boundaries respectively, which restrict x_t to the interval $[a,b]$; see Fig. 4.2.

Combining (4.13) and (4.12c), we get a system of coupled nonlinear integral equations in unknowns $(\varphi_1, \hat{\varphi}_0)$, given by

$$\rho_0(x) = \hat{\varphi}_0(x)\int_{[a,b]} K_\theta(x,y,1)\varphi_1(y)\,\mathrm{d}y, \tag{4.16a}$$

$$\rho_1(x) = \varphi_1(x)\int_{[a,b]} K_\theta(y,x,1)\hat{\varphi}_0(y)\,\mathrm{d}y. \tag{4.16b}$$

Clearly, solving (4.16) is equivalent to solving (4.12). The pair $(\varphi_1, \hat{\varphi}_0)$ can be solved from (4.16) iteratively as a fixed point recursion with guaranteed con-

vergence established through contraction mapping in Hilbert's projective metric; see [53]. The Lemma 9 stated next will be used in the Proposition 10 that follows, showing the existence and uniqueness of the pair $(\varphi_1, \hat{\varphi}_0)$ in (4.16) as well as the fact that the aforesaid fixed point recursion is guaranteed to converge to that pair.

Lemma 9. *For $0 < \theta$, $a < b$, consider the transition probability density $K_\theta(x, y, t)$ in (4.14). Then,*

(i) $K_\theta(x, y, t = 1)$ is continuous on the set $[a, b] \times [a, b]$.

(ii) $K_\theta(x, y, t = 1) > 0$ for all $(x, y) \in [a, b] \times [a, b]$.

Proposition 10. *Given $0 < \theta$, $a < b$, and the endpoint PDFs ρ_0, ρ_1 having compact supports $\subseteq [a, b]$. There exists a unique pair $(\varphi_1, \hat{\varphi}_0)$ that solves (4.16), and equivalently (4.12). Moreover, this unique pair can be computed by the fixed point recursion shown in Fig. 4.1.*

To illustrate how the above results can be used for practical computation, consider solving the RSBP (4.1) with $f \equiv 0$, $\theta = 0.5$, $\overline{\mathcal{X}} = [a, b] \equiv [-4, 4]$, and ρ_0, ρ_1 as (see Fig. 4.3)

$$\rho_0(x) \propto 1 + (x^2 - 16)^2 \exp(-x/2), \tag{4.17a}$$

$$\rho_1(x) \propto 1.2 - \cos(\pi(x + 4)/2), \tag{4.17b}$$

where the supports of (4.17) are restricted to $[-4, 4]$, and the proportionality constants are determined accordingly. For state feedback synthesis enabling this unimodal to bimodal steering over $t \in [0, 1]$, we performed the fixed point recursion over the pair $(\varphi_1, \hat{\varphi}_0)$ using (4.16) with ρ_0, ρ_1 as in (4.17), and K_θ given by (4.14). For numerical implementation, we truncated the infinite sum in (4.14) after the first 100 terms. Fig. 4.4 shows the convergence of this fixed point recursion

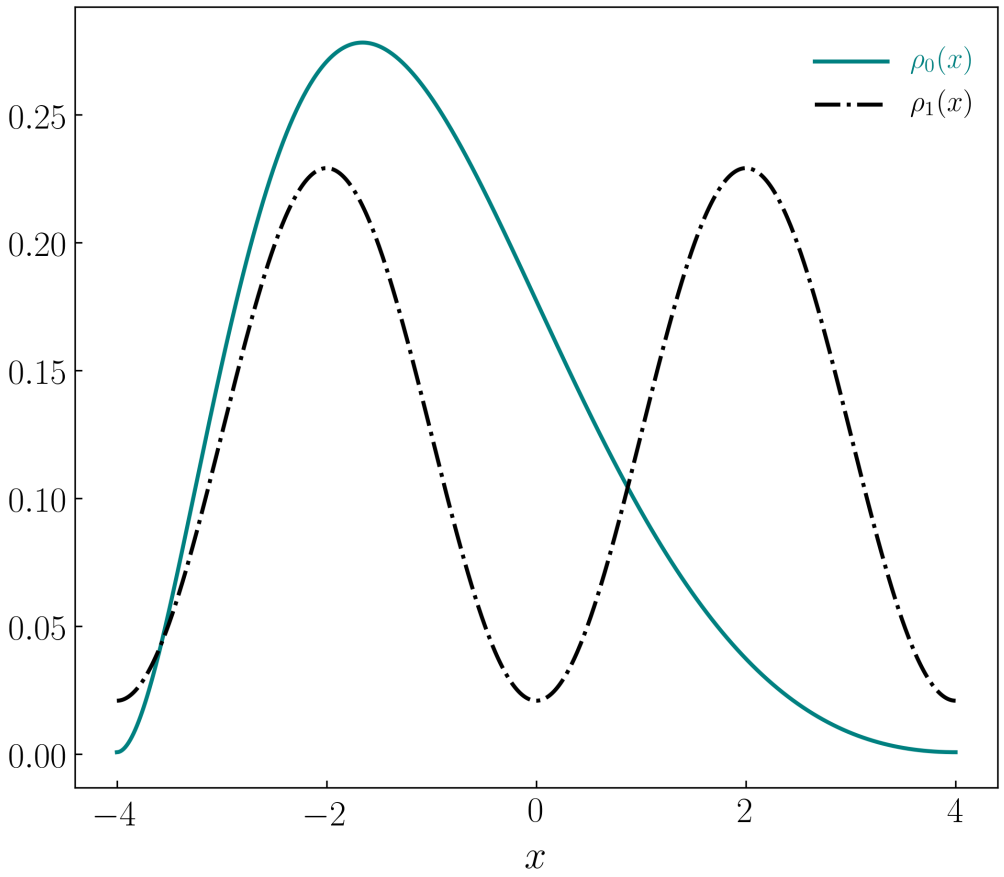

Figure 4.3: The endpoint PDFs ρ_0, ρ_1 shown above are supported on $[-4, 4]$, and are given by (4.17).

w.r.t. Hilbert's projective metric. The converged pair $(\varphi_1, \hat{\varphi}_0)$ is used to compute the transient Schrödinger factors $(\varphi(t, \boldsymbol{x}_t), \hat{\varphi}(t, \boldsymbol{x}_t))$ via (4.13), and then the pair $(\rho^{\mathrm{opt}}(t, \boldsymbol{x}_t^u), \boldsymbol{u}^{\mathrm{opt}}(t, \boldsymbol{x}_t^u))$ via (4.10). Fig. 4.5 depicts the evolution of the optimal controlled transient joint state PDFs $\rho^{\mathrm{opt}}(t, x_t^u)$ as well as 100 sample paths x_t^u of the optimal closed-loop reflected SDE. These sample paths were computed by applying the Euler-Maruyama scheme with time-step size 10^{-3}. Notice from Fig. 4.5 that (i) the closed-loop sample paths satisfy $-4 \leq x_t^u \leq 4$ for all $t \in [0, 1]$, and (ii) in the absence of feedback, the terminal constraint $\rho(1, x_1^u) = \rho_1$ (given by (4.17b)) cannot be satisfied.

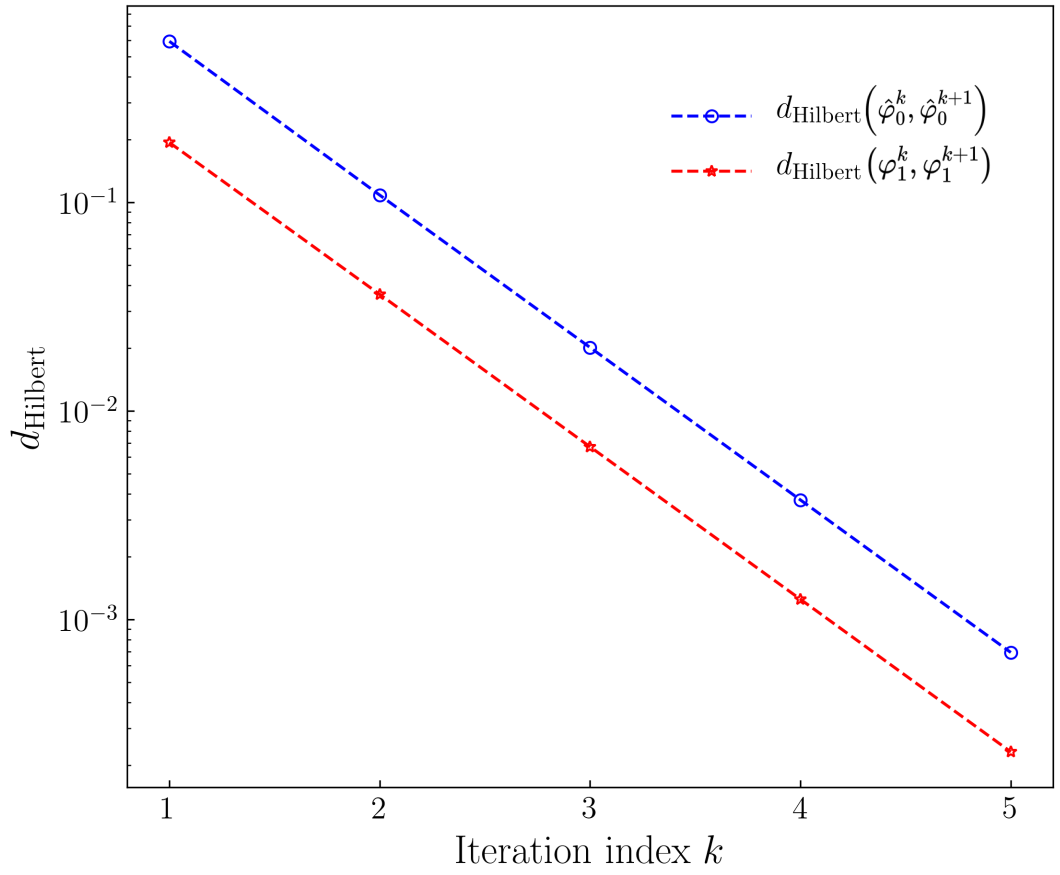

Figure 4.4: Convergence of the fixed point recursion over $(\varphi_1, \hat{\varphi}_0)$ in Hilbert's projective metric d_{Hilbert}.

4.4 RSBP with Prior Drift

For generic $\boldsymbol{f}$, $\overline{\mathcal{X}}$, there is no closed-form expression of the Markov Kernel associated with (4.8)-(4.9b). Hence, unlike the situation in Section 4.3, we cannot explicitly set up coupled integral equations of the form (4.16), thus preventing the numerical implementation of the fixed point recursion (Fig. 4.1) via direct matrix-vector recursion. In this Section, we will show that if $\boldsymbol{f}$ is gradient of a potential, then we can reformulate (4.8)-(4.9) in a way that leads to a variational recursion which in turn enables us to implement the fixed point recursion (Fig. 4.1) in an implicit manner.

4.4.1 Reformulation of the Schrödinger System

Let $\boldsymbol{f}$ be a gradient vector field, i.e., $\boldsymbol{f} = -\nabla V$ for some potential $V \in C^2(\overline{\mathcal{X}})$. The associated Schrödinger system (4.8)-(4.9) becomes

$$\frac{\partial \varphi}{\partial t} = \langle \nabla \varphi, \nabla V \rangle - \theta \Delta \varphi, \tag{4.18a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = \nabla \cdot (\nabla V \hat{\varphi}) + \theta \Delta \hat{\varphi}, \tag{4.18b}$$

$$\varphi_0 \hat{\varphi}_0 = \rho_0, \quad \varphi_1 \hat{\varphi}_1 = \rho_1, \tag{4.18c}$$

$$\langle \nabla \varphi, \boldsymbol{\nu} \rangle|_{\partial \mathcal{X}} = \langle \nabla V \hat{\varphi} + \theta \nabla \hat{\varphi}, \boldsymbol{\nu} \rangle|_{\partial \mathcal{X}} = 0. \tag{4.18d}$$

The idea now is to exploit the structural nonlinearities in (4.18) to design an algorithm that allows computing the Schrödinger factors $(\varphi, \hat{\varphi})$. To that end, the following is a crucial step.

Theorem 4. *Given $V \in C^2(\overline{\mathcal{X}}), \theta > 0$, and $t \in [0,1]$, consider $\varphi(t, \boldsymbol{x}_t)$ in (4.18). Let $s := 1 - t$, and define the mappings $\varphi \mapsto q \mapsto p$ given by $q(s, \boldsymbol{x}_s) := \varphi(t, \boldsymbol{x}_t) =$*

$\varphi(1-s, \boldsymbol{x}_{1-s})$, $p(s, \boldsymbol{x}_s) := q(s, \boldsymbol{x}_s)\exp(-V(\boldsymbol{x}_s)/\theta)$. *Then* $p(s, \boldsymbol{x}_s)$ *solves the PDE initial boundary value problem:*

$$\frac{\partial p}{\partial s} = \nabla \cdot (p\nabla V) + \theta \Delta p, \tag{4.19a}$$

$$p(0, \boldsymbol{x}) = \varphi_1(\boldsymbol{x})\exp(-V(\boldsymbol{x})/\theta), \tag{4.19b}$$

$$\langle \nabla V p + \theta \nabla p, \boldsymbol{\nu} \rangle|_{\partial \mathcal{X}} = 0. \tag{4.19c}$$

Thanks to Theorem 4, solving (4.18) is equivalent to solving

$$\frac{\partial p}{\partial s} = \nabla \cdot (p\nabla V) + \theta \Delta p, \tag{4.20a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} = \nabla \cdot (\nabla V \hat{\varphi}) + \theta \Delta \hat{\varphi}, \tag{4.20b}$$

$$p(s=1, \boldsymbol{x})\exp(V(\boldsymbol{x})/\theta)\hat{\varphi}_0(\boldsymbol{x}) = \rho_0,$$

$$p(s=0, \boldsymbol{x})\exp(V(\boldsymbol{x})/\theta)\hat{\varphi}_1(\boldsymbol{x}) = \rho_1, \tag{4.20c}$$

$$\langle \nabla V p + \theta \nabla p, \boldsymbol{\nu} \rangle|_{\partial \mathcal{X}} = \langle \nabla V \hat{\varphi} + \theta \nabla \hat{\varphi}, \boldsymbol{\nu} \rangle|_{\partial \mathcal{X}} = 0. \tag{4.20d}$$

From (4.20a)-(4.20b), φ and p satisfy the exact same FPK PDE with different initial conditions and integrated in different time coordinates t and s. From (4.20d), φ and p satisfy the same Robin boundary condition. Therefore, a single FPK initial boundary value problem solver can be used to set up the fixed point recursion to solve for $(p_1, \hat{\varphi}_0)$, and hence $(p(s, \boldsymbol{x}_s), \hat{\varphi}(t, \boldsymbol{x}_t))$. From p, we can recover φ as

$$\varphi(t, \boldsymbol{x}_t) = \varphi(1-s, \boldsymbol{x}_{1-s}) = p(s, \boldsymbol{x}_s)\exp(-V(\boldsymbol{x}_s)/\theta).$$

4.4.2 Computation via Wasserstein Proximal Recursion

Building on our previous works [41, 43, 44], we propose proximal recursions to numerically time march the solutions of the PDE initial boundary value problems (4.20) by exploiting certain infinite dimensional gradient descent structure. This enables us to perform the computation associated with the horizontal arrows in Fig. 4.1, and hence the fixed point recursions to solve for $(p, \hat{\varphi})$, and consequently for $(\varphi, \hat{\varphi})$. We give here a brief outline of the ideas behind these proximal recursions.

It is well-known [125, 190] that the flows generated by (4.20a),(4.20b),(4.20d) can be viewed as the gradient descent of the Lyapunov functional

$$F(\varrho) := \int_{\overline{\mathcal{X}}} V(\boldsymbol{x})\varrho(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} \; + \theta \int_{\overline{\mathcal{X}}} \varrho(\boldsymbol{x}) \log \varrho(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} \tag{4.21}$$

w.r.t. the distance metric W referred to as the (quadratic) Wassertein metric [210] on $\mathcal{P}_2(\overline{\mathcal{X}})$. For chosen time-steps τ, σ, this allows us to set up a variational recursion over the discrete time pair $(t_{k-1}, s_{k-1}) := ((k-1)\tau, (k-1)\sigma)$ as

$$\begin{pmatrix} \hat{\phi}_{t_k} \\ \varpi_{s_k} \end{pmatrix} = \begin{pmatrix} \mathrm{prox}^{W^2}_{\tau F}(\hat{\phi}_{t_{k-1}}) \\ \mathrm{prox}^{W^2}_{\sigma F}(\varpi_{s_{k-1}}) \end{pmatrix}, \quad k \in \mathbb{N}, \tag{4.22}$$

wherein the Wasserstein proximal operator

$$\mathrm{prox}^{W^2}_{hF}(\cdot) := \underset{\varrho \in P_2(\overline{X})}{arg\,inf} \frac{1}{2} W^2(\cdot, \varrho) + hF(\varrho), \quad h > 0. \tag{4.23}$$

The sequence of functions generated by the proximal recursions (4.22) approximate the flows $(p(s, \boldsymbol{x}_s), \hat{\varphi}(t, \boldsymbol{x}_t))$ for (4.20a),(4.20b),(4.20d) in the small time step

limit, i.e.,

$$\hat{\phi}_{t_{k-1}} \to \hat{\varphi}(t=(k-1)\tau, \boldsymbol{x}_t) \quad \text{in } L^1(\overline{\mathcal{X}}) \text{ as } \tau \downarrow 0,$$
$$\varpi_{s_{k-1}} \to p(s=(k-1)\sigma, \boldsymbol{x}_s) \quad \text{in } L^1(\overline{\mathcal{X}}) \text{ as } \sigma \downarrow 0.$$

In the numerical example provided next, we solved (4.22) using the algorithm developed in [41].

4.4.3 Numerical Example

We consider an instance of the RSBP with $\overline{\mathcal{X}} = [-4,4]^2$, $\boldsymbol{f} = -\nabla V$, $V(x_1,x_2) := (x_1^2 + x_2^3)/5$. For

$$\rho_0(x_1,x_2) \propto \prod_{i=1,2} \left(1 + (x_i^2 - 16)^2 \exp(-x_i/2)\right), \tag{4.24a}$$
$$\rho_1(x_1,x_2) \propto \prod_{i=1,2} \left(1.2 - \cos(\pi(x_i+4)/2)\right), \tag{4.24b}$$

the optimal controlled joint state PDFs $\rho^{\text{opt}}(t, \boldsymbol{x}_t^{\boldsymbol{u}})$ are shown in Fig. 4.6. The corresponding uncontrolled joint state PDFs $\rho^{\text{unc}}(t, \boldsymbol{x}_t)$ are shown in Fig. 4.7. These results were obtained by solving (4.22) via [41, Sec. III.B] with $\tau = \sigma = 10^{-3}$ to perform the fixed point recursion (Fig. 4.1) applied to (4.20).

4.5 Conclusions

In this paper, we introduced the Reflected Schrödinger Bridge Problem (RSBP) – a stochastic optimal control problem for minimum energy feedback steering of a given joint PDF to another over finite horizon subject to reflecting bound-

ary conditions on the controlled state trajectories. Combining our prior work on Wasserstein proximal recursions with some recent results on contraction mapping associated with the Schrödinger system, we provide a computational pipeline for optimal feedback synthesis. Numerical examples are given to highlight the proposed framework.

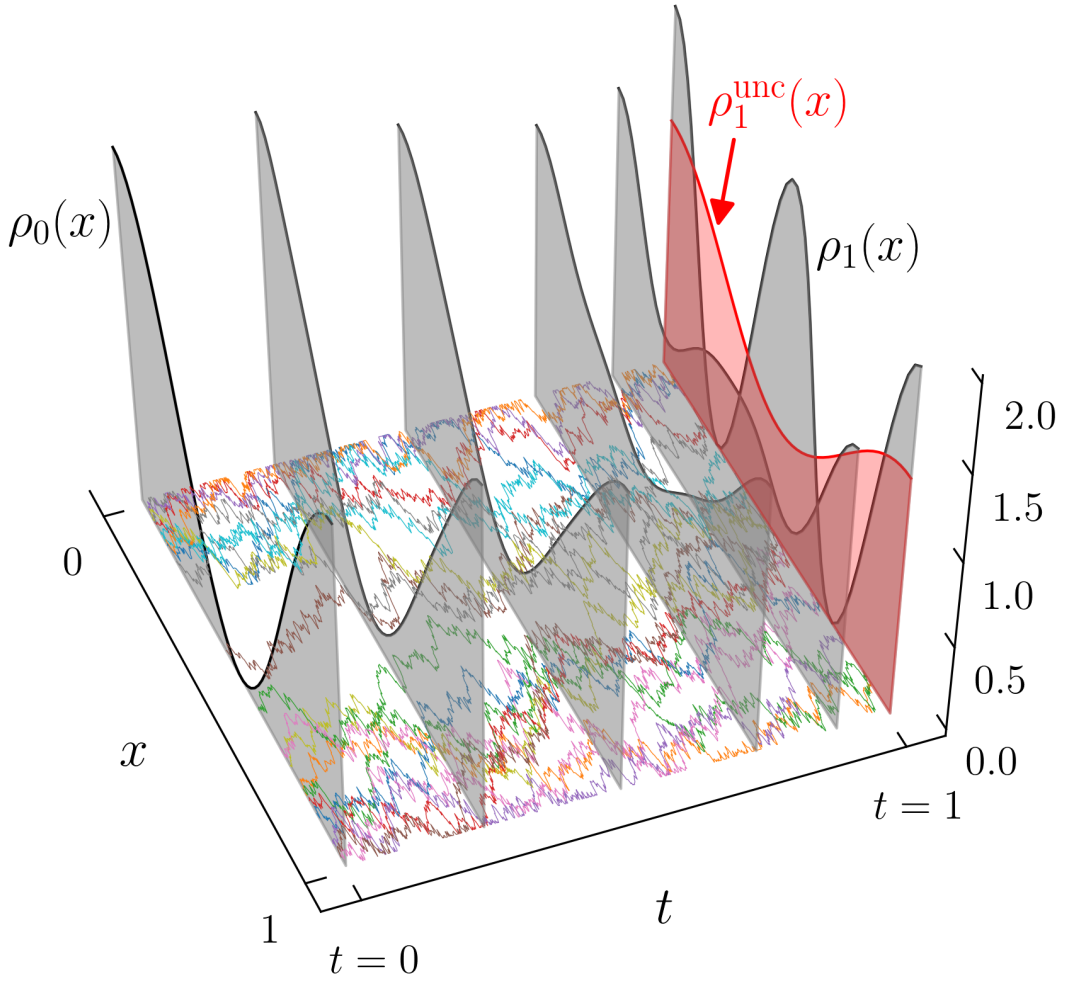

Figure 4.5: Shown as the black curves are the optimal controlled transient joint state PDFs $\rho^{\mathrm{opt}}(t, x_t^u)$ for steering the two-sided reflecting Brownian motion with endpoint PDFs ρ_0, ρ_1 as in Fig. 4.3. The red curve ρ_1^{unc} is the uncontrolled state PDF at $t = 1$, i.e., obtained by setting $u \equiv 0$. Also depicted are the 100 sample paths of the optimally controlled (i.e., closed-loop) reflected SDE. This simulation corresponds to the RSBP (4.1) with problem data $f \equiv 0, [a, b] = [-4, 4], \theta = 0.5$, and ρ_0, ρ_1 given by (4.17).

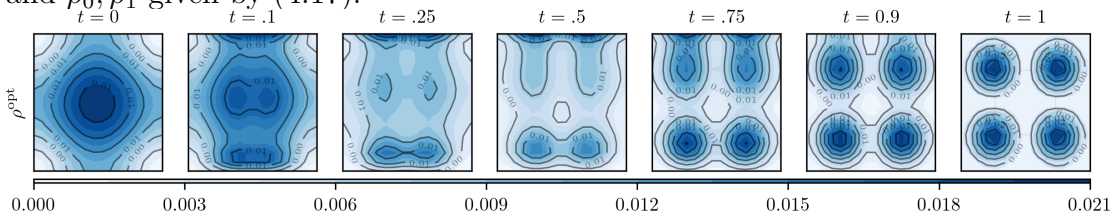

Figure 4.6: For the RSBP in Section 4.4.2, shown here are the contour plots of the optimal controlled joint state PDFs $\rho^{\mathrm{opt}}(t, \boldsymbol{x}_t^{\boldsymbol{u}})$ over $\overline{\mathcal{X}} = [-4, 4]^2$. Each subplot corresponds to a different snapshot of ρ^{opt} in time. The color denotes the joint PDF value; see colorbar (dark hue = high, light hue = low).

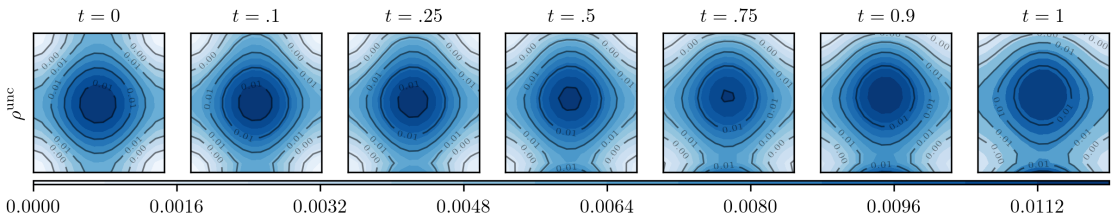

Figure 4.7: For the RSBP in Section 4.4.2, shown here are the contour plots of the uncontrolled joint state PDFs $\rho^{\text{unc}}(t, \boldsymbol{x}_t)$ over $\overline{\mathcal{X}} = [-4, 4]^2$ starting from (4.24a). Each subplot corresponds to a different snapshot of ρ^{unc} in time. The color denotes the joint PDF value; see colorbar (dark hue = high, light hue = low).

Chapter 5

Finite Horizon Density Steering for Multi-input State Feedback Linearizable Systems

In this chapter, we study the feedback synthesis problem for steering the joint state density or ensemble subject to multi-input state feedback linearizable dynamics. This problem is of interest to many practical applications including that of dynamically shaping a robotic swarm. Our results here show that it is possible to exploit the structural nonlinearities to derive the feedback controllers steering the joint density from a prescribed shape to another while minimizing the expected control effort to do so. The developments herein build on our previous work, and extend the theory of the Schrödinger bridge problem subject to feedback linearizable dynamics.

5.1 Introduction

We consider the problem of steering the statistics of the state vector $\boldsymbol{x}(t)$ from a prescribed ensemble or joint density $\rho_0(\boldsymbol{x})$ to another $\rho_1(\boldsymbol{x})$ over a finite time horizon $t \in [0, 1]$, subject to controlled nonlinear dynamics of the form

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{u}, \quad \boldsymbol{x} \in \mathcal{X} \subseteq \mathbb{R}^n, \quad \boldsymbol{u} \in \mathbb{R}^m, \tag{5.1}$$

where $\boldsymbol{f}$ is a smooth vector field on the state space $\mathcal{X} \subseteq \mathbb{R}^n$, and $\boldsymbol{G}$ is an $n \times m$ matrix whose columns consist of the vectors $\boldsymbol{g}_i \in \mathbb{R}^n$ for $i = 1, \ldots, m$, i.e.,

$$\boldsymbol{G}(\boldsymbol{x}) = [\boldsymbol{g}_1(\boldsymbol{x})|\boldsymbol{g}_2(\boldsymbol{x})|\ldots|\boldsymbol{g}_m(\boldsymbol{x})]\,. \tag{5.2}$$

It is of broad practical interest to solve this finite horizon density steering problem while minimizing the average total control effort over the *controlled* state ensemble $\rho(\boldsymbol{x}, t)$.

This problem is motivated by the growing need across science and engineering applications to control a *large population of systems*. Consider for example, shaping the bulk magnetization distribution for Nuclear Magnetic Resonance spectroscopy, controlling heterogeneous (e.g., aerial and ground) robotic swarms [15, 79], strategically synchronizing and desynchronizing a neuronal population to regulate the Parkinsonian tremor [163], and differentially moving the setpoints of a large population of residential air-conditioners by a service provider to make their total energy consumption track the intermittency in supply (e.g., due to stochastic renewable generation) in a privacy-preserving manner [62, 99]. These exemplars concern *population ensemble or density* whose shape is actively controlled over time while preserving the physical mass. The conservation of mass allows an

alternative interpretation of the underlying mathematical problem – instead of steering a large number of dynamical systems, one can think of steering a single system with probabilistic uncertainty in its initial and terminal state, i.e., ρ_0, ρ_1 being joint state probability density functions (PDFs). This too, arises naturally in practice, e.g., in robot motion planning [172], where uncertainties in the initial and terminal states are unavoidable due to process and sensor noise.

From a system-control theoretic viewpoint, finite horizon density steering *via feedback* is a non-classical stochastic optimal control problem. The qualifier "non-classical" points to the fact that finding the feedback policy requires solving an infinite dimensional two-point boundary value problem on the manifold of joint state PDFs. This is an emerging research direction in the systems-control community wherein recent advances [56, 58–60, 104] have uncovered its connections with the theory of optimal mass transport [210, 215] and the Schrödinger Bridge Problem (SBP) [145, 191, 192]. Also, there have been results on the covariance steering problem [115, 170, 195, 196, 224] which concerns steering second order state statistics. With the exception of [2, 84], almost all works have focused on steering the state statistics over a controlled linear system.

Here, we consider finite horizon density steering state feedback linearizable systems of the form (5.1). The nonlinearities in (5.1) induce non-Gaussian statistics even if the endpoint PDFs are both jointly Gaussian. Thus, finding the feedback solution of the density steering problem in a non-parametric sense, is non-trivial. The main contribution of this paper is to show that it is possible to exploit the feedback linearizing transformation for density steering. In particular, we obtain the optimal state feedback policy in terms of the solution of certain Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE). Furthermore, we show that a dynamic stochastic regularization can be used to derive a system of boundary-

coupled linear PDEs, which we refer to as the *Schrödinger system*, whose solutions recover the optimal state feedback and the optimal controlled joint state PDF. We envision that the theoretical developments herein will help design algorithms solving the feedback density steering over nonlinear dynamical systems.

5.2 MIMO Feedback Linearization

We consider multiple-input control system of the form (5.1), and recall some well-known results on feedback linearization [**?**] that will be useful in the sequel.

Definition 5. *(**Full state static feedback linearization**) System (5.1) is said to be* full state static feedback linearizable *around a point $\boldsymbol{x}_0 \in \mathcal{X}$ if there exists a smooth feedback of the form $\boldsymbol{u} = \boldsymbol{\delta}(\boldsymbol{x}) + \boldsymbol{\Gamma}(\boldsymbol{x})\boldsymbol{v}$ defined on $\mathcal{X}$, and a diffeomorphism $\boldsymbol{\tau} : \mathcal{X} \mapsto \mathbb{R}^n$ such that the change of variables $\boldsymbol{z} := \boldsymbol{\tau}(\boldsymbol{x})$ transforms (5.1) into*

$$\dot{\boldsymbol{z}} = \boldsymbol{A}\boldsymbol{z} + \boldsymbol{B}\boldsymbol{v}, \quad \boldsymbol{z} \in \mathbb{R}^n, \quad \boldsymbol{v} \in \mathbb{R}^m, \tag{5.3}$$

wherein the pair $(\boldsymbol{A}, \boldsymbol{B})$ is controllable.

In other words, (5.1) is full state static feedback linearizable if there exists a triple $(\boldsymbol{\delta}(\boldsymbol{x}), \boldsymbol{\Gamma}(\boldsymbol{x}), \boldsymbol{\tau}(\boldsymbol{x}))$ such that

$$\begin{aligned} \left(\nabla \boldsymbol{\tau}\left(\boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{\delta}(\boldsymbol{x})\right)\right)_{\boldsymbol{x}=\boldsymbol{\tau}^{-1}(\boldsymbol{z})} &= \boldsymbol{A}\boldsymbol{z}, \\ \left(\nabla \boldsymbol{\tau}\left(\boldsymbol{G}(\boldsymbol{x})\boldsymbol{\Gamma}(\boldsymbol{x})\right)\right)_{\boldsymbol{x}=\boldsymbol{\tau}^{-1}(\boldsymbol{z})} &= \boldsymbol{B}, \end{aligned} \tag{5.4}$$

where the pair $(\boldsymbol{A}, \boldsymbol{B})$ satisfies

$$\operatorname{rank}\left[\boldsymbol{B}, \boldsymbol{A}\boldsymbol{B}, \boldsymbol{A}^2\boldsymbol{B}, \ldots \boldsymbol{A}^{n-1}\boldsymbol{B}\right] = n. \tag{5.5}$$

Definition 6. *[?, p. 220] (**Vector relative degree**) Consider the Multi-Input Multi-Output (MIMO) system:*

$$\begin{aligned} \dot{\boldsymbol{x}} &= \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{u}, \\ \boldsymbol{y} &= \boldsymbol{h}(\boldsymbol{x}), \end{aligned} \tag{5.6}$$

where $\boldsymbol{x} \in \mathcal{X} \subseteq \mathbb{R}^n, \boldsymbol{u} \in \mathbb{R}^m$ as before. Furthermore, $\boldsymbol{h}(\boldsymbol{x}) := (h_1(\boldsymbol{x}), \dots h_m(\boldsymbol{x})) \in \mathbb{R}^m$, where h_j are smooth scalar-valued functions for all $j = 1 \dots, m$. The input-output system (5.6) is said to have vector relative degree *$\boldsymbol{\pi} = (\pi_1, \pi_2, \dots, \pi_m)$ at $\boldsymbol{x}_0 \in \mathcal{X}$, if*

$$L_{\boldsymbol{g}_j} L_{\boldsymbol{f}}^k h_i(\boldsymbol{x}) \equiv 0, \quad 1 \leq i, j \leq m, \quad 1 \leq k < \pi_i - 1, \tag{5.7}$$

and the $m \times m$ matrix

$$\boldsymbol{C}(\boldsymbol{x}) := \begin{bmatrix} L_{\boldsymbol{g}_1} L_{\boldsymbol{f}}^{\pi_1 - 1} h_1(\boldsymbol{x}) & \dots & L_{\boldsymbol{g}_m} L_{\boldsymbol{f}}^{\pi_1 - 1} h_1(\boldsymbol{x}) \\ \vdots & \ddots & \vdots \\ L_{\boldsymbol{g}_1} L_{\boldsymbol{f}}^{\pi_m - 1} h_m(\boldsymbol{x}) & \dots & L_{\boldsymbol{g}_m} L_{\boldsymbol{f}}^{\pi_m - 1} h_m(\boldsymbol{x}) \end{bmatrix} \tag{5.8}$$

evaluated at $\boldsymbol{x} = \boldsymbol{x}_0$, is non-singular.

Here, $\pi_j \in \mathbb{N}$, $j = 1, \dots, m$, is the number of times one has to differentiate the j^{th} output y_j w.r.t. t such that *at least* one of the m input components appears explicitly in the expression for $y_j^{(\pi_j)}$. In other words, π_j is the number of integrators between the input and the j^{th} output.

Remark 7. *It is known [?, p. 230] that given an n-dimensional vector field $\boldsymbol{f}$, and a matrix $\boldsymbol{G}(\boldsymbol{x})$ of rank m, the system (5.1) is full state static feedback linearizable if and only if:*

(i) there exist functions $h_1(\boldsymbol{x}), h_2(\boldsymbol{x}), \dots, h_m(\boldsymbol{x})$, such that the input-output system

(5.6) has relative degree $\boldsymbol{\pi}$ *at* $\boldsymbol{x}_0 \in \mathcal{X}$*,*

and

(ii) the relative degree $\boldsymbol{\pi}$ *is such that* $\pi_1 + \pi_2 + \cdots + \pi_m = n$*, where* n *is the dimension of the state vector* $\boldsymbol{x}$*.*

The output function $\boldsymbol{h}(\boldsymbol{x})$ play an important role in transforming (5.1) into a controllable linear system. If we can find $\boldsymbol{h}(\cdot)$ satisfying conditions (i)-(ii) in Remark 7, then we can use the same to construct a state feedback law and a desired change of coordinates. Explicitly, this feedback law can be obtained as

$$\boldsymbol{u} = \underbrace{-(\boldsymbol{C}(\boldsymbol{x}))^{-1}\boldsymbol{d}(\boldsymbol{x})}_{:=\boldsymbol{\delta}(\boldsymbol{x})} + \underbrace{(\boldsymbol{C}(\boldsymbol{x}))^{-1}}_{:=\boldsymbol{\Gamma}(\boldsymbol{x})}\boldsymbol{v}, \tag{5.9}$$

where $\boldsymbol{C}(\boldsymbol{x})$ is as in (5.8), and

$$\boldsymbol{d}(\boldsymbol{x}) := \left(L_{\boldsymbol{f}}^{\pi_1} h_1(\boldsymbol{x}),\ L_{\boldsymbol{f}}^{\pi_2} h_2(\boldsymbol{x}),\ \ldots,\ L_{\boldsymbol{f}}^{\pi_m} h_m(\boldsymbol{x})\right)^\top, \tag{5.10}$$

The linearizing coordinates $\boldsymbol{z} := \boldsymbol{\tau}(\boldsymbol{x})$ are subdivided as

$$\boldsymbol{z} = \begin{pmatrix} \boldsymbol{z}^1 \\ \boldsymbol{z}^2 \\ \vdots \\ \boldsymbol{z}^m \end{pmatrix}, \quad \boldsymbol{\tau}(\boldsymbol{x}) = \begin{pmatrix} \boldsymbol{\tau}^1(\boldsymbol{x}) \\ \boldsymbol{\tau}^2(\boldsymbol{x}) \\ \vdots \\ \boldsymbol{\tau}^m(\boldsymbol{x}) \end{pmatrix}, \tag{5.11}$$

where each $\boldsymbol{z}^i, \boldsymbol{\tau}^i \in \mathbb{R}^{\pi_i}$, $i = 1, \ldots, m$, have components

$$z_k^i = \tau_k^i(\boldsymbol{x}) := L_{\boldsymbol{f}}^{k-1} h_i(\boldsymbol{x}), \quad k = 1, \ldots, \pi_i. \tag{5.12}$$

The feedback law (5.9) and the change of coordinates (5.11), together transform

(5.1) which is in state-control pair $(\boldsymbol{x}, \boldsymbol{u})$, into the Brunovsky canonical form in the state-control pair $(\boldsymbol{z}, \boldsymbol{v})$, given by

$$\dot{\boldsymbol{z}} = \boldsymbol{A}\boldsymbol{z} + \boldsymbol{B}\boldsymbol{v}, \quad \boldsymbol{z} \in \mathbb{R}^n, \quad \boldsymbol{v} \in \mathbb{R}^m, \tag{5.13}$$

where $\boldsymbol{A}, \boldsymbol{B}$ are *block diagonal* matrices

$$\boldsymbol{A} := \operatorname{diag}(\boldsymbol{A}_1, \boldsymbol{A}_2, \ldots, \boldsymbol{A}_m), \quad \boldsymbol{B} := \operatorname{diag}(\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_m),$$

wherein for each $i = 1, \ldots, m$, we have

$$\boldsymbol{A}_i := [\boldsymbol{0}|\boldsymbol{e}_1|\boldsymbol{e}_2|\ldots|\boldsymbol{e}_{\pi_i - 1}] \in \mathbb{R}^{\pi_i \times \pi_i}, \quad \boldsymbol{b}_i := \boldsymbol{e}_{\pi_i} \in \mathbb{R}^{\pi_i}.$$

Remark 8. *Since static state feedback linearization is equivalent to Remark 7, hence the matrix (5.8) is invertible at $\boldsymbol{x} = \boldsymbol{x}_0$. This guarantees that $\boldsymbol{\Gamma}(\boldsymbol{x})$ and $\boldsymbol{\delta}(\boldsymbol{x})$ in (5.9) are well-defined at $\boldsymbol{x}_0$.*

As seen above, the existence of the (fictitious) output $\boldsymbol{h}(\cdot)$ is a necessary and sufficient condition for full state feedback linearization. The following result allows us to establish the existence of the $\boldsymbol{h}(\cdot)$ under suitable conditions on the vector fields $\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{g}_1(\boldsymbol{x}), \ldots, \boldsymbol{g}_m(\boldsymbol{x})$. Thus, the conditions for full state feedback linearization can be restated as the following.

Proposition 11. *[?, p. 232] Consider the system (5.1) where* $\operatorname{rank}(\boldsymbol{G}(\boldsymbol{x}_0)) = m$*, and for $i = 0, 1, \ldots, n-1$, let*

$$\Delta_i(\boldsymbol{x}) := \operatorname{span}\{\operatorname{ad}_{\boldsymbol{f}}^k \boldsymbol{g}_k : 0 \leq k \leq i, 1 \leq j \leq m\}.$$

Then, there exist scalar-valued functions $h_1(\boldsymbol{x}), \ldots, h_m(\boldsymbol{x})$, defined on $\mathcal{X}$ such that

(5.6) has relative degree $\boldsymbol{\pi}$ *at* $\boldsymbol{x}_0$*, with* $\pi_1 + \pi_2 + \cdots + \pi_m = n$*, iff:*

(i) Δ_i *has constant dimension near* $\boldsymbol{x}_0$ *for each* $i = 0, 1, \ldots, n-1$*,*

(ii) Δ_{n-1} *has dimension* n*,*

(iii) Δ_i *is involutive for each* $i = 0, 1, \ldots, n-2$*.*

Proposition 11 helps to verify if a given system of the form (5.1) is full-state feedback linearizable. For the construction of the functions h_i in Proposition 1, we refer the readers to [**?**].

Example 1. *Let us consider a system of the form (5.1) defined on a neighborhood of* $\boldsymbol{x}_0 = \mathbf{0}$*, given by*

$$\dot{\boldsymbol{x}} = \underbrace{\begin{pmatrix} x_2 + x_2^2 \\ x_3 - x_1 x_4 + x_4 x_5 \\ x_2 x_4 + x_1 x_5 - x_5^2 \\ x_5 \\ x_2^2 \end{pmatrix}}_{\boldsymbol{f}(\boldsymbol{x})} + \underbrace{\begin{pmatrix} 0 \\ 0 \\ \cos(x_1 - x_5) \\ 0 \\ 0 \end{pmatrix}}_{\boldsymbol{g}_1(\boldsymbol{x})} u_1 + \underbrace{\begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 1 \end{pmatrix}}_{\boldsymbol{g}_2(\boldsymbol{x})} u_2. \tag{5.14}$$

A direct computation verifies that (5.14) satisfies the conditions (i)-(iii) of Proposition 11, implying the existence of output functions $h_1(\boldsymbol{x}), h_2(\boldsymbol{x})$*. Following the*

*constructive steps in [**?**, p. 232], these output functions can be obtained as*

$$y_1 = h_1 = x_1 - x_5, y_2 = h_2 = x_4. \tag{5.15}$$

Notice that

$$\begin{aligned} &L_{\boldsymbol{g}_1} h_1(\boldsymbol{x}) \equiv L_{\boldsymbol{g}_2} h_1(\boldsymbol{x}) \equiv L_{\boldsymbol{g}_1} L_{\boldsymbol{f}} h_1(\boldsymbol{x}) \equiv L_{\boldsymbol{g}_2} L_{\boldsymbol{f}} h_1(\boldsymbol{x}) \equiv 0, \\ &L_{\boldsymbol{g}_1} h_2(\boldsymbol{x}) \equiv L_{\boldsymbol{g}_2} h_2(\boldsymbol{x}) \equiv 0, \end{aligned}$$

and that the matrix

$$\begin{aligned} \boldsymbol{C}(\boldsymbol{x}_0) &= \begin{pmatrix} L_{\boldsymbol{g}_1} L_{\boldsymbol{f}}^2 h_1(\boldsymbol{x}) & L_{\boldsymbol{g}_2} L_{\boldsymbol{f}}^2 h_1(\boldsymbol{x}) \\ L_{\boldsymbol{g}_1} L_{\boldsymbol{f}} h_2(\boldsymbol{x}) & L_{\boldsymbol{g}_2} L_{\boldsymbol{f}} h_2(\boldsymbol{x}) \end{pmatrix} \Big|_{\boldsymbol{x}=\boldsymbol{x}_0} \\ &= \begin{pmatrix} \cos(x_1 - x_5) & 1 \\ 0 & 1 \end{pmatrix} \Big|_{\boldsymbol{x}=\boldsymbol{x}_0} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \end{aligned}$$

is non-singular. Therefore, (5.14) with output (5.15) has vector relative degree $\boldsymbol{\pi} = (\pi_1, \pi_2) = (3, 2)$ satisfying $\pi_1 + \pi_2 = 3 + 2 = 5$, which is indeed the dimension of the state space. From (5.12), we obtain the change of coordinates

$$\boldsymbol{\tau}(\boldsymbol{x}) := \begin{pmatrix} h_1(\boldsymbol{x}) \\ L_{\boldsymbol{f}} h_1(\boldsymbol{x}) \\ L_{\boldsymbol{f}}^2 h_1(\boldsymbol{x}) \\ h_2(\boldsymbol{x}) \\ L_{\boldsymbol{f}} h_2(\boldsymbol{x}) \end{pmatrix} = \begin{pmatrix} x_1 - x_5 \\ x_2 \\ x_3 - x_1 x_4 + x_4 x_5 \\ x_4 \\ x_5 \end{pmatrix}. \tag{5.16}$$

In this case, (5.8) and (5.10) yields

$$\boldsymbol{d}(\boldsymbol{x}) = \begin{pmatrix} L_{\boldsymbol{f}}^3 h_1(\boldsymbol{x}) \\ L_{\boldsymbol{f}}^2 h_2(\boldsymbol{x}) \end{pmatrix} = \begin{pmatrix} 0 \\ x_2^2 \end{pmatrix},$$

$$(\boldsymbol{C}(\boldsymbol{x}))^{-1} = \begin{pmatrix} 1/\cos(x_1 - x_5) & -1/\cos(x_1 - x_5) \\ 0 & 1 \end{pmatrix},$$

which, following (5.9), result in the feedback law

$$\begin{aligned} \boldsymbol{u} = & \underbrace{\begin{pmatrix} -x_2^2/\cos(x_1 - x_5) \\ x_2^2 \end{pmatrix}}_{\boldsymbol{\delta}(\boldsymbol{x})} \\ & + \underbrace{\begin{pmatrix} 1/\cos(x_1 - x_5) & -1\cos(x_1 - x_5) \\ 0 & 1 \end{pmatrix}}_{\boldsymbol{\Gamma}(\boldsymbol{x})} \boldsymbol{v}. \end{aligned} \tag{5.17}$$

Hence, we have constructed a triple $(\boldsymbol{\delta}(\boldsymbol{x}), \boldsymbol{\Gamma}(\boldsymbol{x}), \boldsymbol{\tau}(\boldsymbol{x}))$ given by (5.16)-(5.17), that transform (5.14) into

$$\dot{\boldsymbol{z}} = \underbrace{\begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}}_{\boldsymbol{A}} \boldsymbol{z} + \underbrace{\begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}}_{\boldsymbol{B}} \boldsymbol{v}. \tag{5.18}$$

5.3 Minimum Energy Density Control

5.3.1 Stochastic Optimal Control Problem

Given system (5.1), and two prescribed endpoint PDFs $\rho_0(\boldsymbol{x}), \rho_1(\boldsymbol{x})$, we consider the following minimum energy finite horizon stochastic optimal control problem:

$$\inf_{\boldsymbol{u}\in U} \quad \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|\boldsymbol{u}(x,t)\|_2^2 \,\mathrm{d}t\right\}, \tag{5.19a}$$

$$\text{subject to} \quad \dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{u}, \tag{5.19b}$$

$$\boldsymbol{x}(0) \sim \rho_0(\boldsymbol{x}) \quad \boldsymbol{x}(1) \sim \rho_1(\boldsymbol{x}), \tag{5.19c}$$

where the state space is $\mathcal{X} \subseteq \mathbb{R}^n$, $\boldsymbol{u} \in \mathbb{R}^m$ and (5.19b) is *feedback linearizable*. The infimum is taken over the set of admissible controls with finite energy, i.e., $\mathcal{U} := \{\boldsymbol{u} : \mathbb{R}^n \times [0,1] \mapsto \mathbb{R}^m | \ \|\boldsymbol{u}\|_2^2 < \infty\}$, and the expectation operator $\mathbb{E}\{\cdot\}$ in (5.19a) is w.r.t. the controlled joint state PDF $\rho(\boldsymbol{x}, t)$ satisfying endpoint conditions (5.19c). The objective is to steer the joint PDF $\rho(\boldsymbol{x}, t)$ from the given initial PDF ρ_0 at $t = 0$ to a terminal PDF ρ_1 at $t = 1$ while minimizing the expected control effort.

The problem (5.19) can be recast into a "fluid dynamics" version [21], which is the following variational problem:

$$\inf_{\rho,\boldsymbol{u}} \quad \int_0^1 \int_{\mathcal{X}} \frac{1}{2}\|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2 \,\rho(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}t, \tag{5.20a}$$

$$\text{subject to} \quad \frac{\partial \rho}{\partial t} + \nabla_{\boldsymbol{x}} \cdot (\rho(\boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{u})) = 0, \tag{5.20b}$$

$$\rho(\boldsymbol{x}, t=0) = \rho_0, \quad \rho(\boldsymbol{x}, t=1) = \rho_1. \tag{5.20c}$$

Here, the infimum is taken over $\mathcal{P}(\mathcal{X}) \times \mathcal{U}$, where $\mathcal{P}(\mathcal{X})$ denotes the space of

all joint PDFs supported on $\mathcal{X}$. We note that (5.20b) is the Liouville PDE [34] associated to the dynamical system (5.19b).

5.3.2 Reformulation in Feedback Linearized Coordinates

In our recent work [38], we considered the problem (5.19) for the single-input case, i.e., the case $\boldsymbol{G}(\boldsymbol{x}) \equiv [\boldsymbol{g}_1(\boldsymbol{x})] \in \mathbb{R}^n$, and the input u is scalar-valued. The main idea in [38] was to recast (5.20), which is in state-control pair $(\boldsymbol{x}, \boldsymbol{u})$, into an equivalent formulation in feedback linearized state-control pair $(\boldsymbol{z}, \boldsymbol{v})$. This was made possible by using the diffeomorphism $\boldsymbol{\tau} : \mathcal{X} \mapsto \mathcal{Z}$ to pushforward the endpoint PDFs ρ_0, ρ_1 to PDFs σ_0, σ_1 supported on the feedback linearized state space $\mathcal{Z}$. Specifically,

$$\sigma_i(\boldsymbol{z}) := \boldsymbol{\tau}_\sharp \rho_i = \frac{\rho_i(\boldsymbol{\tau}^{-1}(\boldsymbol{z}))}{|\det(\nabla_{\boldsymbol{x}} \boldsymbol{\tau}_{\boldsymbol{x}=\tau^{-1}(\boldsymbol{z})}|}, \quad i \in \{0, 1\}, \tag{5.21}$$

and $\mathcal{Z} := \{\boldsymbol{z} \in \mathbb{R}^n | \boldsymbol{z} = \tau(\boldsymbol{x}), \boldsymbol{x} \in \mathcal{X}\}$.

Since $\boldsymbol{\tau}$ is a diffeomorphism, the PDFs $\{\sigma_i\}_{i=0,1}$ supported on the feedback linearized state space $\mathcal{Z}$, are well defined, i.e., $\text{spt}(\sigma_i) \subseteq \mathcal{Z}$ provided that $\text{spt}(\rho_i) \subseteq \mathcal{X}$.

To generalize the reformulation in [38, Sec. III.B] for the multi-input case, we proceed by setting

$$\boldsymbol{\delta}_{\boldsymbol{\tau}} := \boldsymbol{\delta} \circ \boldsymbol{\tau}^{-1}, \quad \boldsymbol{\Gamma}_{\boldsymbol{\tau}} := \boldsymbol{\Gamma} \circ \boldsymbol{\tau}^{-1}, \tag{5.22}$$

where $\boldsymbol{\delta}$ and $\boldsymbol{\Gamma}$ are as in (5.9). Using $\boldsymbol{u}(\boldsymbol{z}) = \boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) + \boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z})\boldsymbol{v}$, we now transcribe

(5.20) into

$$\inf_{\sigma,\boldsymbol{v}} \quad \int_0^1 \int_{\mathcal{Z}} \frac{1}{2}\mathcal{L}(\boldsymbol{z},\boldsymbol{v})\,\sigma(\boldsymbol{z},t)\,\mathrm{d}\boldsymbol{z}\,\mathrm{d}t, \tag{5.23a}$$

$$\text{subject to} \quad \frac{\partial\sigma}{\partial t} + \nabla_{\boldsymbol{z}}\cdot((\boldsymbol{A}\boldsymbol{z}+\boldsymbol{B}\boldsymbol{v})\sigma) = 0, \tag{5.23b}$$

$$\sigma(\boldsymbol{x},t=0) = \sigma_0, \quad \sigma(\boldsymbol{x},t=1) = \sigma_1, \tag{5.23c}$$

where

$$\mathcal{L}(\boldsymbol{z},\boldsymbol{v}) := \|\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) + \boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z})\boldsymbol{v}\|_2^2. \tag{5.24}$$

The infimum in (5.23) is taken over the pair of transformed PDFs and admissible controls $(\sigma,\boldsymbol{v}) \in \mathcal{P}(\mathcal{Z}) \times \mathcal{V}$ where $\mathcal{V} := \{v : \mathcal{Z}\times[0,1] \mapsto \mathbb{R}^m |\ \|v\|_2^2 < \infty\}$.

Remark 9. *The solution pair $(\rho^{\mathrm{opt}}, \boldsymbol{u}^{\mathrm{opt}})$ for (5.20) can be recovered from the optimal solution $(\sigma^{\mathrm{opt}}, \boldsymbol{v}^{\mathrm{opt}})$ of (5.23) via the transformations*

$$\rho^{\mathrm{opt}}(\boldsymbol{x},t) = \sigma^{\mathrm{opt}}(\boldsymbol{\tau}(\boldsymbol{x}),t)|\mathrm{det}\nabla_{\boldsymbol{x}}\boldsymbol{\tau}_{\boldsymbol{x}}(\boldsymbol{x})|, \tag{5.25a}$$

$$\boldsymbol{u}^{\mathrm{opt}}(\boldsymbol{x},t) = \boldsymbol{\delta}(\boldsymbol{x}) + \boldsymbol{\Gamma}(\boldsymbol{x})\boldsymbol{v}^{\mathrm{opt}}(\boldsymbol{\tau}^{-1}(\boldsymbol{x}),t) \tag{5.25b}$$

for $\boldsymbol{x} \in \mathcal{X}$, and $t \in [0,1]$.

Example 2. *To illustrate the reformulation (5.23), let us reconsider the system (5.14). In this case, the inverse mapping of (5.16) is given by*

$$\boldsymbol{x} = \boldsymbol{\tau}^{-1}(\boldsymbol{z}) := \begin{pmatrix} z_1 + z_5 & z_2 & z_3 + z_1 z_4 & z_4 & z_5 \end{pmatrix}^\top. \tag{5.26}$$

Here, the determinant of the Jacobian of (5.16) is non-zero for all vectors in $\mathbb{R}^n$,

i.e., $\mathcal{Z} = \mathbb{R}^n$. *From (5.17) and (5.22), we have*

$$\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) = \begin{pmatrix} -x_2^2/\cos(z_1) \\ z_2^2 \end{pmatrix}, \tag{5.27}$$

and

$$\boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z}) = \begin{pmatrix} 1/\cos(z_1) & -1/\cos(z_1) \\ 0 & 1 \end{pmatrix}. \tag{5.28}$$

The functional $\mathcal{L}(\boldsymbol{z}, \boldsymbol{v})$ *in (5.24) equals*

$$\begin{aligned} \boldsymbol{v}^\top \begin{pmatrix} 2/\cos^2(z_1) & -1/\cos(z_1) \\ -1/\cos(z_1) & 1 \end{pmatrix} \boldsymbol{v} + \langle \Big(z_2^2/\cos^2(z_1), \\ -z_2^2/\cos(z_1) + z_2^2 \Big)^\top, \boldsymbol{v} \rangle + x_2^4/\cos^2(z_1) + z_2^4. \end{aligned} \tag{5.29}$$

Remark 10. *Because feedback linearization guarantees that the matrix pair* $(\boldsymbol{A}, \boldsymbol{B})$ *is controllable, any vector* $\boldsymbol{z}_1 \in \mathcal{Z}$ *is reachable from any other vector* $\boldsymbol{z}_0 \in \mathcal{Z}$ *for all* $t \in [0, 1]$ *via the flow of (5.13). This ensures that in (5.23c), the initial PDF* $\sigma_0(\boldsymbol{z})$ *can be steered to* $\sigma_1(\boldsymbol{z})$ *via the flow* $\sigma(\boldsymbol{z}, t)$ *of the controlled Liouville PDE (5.23b). Thus, the constraint set of (5.23) is non-empty, and the problem is feasible.*

5.3.3 Optimality

To show the existence and uniqueness of minimizer for (5.23), we set $\boldsymbol{m} := \sigma \boldsymbol{v}$, and consider the change of variable $(\sigma, \boldsymbol{v}) \mapsto (\sigma, \boldsymbol{m})$, transforming (5.23) into

$$\begin{aligned} &\inf_{\sigma, \boldsymbol{m}} && \int_0^1 \int_{\mathcal{Z}} \mathcal{J}(\sigma, \boldsymbol{m}) \, \mathrm{d}\boldsymbol{z} \, \mathrm{d}t, && \text{(5.30a)} \\ &\text{subject to} && \frac{\partial \sigma}{\partial t} + \nabla_{\boldsymbol{z}} \cdot (\boldsymbol{A}\boldsymbol{z}\sigma + \boldsymbol{B}\boldsymbol{m}) = 0, && \text{(5.30b)} \\ & && \sigma(\boldsymbol{x}, t=0) = \sigma_0, \quad \sigma(\boldsymbol{x}, t=1) = \sigma_1, && \text{(5.30c)} \end{aligned}$$

where

$$\mathcal{J}(\sigma, \boldsymbol{m}) := \begin{cases} \frac{1}{2} \| \boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) + \boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z}) \frac{\boldsymbol{m}}{\sigma} \|_2^2 \, \sigma & \text{if } \sigma > 0, \\ 0 & \text{if } (\sigma, \boldsymbol{m}) = (0,0), \\ +\infty & \text{otherwise.} \end{cases} \tag{5.31}$$

We note that $\mathcal{J}(\sigma, \boldsymbol{m})$ is the perspective function of the strictly convex map $\boldsymbol{m} \mapsto \| \boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z})\sigma + \boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z})\boldsymbol{m} \|_2^2$; therefore, $\mathcal{J}$ is jointly strictly convex in $(\sigma, \boldsymbol{m})$. The constraints (5.30b)-(5.30c) are linear in $(\sigma, \boldsymbol{m})$. Hence, (5.30) admits a unique minimizing pair, and equivalently, so does (5.23). The following theorem summarizes how this optimal pair for (5.23), denoted hereafter as $(\sigma^{\mathrm{opt}}, \boldsymbol{v}^{\mathrm{opt}})$, can be obtained.

Theorem 12. *(**Optimal control for (5.23)**) The optimal control $\boldsymbol{v}^{\mathrm{opt}}$ for the problem (5.23), is given by*

$$\boldsymbol{v}^{\mathrm{opt}}(\boldsymbol{z}, t) = (\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{\top} \boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z}))^{-1} \boldsymbol{B}^{\top} \nabla_{\boldsymbol{z}} \psi - \boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1}(\boldsymbol{z}) \boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}), \tag{5.32}$$

where ψ solves the Hamilton-Jacobi-Bellman (HJB) PDE

$$\begin{aligned}&\frac{\partial \psi}{\partial t}+\langle \nabla_{\boldsymbol{z}}\psi, \boldsymbol{A}\boldsymbol{z}\rangle-\langle \nabla_{\boldsymbol{z}}\psi, \boldsymbol{B}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1}(\boldsymbol{z})\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z})\rangle\\&+\frac{1}{2}\langle \nabla_{\boldsymbol{z}}\psi, \boldsymbol{B}\left(\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{\top}(\boldsymbol{z})\boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z})\right)^{-1}\boldsymbol{B}^{\top}\nabla_{\boldsymbol{z}}\psi\rangle=0.\end{aligned}\tag{5.33}$$

Furthermore, if the optimal joint state PDF σ^{opt} is a solution to the Liouville PDE

$$\frac{\partial \sigma^{\text{opt}}}{\partial t}+\nabla_{\boldsymbol{z}}\cdot\left(\left(\boldsymbol{A}\boldsymbol{z}+\boldsymbol{B}\boldsymbol{v}^{\text{opt}}\right)\right)=0,\tag{5.34}$$

with boundary conditions $\sigma^{\text{opt}}(\boldsymbol{z},0)=\sigma_0^{\text{opt}}(\boldsymbol{z})$, and $\sigma^{\text{opt}}(\boldsymbol{z},1)=\sigma_1^{\text{opt}}(\boldsymbol{z})$, then the pair $(\sigma^{\text{opt}},\boldsymbol{v}^{\text{opt}})$ solves the problem (5.23).

Proof. The Lagrangian associated with (5.23) is

$$\begin{aligned}\mathscr{L}(\sigma,\psi,\boldsymbol{v})=&\int_0^1\int_{\mathcal{Z}}\frac{1}{2}\mathcal{L}(\boldsymbol{z},\boldsymbol{v})\sigma(\boldsymbol{z},t)\,\mathrm{d}\boldsymbol{z}\mathrm{d}t\\&+\underbrace{\int_{\mathcal{Z}}\int_0^1\psi(\boldsymbol{z},t)\frac{\partial\sigma}{\partial t}\mathrm{d}t\,\mathrm{d}\boldsymbol{z}}_{\text{term 1}}\\&+\underbrace{\int_0^1\int_{\mathcal{Z}}\psi(\boldsymbol{z},t)\nabla_{\boldsymbol{z}}\cdot\left(\left(\boldsymbol{A}\boldsymbol{z}+\boldsymbol{B}\boldsymbol{v}\right)\sigma\right)\,\mathrm{d}\boldsymbol{z}\mathrm{d}t}_{\text{term 2}}.\end{aligned}\tag{5.35}$$

In (5.35), we interchange the order of integration and perform integration by parts w.r.t. t in term 1, and w.r.t. $\boldsymbol{z}$ in term 2. Since $\sigma(\boldsymbol{z},t)\to 0$ as $z\to\partial\mathcal{Z}$, we can express $\mathscr{L}$ as

$$\int_0^1\int_{\mathcal{Z}}\left\{\frac{1}{2}\mathcal{L}(\boldsymbol{z},\boldsymbol{v})-\frac{\partial\psi}{\partial t}-\langle\nabla_{\boldsymbol{z}}\psi,\boldsymbol{A}\boldsymbol{z}+\boldsymbol{B}\boldsymbol{v}\rangle\right\}\sigma(\boldsymbol{z},t)\mathrm{d}\boldsymbol{z}\mathrm{d}t.\tag{5.36}$$

Performing pointwise minimization of the above w.r.t. $\boldsymbol{v}$ while fixing σ, we obtain

$$\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{\top}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}\boldsymbol{v}^{\mathrm{opt}}(\boldsymbol{z},t) = \boldsymbol{B}^{\top}\nabla_{z}\psi(\boldsymbol{z},t) - \boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{\top}(\boldsymbol{z})\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}). \tag{5.37}$$

Taking the matrix inverse on both sides yield (5.32). Substituting $\boldsymbol{v}^{\mathrm{opt}}$ back into (5.36) and equating to zero, we then get

$$\begin{aligned}&\int_0^1\int_{\mathcal{Z}}\Big\{-\frac{\partial\psi}{\partial t} - \langle\nabla_{z}\psi, \boldsymbol{A}\boldsymbol{z}\rangle + \langle\nabla_{z}\psi, \boldsymbol{B}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1}(\boldsymbol{z})\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z})\rangle\\ &-\frac{1}{2}\langle\nabla_{z}\psi, \boldsymbol{B}\left(\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{\top}(\boldsymbol{z})\boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z})\right)^{-1}\boldsymbol{B}^{\top}\nabla_{z}\psi\rangle\Big\}\sigma(\boldsymbol{z},t)\mathrm{d}\boldsymbol{z}\mathrm{d}t = 0.\end{aligned} \tag{5.38}$$

Since (5.38) holds for arbitrary σ, we arrive at (5.33). ■

Example 3. *(**HJB for (5.29)**) From (5.27)-(5.28), we have*

$$(\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{\top}(\boldsymbol{z})\boldsymbol{\Gamma}_{\boldsymbol{\tau}}(\boldsymbol{z}))^{-1} = \begin{pmatrix}\cos^2(z_1) & \cos(z_1)\\ \cos(z_1) & 2\end{pmatrix} \tag{5.39}$$

and

$$\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1}(\boldsymbol{z})\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) = \begin{pmatrix}0\\ z_2^2\end{pmatrix}. \tag{5.40}$$

Substituting (5.39)-(5.40) into (5.33), and using the pair $(\boldsymbol{A},\boldsymbol{B})$ from (5.18), gives the HJB PDE

$$\begin{aligned}&\frac{\partial\psi}{\partial t} + z_2\frac{\partial\psi}{\partial z_1} + z_3\frac{\partial\psi}{\partial z_2} + z_5\frac{\partial\psi}{\partial z_4} - z_2^2\frac{\partial\psi}{\partial z_5}\\ &+\frac{1}{2}\Big[\cos^2(z_1)\left(\frac{\partial\psi}{\partial z_3}\right)^2 + \cos(z_1)\frac{\partial\psi}{\partial z_3}\frac{\partial\psi}{\partial z_5} + 2\left(\frac{\partial\psi}{\partial z_5}\right)^2\Big] = 0.\end{aligned} \tag{5.41}$$

Remark 11. *Computing the pair $(\sigma^{\mathrm{opt}}, \boldsymbol{v}^{\mathrm{opt}})$ in Theorem 12 is challenging in general since it calls for solving a system of coupled nonlinear PDEs (5.33)-(5.34) with atypical boundary conditions. In the following Sections, we will provide fur-*

ther reformulations of (5.23) to make it computationally amenable.

5.4 Stochastic Density Steering: Reformulation into Schrödinger System

Motivated by [**?**], we consider a generalized version of (5.23) by adding a diffusion term to (5.23b):

$$\begin{aligned}
\inf_{\sigma,\boldsymbol{v}} \quad & \int_0^1 \int_{\mathcal{Z}} \frac{1}{2}\mathcal{L}(\boldsymbol{z},\boldsymbol{v})\sigma(\boldsymbol{z},t)\,\mathrm{d}\boldsymbol{z}\,\mathrm{d}t, && \text{(5.42a)}\\
\text{subject to} \quad & \frac{\partial\sigma}{\partial t} + \nabla_{\boldsymbol{z}}\cdot((\boldsymbol{A}\boldsymbol{z}+\boldsymbol{B}\boldsymbol{v})\sigma) \\
& \qquad = \epsilon \mathbf{1}^\top \left(\boldsymbol{D}(\boldsymbol{z}) \odot \mathrm{Hess}(\sigma)\right)\mathbf{1}, && \text{(5.42b)}\\
& \sigma(\boldsymbol{z}, t=0) = \sigma_0, \quad \sigma(\boldsymbol{z}, t=1) = \sigma_1, && \text{(5.42c)}
\end{aligned}$$

where $\boldsymbol{D}(\boldsymbol{z}) := \boldsymbol{B}\boldsymbol{\Gamma}_\tau^{-1}(\boldsymbol{z})(\boldsymbol{B}\boldsymbol{\Gamma}_\tau^{-1}(\boldsymbol{z}))^\top$. In particular, the controlled Liouville PDE in (5.23b) is now replaced by a Fokker-Planck-Kolmogorov PDE in (5.42b), having an additional diffusion term $\sqrt{2\epsilon}\boldsymbol{B}\boldsymbol{\Gamma}_\tau^{-1}(\boldsymbol{z})$, where the parameter $\epsilon > 0$ (not necessarily small). Formally, this generalization is equivalent to adding a stochastic perturbation to the controlled sample path ODE $\dot{\boldsymbol{z}} = \boldsymbol{A}\boldsymbol{z} + \boldsymbol{B}\boldsymbol{v}$, resulting in the Itô SDE

$$\mathrm{d}\boldsymbol{z} = (\boldsymbol{A}\boldsymbol{z}+\boldsymbol{B}\boldsymbol{v})\,\mathrm{d}t + \sqrt{2\epsilon}\boldsymbol{B}\boldsymbol{\Gamma}_\tau^{-1}(\boldsymbol{z})\,\mathrm{d}\boldsymbol{w}, \tag{5.43}$$

where $\boldsymbol{w}(t) \in \mathbb{R}^m$ is standard Wiener process. In the special case $\boldsymbol{\delta}_\tau(\boldsymbol{z}) \equiv 0$ and $\boldsymbol{\Gamma}_\tau(\boldsymbol{z}) \equiv \boldsymbol{I}$, problem (5.42) reduces to the Schrödinger bridge problem with linear prior dynamics [**?**, equation (49)]. Thus, (5.42) is a Schrödinger bridge-like problem with a prior dynamics that has linear drift and nonlinear diffusion

coefficient.

The following Theorem characterizes the minimizing pair $(\sigma^{\text{opt}}, \boldsymbol{v}^{\text{opt}})$ for problem (5.42).

Theorem 13. *(**Optimal control for (5.42)**) The optimal control $\boldsymbol{v}^{\text{opt}}(\boldsymbol{z}, t)$ for (5.42) is given by (5.32), where ψ solves the HJB PDE*

$$\begin{aligned}&\frac{\partial \psi}{\partial t} + \langle \nabla_{\boldsymbol{z}} \psi, \boldsymbol{A}\boldsymbol{z} \rangle - \langle \nabla_{\boldsymbol{z}} \psi, \boldsymbol{B}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1}(\boldsymbol{z})\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) \rangle \\ &+ \frac{1}{2}\langle \nabla_{\boldsymbol{z}} \psi, \boldsymbol{D}(\boldsymbol{z}) \nabla_{\boldsymbol{z}} \psi \rangle + \epsilon \langle \boldsymbol{D}(\boldsymbol{z}), \text{Hess}(\psi) \rangle = 0, \end{aligned} \tag{5.44}$$

and the optimal joint state PDF $\sigma^{\text{opt}}(\boldsymbol{z}, t)$ solves the controlled Fokker-Planck-Kolmogorov PDE

$$\begin{aligned}&\frac{\partial \sigma^{\text{opt}}}{\partial t} + \nabla_{\boldsymbol{z}} \cdot ((\boldsymbol{A}\boldsymbol{z} + \boldsymbol{B}\boldsymbol{v}^{\text{opt}})\sigma^{\text{opt}}) \\ &\qquad - \epsilon \mathbf{1}^{\top} \left(\boldsymbol{D}(\boldsymbol{z}) \odot \text{Hess}(\sigma^{\text{opt}}) \right) \mathbf{1} = 0, \end{aligned} \tag{5.45}$$

with boundary conditions

$$\sigma^{\text{opt}}(\boldsymbol{z}, t = 0) = \sigma_0^{\text{opt}}, \quad \sigma^{\text{opt}}(\boldsymbol{z}, t = 1) = \sigma_1^{\text{opt}}. \tag{5.46}$$

Proof. The proof proceeds similarly as in Theorem 12 except that we now have an additional term in the Lagrangian (5.35) which we refer to as "term 3", given by

$$\underbrace{- \epsilon \int_0^1 \int_{\mathcal{Z}} \psi(\boldsymbol{z}, t) \mathbf{1}^{\top} \left(\boldsymbol{D}(\boldsymbol{z}) \odot \text{Hess}(\sigma^{\text{opt}}) \right) \mathbf{1} \mathrm{d}\boldsymbol{z} \mathrm{d}t.}_{\text{term 3}} \tag{5.47}$$

From the following chain of equalities:

$$\begin{aligned}
&\int_{\mathcal{Z}} \langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(\psi)\rangle \sigma^{\mathrm{opt}}(\boldsymbol{z},t)\mathrm{d}\boldsymbol{z} \\
&= \int_{\mathcal{Z}} \sum_{i,j=1}^{n} \boldsymbol{D}_{ij}(\boldsymbol{z}) \frac{\partial \psi(\boldsymbol{z},t)}{\partial z_i \partial z_j} \sigma^{\mathrm{opt}}(\boldsymbol{z},t)\mathrm{d}\boldsymbol{z} \\
&= \sum_{i,j=1}^{n} \int_{\mathcal{Z}} \boldsymbol{D}_{ij}(\boldsymbol{z}) \frac{\partial \psi(\boldsymbol{z},t)}{\partial z_i \partial z_j} \sigma^{\mathrm{opt}}(\boldsymbol{z},t)\mathrm{d}\boldsymbol{z} \\
&= -\sum_{i,j=1}^{n} \int_{\mathcal{Z}} \frac{\partial \psi(\boldsymbol{z},t)}{\partial z_j} \frac{\partial (\boldsymbol{D}_{ij}\sigma^{\mathrm{opt}}(\boldsymbol{z},t))}{\partial z_i}\mathrm{d}\boldsymbol{z} \\
&= \int_{\mathcal{Z}} \psi(\boldsymbol{z},t) \sum_{i,j=1}^{n} \frac{\partial (\boldsymbol{D}(\boldsymbol{z})_{ij}\sigma^{\mathrm{opt}}(\boldsymbol{z},t))}{\partial z_j \partial z_i}\mathrm{d}\boldsymbol{z} \\
&= \int_{\mathcal{Z}} \psi(\boldsymbol{z},t) \mathbf{1}^{\top} \Big(\boldsymbol{D}(\boldsymbol{z}) \odot \mathrm{Hess}(\sigma^{\mathrm{opt}}) \Big) \mathbf{1}\mathrm{d}\boldsymbol{z} \qquad (5.48)
\end{aligned}$$

we deduce that (5.47) is equal to

$$-\epsilon \int_0^1 \int_{\mathcal{Z}} \langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(\psi)\rangle \sigma^{\mathrm{opt}}(\boldsymbol{z},t)\mathrm{d}\boldsymbol{z}\mathrm{d}t. \qquad (5.49)$$

So, the expression inside the curly braces in (5.36), now will have an additional term $-\epsilon\langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(\psi)\rangle$ that is independent of $\boldsymbol{v}$. Therefore, pointwise minimization of (5.36) with this additional term w.r.t. $\boldsymbol{v}$, gives (5.32), and the associated HJB PDE becomes (5.44). ■

Next, we show that the so-called Hopf-Cole transform [68, 114] allows to reduce the system of *nonlinear* PDEs (5.44)-(5.45) with boundary conditions (5.46), into a system of boundary-coupled *linear* PDEs, which we refer as the "Schrödinger System".

Theorem 14. ***(Schrödinger System)*** *Consider the Hopf-Cole transformation*

$(\sigma^{\mathrm{opt}}, \psi) \mapsto (\varphi, \hat{\varphi})$*:*

$$\varphi(\boldsymbol{z}, t) := \exp(\psi(\boldsymbol{z}, t)/2\epsilon), \tag{5.50a}$$

$$\hat{\varphi}(\boldsymbol{z}, t) := \sigma^{\mathrm{opt}}(\boldsymbol{z}, t) \exp(-\psi(\boldsymbol{z}, t)/2\epsilon), \tag{5.50b}$$

applied to the system of coupled nonlinear PDEs (5.44)-(5.45). The pair $(\varphi, \hat{\varphi})$ *satisfies the following system of linear PDEs:*

$$\frac{\partial \varphi}{\partial t} + \langle \nabla_{\boldsymbol{z}} \varphi, \boldsymbol{A}\boldsymbol{z} - \boldsymbol{B}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1} \boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) \rangle + \epsilon \langle \boldsymbol{D}, \mathrm{Hess}(\varphi) \rangle = 0, \tag{5.51a}$$

$$\begin{aligned} \frac{\partial \hat{\varphi}}{\partial t} + \nabla_{\boldsymbol{z}} \cdot \Big(\Big(\boldsymbol{A}\boldsymbol{z} - \boldsymbol{B}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1} \boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z}) \Big) \hat{\varphi} \Big) \\ - \epsilon \mathbf{1} \left(\boldsymbol{D}(\boldsymbol{z}) \odot \mathrm{Hess}(\hat{\varphi}) \right) \mathbf{1} = 0, \end{aligned} \tag{5.51b}$$

with coupled boundary conditions

$$\varphi_0(\boldsymbol{z}) \hat{\varphi}_0(\boldsymbol{z}) = \sigma_0^{\mathrm{opt}}(\boldsymbol{z}), \quad \varphi_1(\boldsymbol{z}) \hat{\varphi}_1(\boldsymbol{z}) = \sigma_1^{\mathrm{opt}}(\boldsymbol{z}). \tag{5.52}$$

Proof. From (5.50a), $\psi = 2\epsilon \log \varphi$, which yields

$$\frac{\partial \psi}{\partial t} = \frac{2\epsilon}{\varphi} \frac{\partial \varphi}{\partial t}, \quad \nabla_{\boldsymbol{z}} \psi = \frac{2\epsilon}{\varphi} \nabla_{\boldsymbol{z}} \varphi. \tag{5.53}$$

On the other hand, notice that

$$
\begin{aligned}
&\epsilon\langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(\psi)\rangle \\
&= \epsilon\langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(2\epsilon\log\varphi)\rangle \\
&= 2\epsilon^2 \sum_{i,j=1}^{n} \boldsymbol{D}_{ij}(\boldsymbol{z}) \frac{\partial^2 \log\varphi}{\partial z_i \partial z_j} \\
&= 2\epsilon^2 \sum_{i,j=1}^{n} \boldsymbol{D}_{ij}(\boldsymbol{z}) \left(\frac{1}{\varphi}\frac{\partial^2\varphi}{\partial z_i z_j} - \frac{1}{\varphi^2}\frac{\partial\varphi}{\partial z_i}\frac{\partial\varphi}{\partial z_j} \right) \\
&= \frac{2\epsilon^2}{\varphi}\langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(\varphi)\rangle - \frac{2\epsilon^2}{\varphi^2}\langle \nabla_{\boldsymbol{z}}\varphi, \boldsymbol{D}(\boldsymbol{z})\nabla_{\boldsymbol{z}}\varphi\rangle. \qquad (5.54)
\end{aligned}
$$

Substituting (5.53) and (5.54) into (5.44) yields

$$
\begin{aligned}
&\frac{2\epsilon}{\varphi}\frac{\partial\varphi}{\partial t} + \frac{2\epsilon}{\varphi}\langle \nabla_{\boldsymbol{z}}\varphi, \boldsymbol{A}\boldsymbol{z} - \boldsymbol{B}\boldsymbol{\Gamma}_{\boldsymbol{\tau}}^{-1}(\boldsymbol{z})\boldsymbol{\delta}_{\boldsymbol{\tau}}(\boldsymbol{z})\rangle \\
&+ \frac{1}{2}\frac{4\epsilon^2}{\varphi^2}\langle \nabla_{\boldsymbol{z}}\varphi, \boldsymbol{D}(\boldsymbol{z})\nabla_{\boldsymbol{z}}\varphi\rangle + \frac{2\epsilon^2}{\varphi}\langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(\varphi)\rangle \\
&- \frac{2\epsilon^2}{\varphi^2}\langle \nabla_{\boldsymbol{z}}\varphi, \boldsymbol{D}(\boldsymbol{z})\nabla\varphi\rangle = 0,
\end{aligned}
$$

which gives (5.51a).

Next, let $\boldsymbol{\omega}(\boldsymbol{z}) := \boldsymbol{A}\boldsymbol{z} - \boldsymbol{B}\boldsymbol{\Gamma}_{\tau}^{-1}(\boldsymbol{z})\boldsymbol{\delta}_{\tau}(\boldsymbol{z})$. We then have

$$
\begin{aligned}
&\nabla_{\boldsymbol{z}} \cdot (\hat{\varphi}\boldsymbol{\omega}(\boldsymbol{z})) = \langle \nabla_{\boldsymbol{z}}\hat{\varphi}, \boldsymbol{\omega}(\boldsymbol{z})\rangle + \hat{\varphi}\nabla_{\boldsymbol{z}} \cdot \boldsymbol{\omega}(\boldsymbol{z}) \\
&= \exp\left(-\psi/2\epsilon\right) \Big(\langle \nabla\sigma^{\mathrm{opt}}, \boldsymbol{\omega}(\boldsymbol{z})\rangle + \sigma^{\mathrm{opt}}\nabla_{\boldsymbol{z}} \cdot (\boldsymbol{\omega}(\boldsymbol{z})) \\
&- \frac{\sigma^{\mathrm{opt}}}{2\epsilon}\langle \nabla_{\boldsymbol{z}}\psi, \boldsymbol{\omega}(\boldsymbol{z})\rangle \Big), \qquad (5.55)
\end{aligned}
$$

and

$$\epsilon \mathbf{1}^\top (\boldsymbol{D} \odot \mathrm{Hess}(\hat{\varphi}))\mathbf{1} = \exp\left(-\frac{\psi}{2\epsilon}\right)\Big(\epsilon \mathbf{1}^\top (\boldsymbol{D} \odot \mathrm{Hess}(\sigma^{\hat{\mathrm{opt}}}))\mathbf{1}$$
$$- \frac{\partial(\boldsymbol{D}_{ij}(\boldsymbol{z})\sigma^{\mathrm{opt}})}{\partial x_j}\frac{\partial \psi}{\partial z_j} - \frac{1}{2}\sigma^{\mathrm{opt}}\langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(\psi)\rangle$$
$$+ \frac{\sigma^{\mathrm{opt}}}{4\epsilon}\langle \nabla_{\boldsymbol{z}}\psi, \boldsymbol{D}(\boldsymbol{z})\nabla_{\boldsymbol{z}}\psi\rangle\Big). \tag{5.56}$$

In (5.50b), taking the partial derivative of $\hat{\varphi}$ w.r.t. t, and using (5.44)-(5.45) together with (5.55)-(5.56), we get

$$\begin{aligned}
\frac{\partial \hat{\varphi}}{\partial t} &= \exp\left(-\psi/2\epsilon\right)\left(\frac{\partial \sigma^{\mathrm{opt}}}{\partial t} - \frac{\sigma^{\mathrm{opt}}}{2\epsilon}\frac{\partial \psi}{\partial t}\right)\\
&= \exp\left(-\psi/2\epsilon\right)\Big(- \nabla_{\boldsymbol{z}}(\sigma^{\mathrm{opt}}\boldsymbol{w}(\boldsymbol{z})) - \nabla_{\boldsymbol{z}}(\sigma^{\mathrm{opt}}\boldsymbol{D}\nabla\psi)\\
&+ \epsilon \mathbf{1}(\boldsymbol{D}(\boldsymbol{z}) \odot \mathrm{Hess}(\sigma^{\mathrm{opt}}))\mathbf{1} + \frac{\sigma^{\mathrm{opt}}}{2\epsilon}\langle \nabla_{\boldsymbol{z}}\psi, \boldsymbol{w}(\boldsymbol{z})\rangle\\
&+ \frac{\sigma^{\mathrm{opt}}}{4\epsilon}\langle \nabla_{\boldsymbol{z}}\psi, \boldsymbol{D}(\boldsymbol{z})\nabla_{\boldsymbol{z}}\psi\rangle + \frac{\sigma^{\mathrm{opt}}}{2}\langle \boldsymbol{D}(\boldsymbol{z}), \mathrm{Hess}(\psi)\rangle\Big)\\
&= -\nabla_{\boldsymbol{z}} \cdot (\hat{\varphi}\boldsymbol{w}(\boldsymbol{z})) + \epsilon \mathbf{1}^\top (\boldsymbol{D} \odot \mathrm{Hess}(\hat{\varphi}))\mathbf{1},
\end{aligned}$$

which is indeed (5.51b). The boundary conditions (5.46) follow directly from (5.50a)-(5.50b). ■

Theorem 14 in principle allows solving problem (5.42) in the following manner. Let $(\varphi_1, \hat{\varphi}_0) := (\varphi(\boldsymbol{z}, t=1), \hat{\varphi}(\boldsymbol{z}, t=0))$ denote the terminal-initial condition pair for the system (5.51a)-(5.51b). By making an arbitrary guess for the pair $(\varphi_1, \hat{\varphi}_0)$, one can perform a fixed point recursion on the Schrödinger system (5.51)-(5.52), and the converged pair $(\varphi_1, \hat{\varphi}_0)$ can then be used to compute the transient pair $(\varphi(\boldsymbol{z}, t), \hat{\varphi}(\boldsymbol{z}, t))$. Then, by (5.50), we recover $(\sigma^{\mathrm{opt}}, \psi)$, and thus $(\sigma^{\mathrm{opt}}, \boldsymbol{v}^{\mathrm{opt}})$ from (5.32). Notice that this procedure with small $\epsilon > 0$ will yield the pair $(\sigma^{\mathrm{opt}}, \boldsymbol{v}^{\mathrm{opt}})$

solving problem (5.23). Finally, the mapping (5.25) in Remark 9 recovers the solution $(\rho^{\text{opt}}, \boldsymbol{u}^{\text{opt}})$ for problem (5.20). This algorithmic framework and its convergence will be the topic of our future research.

Conclusions

We considered the minimum energy joint state PDF steering problem over finite time horizon subject to the multi-input state feedback linearizable dynamics. We showed that the density steering problem can be made amenable in the feedback linearized coordinates. We derived the state feedback controller in terms of the solutions of a pair of coupled HJB and Fokker-Planck-Kolmogorov PDEs. Furthermore, we reduced this system of coupled nonlinear PDEs to a system of boundary-coupled linear PDEs. Our results are expected to lay the foundation for developing computational algorithms solving the density steering problem.

Chapter 6

Global Convergence of Second-Order Dynamics in Two-Layer Neural Networks

Recent results have shown that for two-layer fully connected neural networks, gradient flow converges to a global optimum in the infinite width limit, by making a connection between the mean field dynamics and the Wasserstein gradient flow. These results were derived for first-order gradient flow, and a natural question is whether second-order dynamics, i.e., dynamics with momentum, exhibit a similar guarantee. We show that the answer is positive for the heavy ball method. In this case, the resulting integro-PDE is a nonlinear kinetic Fokker Planck equation, and unlike the first-order case, it has no apparent connection with the Wasserstein gradient flow. Instead, we study the variations of a Lyapunov functional along the solution trajectories to characterize the stationary points and to prove convergence. While our results are asymptotic in the mean field limit, numerical

simulations indicate that global convergence may already occur for reasonably small networks.

6.1 Introduction

The empirical success of neural network models has prompted several theoretical studies that attempt to shed some light on their performance, and provide guarantees under suitable assumptions. For fully connected networks, universal approximation results such as [16, 74] provide a partial explanation for this empirical success, by proving that a large enough network can approximate any continuous function on a compact set, though such results do not address the dynamics of learning, i.e., whether local search algorithms such as gradient descent can find global solutions. Recent works [65, 158] have tackled this question for two-layer networks, and proved convergence to global solutions, by studying the dynamics in the space of distributions over parameters. They make the observation that (Euclidean) gradient flow in the parameter space is equivalent to a Wasserstein gradient flow in the distribution space. This allows for an analysis of the long-time behavior of the dynamics in the mean field limit, i.e., when the width of the network tends to infinity.

To the best of our knowledge, previous works in this setting, such as [65, 158, 188, 194], have only considered first-order gradient dynamics, and a natural question is whether similar guarantees hold for second-order dynamics, i.e., dynamics with momentum. This is the subject of our investigation. Momentum methods such as the heavy ball method [185], Nesterov's method [165], or the Adam method [132], are widely used in practice [206] and have received significant attention in the

optimization literature. Their continuous-time counterpart is given by a family of second-order differential equations, which can be interpreted as damped nonlinear oscillators [9, 37, 87, 91]. For example, [204, 218] studied the continuous-time limit of Nesterov's method, which is an instance in this family with a particular form of damping.

In this paper, our analysis will focus on the heavy ball method–perhaps the simplest and the earliest instance of second-order optimization dynamics. It corresponds to a constant damping coefficient, making the analysis more tractable. Even in this relatively simple setting, the distribution dynamics for a two-layer neural network is given by a nonlinear kinetic Fokker-Planck equation [214], and unlike the first-order case, there is no apparent connection with the Wasserstein gradient flow. Hence, our approach to analyze the mean field dynamics will be somewhat different, even though the tools we use are similar. Our analysis takes inspiration from previous works in the first-order case [65, 158], and also from the study of kinetic Fokker-Planck equations [31, 200].

6.1.1 Two-layer neural networks

We describe the problem setting before summarizing our results. We seek to learn a function $\psi \in \mathcal{F}$, where $\mathcal{F}$ is a Hilbert space equipped with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{F}}$. The model is parameterized by $(\theta_1, \ldots, \theta_n) \in \Theta^n$, and its output is given by

$$\psi_{\theta_1,\ldots,\theta_n}(x) := \frac{1}{n} \sum_{i=1}^{n} \Psi(\theta_i)(x), \tag{6.1}$$

where n is the number of neurons (also referred to as the width of the network), x is the input vector, and $\Psi(\theta_i) \in \mathcal{F}$. In the two-layer neural network setting, we

take $\Theta = \mathbb{R}^d$, $\theta_i = (a_i, b_i) \in \mathbb{R}^{d-1} \times \mathbb{R}$, and $\Psi(\theta_i)(x) = b_i s(\langle a_i, x \rangle)$, where a_i, b_i are the weights of the first and the second layer, respectively, and $s : \mathbb{R} \mapsto \mathbb{R}$ is an activation function. While (6.1) is perhaps an unusual way to describe the output of a neural network, it highlights a structure that lends itself to mean field analysis: the model can be viewed as an average of "basis functions" $\Psi(\theta)(\cdot)$, parameterized by the vector $\theta \in \Theta = \mathbb{R}^d$. This point of view allows us to further rewrite ψ as the integral

$$\psi_\mu = \langle \Psi, \mu \rangle := \int_\Theta \Psi(\theta) \, \mathrm{d}\mu(\theta), \tag{6.2}$$

where μ is a probability measure on Θ, encoding the parameter distribution. When μ is an average of Dirac masses, i.e., $\mu = \frac{1}{n} \sum_{i=1}^n \delta_{\theta_i}$, the integral (6.2) reduces to the summation (6.1).

We are given a convex, Fréchet differentiable functional $R : \mathcal{F} \mapsto \mathbb{R}_+$, referred to as the risk functional, which measures the expected loss of the model. For instance, in the quadratic loss case, $R(\psi) = \frac{1}{2}\mathbb{E}_{(x,y)}|\psi(x) - y|^2 = \frac{1}{2} \langle \psi - y, \psi - y \rangle_{\mathcal{F}}$, where the inner product is taken to be $\langle \psi, \phi \rangle_{\mathcal{F}} := \mathbb{E}_{(x,y)}[\psi\phi]$, and the features and labels (x, y) are sampled from a data distribution D. We are also given a regularization function $g : \Theta \to \mathbb{R}_+$, and we consider the regularized risk $R(\Psi_{\theta_1,\ldots,\theta_n}) + \frac{1}{n} \sum_{i=1}^n g(\theta_i)$. While this is, in general, a non-convex function of θ, when lifted to the space of probability measures, it becomes

$$F(\mu) := R(\langle \Psi, \mu \rangle) + \langle g, \mu \rangle \, . \tag{6.3}$$

The functional F is convex and Fréchet differentiable. Thus, the learning problem can be recast as a measure-valued convex optimization problem, $\inf_\mu F(\mu)$. This is the point of view taken in [65, 158], as well as in earlier works such as [10, 24].

We note that while our motivation is the study of neural network dynamics, this setting applies to other problems, see [10, 65].

6.1.2 Particle and distribution dynamics

We give an informal overview of the general strategy used to study the mean field limit. The first step is to make a connection between the dynamics of the particles $(\theta_1, \ldots, \theta_n) \in \Theta^n$, and the dynamics of their distribution μ supported over Θ. Suppose the particles move following a (time-varying) vector field $v_t : \Theta \mapsto \mathbb{R}^d$, where t indexes time. That is, the trajectories are solutions to the following ordinary differential equation (ODE):

$$\dot{\theta}_t = v_t(\theta_t). \tag{6.4}$$

Then at time t, the distribution μ_t of these particles – more precisely, the push-forward of an initial distribution μ_0 by the flow of the ODE (6.4) – is given by the solution to a partial differential equation (PDE) known as the continuity equation:

$$\partial_t \mu_t + \nabla \cdot (\mu_t v_t) = 0, \tag{6.5}$$

where $\nabla\cdot$ stands for the divergence operator. When μ_t does not have a density, (6.5) should be interpreted distributionally. Let μ_t^n be the solution of (6.5) initialized at $\frac{1}{n}\sum_{i=1}^n \delta_{\theta_i}$. One can show that if the initial positions $(\theta_1, \ldots, \theta_n)$ are drawn from a fixed distribution μ_0, then as the number of particles n tends to infinity, the solutions μ_t^n, weakly converge to the solution μ_t of (6.5) initialized at μ_0. One can then focus on studying the dynamics of the mean field limit μ_t.

The vector field v_t in (6.4) and (6.5), which describes the movement of the parti-

cles, is determined by the particular learning dynamics under consideration. The choice of v_t in [65,158,188,194] corresponds to the first-order gradient flow in θ. In this case, the PDE (6.5) has additional structure: it corresponds to a Wasserstein gradient flow of the functional F in (6.3), as is well known in the optimal transportation literature [6,125]. We study a different choice of v_t that corresponds to the heavy ball method. Our primary goal is to provide, under suitable assumptions, convergence guarantees to the global minimizers of F, which is similar in spirit with [65,158]. The latter two results differ: in [65], the authors study *deterministic* gradient flow under a homogeneity assumption on the objective, and prove that *assuming* μ_t *converges*, it can only converge to a global minimum. In [158], the authors study *noisy* gradient flow for the quadratic loss, and prove that μ_t converges arbitrarily close to the global minimum (depending on the magnitude of the noise). Our approach is closer to the latter: we study convergence of the *noisy* heavy ball method. The convergence of the noiseless second-order dynamics in the mean field limit remains an open question.

6.1.3 Summary of contributions

We start by deriving the distributional PDE associated with the heavy ball method (Section 6.2). To study the dynamics in the mean field limit, we define a Lyapunov functional in Section 6.3 and bound its variations along solution trajectories. This relies on a general criterion given in Lemma 15, which reveals a close connection between the Lyapunov functions for dynamics with no particle interaction (as in convex optimization) and the Lyapunov functionals for mean field dynamics.

Equipped with this result, in Section 6.4, we characterize the stationary solutions in Theorem 6, and show that they must satisfy a Boltzmann fixed point equation,

for which we prove the existence and uniqueness of a solution in Proposition 17. Furthermore, we show in Theorem 7 that the solution trajectory converges to this unique stationary point. Finally, we show in Theorem 8 that by using vanishingly small noise, the limit can be made arbitrarily close to the global infimum of F.

In Section 6.5, we illustrate these results with numerical experiments that include other variants of second-order dynamics beyond the heavy ball method. The experiments suggest that the convergence may already occur with a reasonably small number of particles.

The proofs are deferred to the appendix.

6.2 Mean field second-order dynamics

6.2.1 Assumptions

Let $F_0(\mu) := R(\langle \Psi, \mu \rangle)$ denote the unregularized loss, and $F(\mu) = F_0(\mu) + \langle g, \mu \rangle$. We make the following assumptions:

(A1) $R : \mathcal{F} \mapsto \mathbb{R}_+$ is convex, Fréchet differentiable.

(A2) $\Psi : \Theta \mapsto \mathcal{F}$ is Fréchet differentiable.

(A3) $\nabla F'(\rho) \in L^\infty(\Theta)$ for all $\rho \in \mathcal{P}(\Theta)$.

(A4) $\{F_0'(\rho) : \rho \in L^1(\Theta), \|\rho\|_1 \leq 1\}$ is uniformly equicontinuous and uniformly bounded, and $g : \Theta \to \mathbb{R}_+$ is differentiable and confining, i.e. $\lim_{|\theta| \to \infty} g(\theta) = \infty$ and $\exp(-\beta g)$ is integrable for all $\beta > 0$.

We discuss some of the implications of these assumptions. (A1) and (A2) are basic

regularity assumptions implying that F is Fréchet differentiable and $F'(\rho) : \Theta \to \mathbb{R}$ is a differentiable function of θ, so the gradient of the loss in the parameter space is well-defined. (A3) is used to prove the existence and uniqueness of solutions of the PDE. The assumption that the loss is regularized by g, together with the condition (A4), are important to guarantee existence of a stationary solution, as is common in the literature [31, 158, 180]. In particular, the assumption that the regularizer g is confining is rather mild; it essentially requires $g(\theta)$ to grow sufficiently fast when $|\theta|$ tends to infinity. One simple choice is to take $g(\theta) = |\theta|$ (except near 0, since we also require differentiability of g). In Appendix A.24, we explicate the foregoing assumptions in the context of the quadratic loss, and show that they are implied by the assumptions made in previous work.

6.2.2 Second-order dynamics

Given a differentiable $f : \mathbb{R}^d \mapsto \mathbb{R}$, a broad family of second-order dynamics is described by the ODE

$$\ddot{\theta}_t = -\nabla_\theta f(\theta) - \gamma_t \dot{\theta}_t, \tag{6.6}$$

which can be interpreted as a dissipative nonlinear oscillator with potential f, and damping coefficient γ_t, see [9, 37, 106, 204]. Under certain assumptions, such as convexity of f, it can be shown that the solutions converge to global minimizers of f, see [37]. When $\gamma_t \equiv \gamma/t$ for some positive constant γ, this corresponds to Nesterov's method in continuous-time [204], and when $\gamma_t \equiv \gamma$ is a time-independent positive constant, it corresponds to the heavy ball method [9, 87].

As will become clear shortly, the potential f in our setting is time-varying due to the interaction between particles. Recall from (6.3) that the objective functional

is $F(\mu) = R(\langle \Psi, \mu \rangle) + \langle g, \mu \rangle$, where μ is a distribution over parameters $\theta \in \Theta$. The gradient of the objective in the parameter space Θ is given by $\nabla F'(\mu)(\theta)$, see Appendix A.12 for a detailed derivation. In particular, setting $v_t(\theta) \equiv -\nabla F'(\mu_t)(\theta)$ in (6.4)-(6.5) corresponds to the first-order gradient flow, as in [65, 158].

In the second-order case, it is convenient to write (6.6) as a system of two first-order equations describing the evolution of position-velocity pair $(\theta, r) \in \mathcal{T}\Theta$ (the tangent bundle). Then $\mu_t \in \mathcal{M}(\mathcal{T}\Theta)$ denotes the joint distribution over $\mathcal{T}\Theta$ at time t, and $[\mu_t]^\theta$ is the corresponding marginal measure over Θ. We suppose that a Brownian motion is applied to the velocity (or rate) r, resulting in the following Itô stochastic differential equation (SDE):

$$\mathrm{d}\begin{pmatrix} \theta \\ r \end{pmatrix} = \begin{pmatrix} r \\ -\nabla F'([\mu_t]^\theta)(\theta) - \gamma r \end{pmatrix} \mathrm{d}t + \begin{pmatrix} 0 \\ \sqrt{2\gamma\beta^{-1}}\, \mathrm{d}W_t \end{pmatrix}, \tag{6.7}$$

where the parameter $\beta > 0$ is referred to as the inverse temperature, $\gamma > 0$ is the constant damping coefficient, and W_t is the standard Wiener process in the tangent space of Θ. Eq. (6.7) is an underdamped Langevin equation with interaction potential $F'([\mu_t]^\theta)$. It describes the stochastic heavy ball method [87] in the parameter space. The dependence of the potential on μ_t reflects the fact that the output of the neural network (and its loss) depend not on a single particle, but on the distribution of all particles. Note that the dependence is on the marginal $[\mu_t]^\theta$, since the loss only depends on positions, and not velocities.

The distribution dynamics corresponding to (6.7) is given by

$$\partial_t \mu_t = -\nabla . \left[\mu_t . \begin{pmatrix} r \\ -\nabla F'([\mu_t]^\theta) - \gamma r \end{pmatrix} \right] + \gamma\beta^{-1} \Delta_r \mu_t, \tag{6.8}$$

where Δ_r denotes the Laplacian operator w.r.t. the r variable, and corresponds to the Brownian motion applied to r. The integro-PDE (6.8) is a nonlinear kinetic Fokker-Planck equation.

Consistency of the mean field limit We now provide a consistency result between the second-order particle dynamics and the mean field limit.

Theorem 5. *Let $\mu_0 \in \mathcal{M}(\mathcal{T}\Theta)$. Consider a set of n interacting particles with states $\{(\theta_t^i, r_t^i)\}_{i=1}^n$ where i denotes the i-th particle. Suppose these particles solve copies of the SDE* (6.7) *indexed by $i = 1, \dots, n$, in which μ_t is replaced by the empirical distribution $\mu_t^n := \frac{1}{n}\sum_{i=1}^n \delta_{(\theta_t^i, r_t^i)}$, and with initial states $\{(\theta_0^i, r_0^i)\}_{i=1}^n$ sampled independently from μ_0. Then there exists $\mu_t \in \mathcal{M}(\mathcal{T}\Theta)$ such that, almost surely, $\mu_t^n \to \mu_t$ weakly, as $n \to \infty$. Furthermore, μ_t solves* (6.8) *in the weak distributional sense with initial condition μ_0.*

This result motivates the study of the long-time behavior of the mean field limit μ_t. One of the advantages of the mean-field setting is that one can work with absolutely continuous distributions, which simplifies the analysis.

Existence and uniqueness of solutions We make the following assumption on the initial condition μ_0, both to obtain existence and uniqueness of a solution, and to guarantee finiteness of the free energy, introduced in the next section.

(A5) μ_0 is absolutely continuous, and the associated PDF ρ_0 satisfies $\langle g(\theta) + |r|^2/2, \rho_0\rangle < \infty$, $\langle \log^+ \rho_0, \rho_0\rangle < \infty$, and $\int |\nabla F_0'([\rho_0]^\theta)(\theta)|^2 \mathrm{d}\theta < \infty$,

where $\log^+ \rho := \max\{\log \rho, 0\}$. It is known (see [31, 209]) that if F satisfies (A3) and the initial condition μ_0 satisfies (A5), then (6.8) admits a unique solution $\mu_t \in C([0,\infty), \mathcal{M}(\mathcal{T}\Theta))$. That μ_t remains absolutely continuous for all $t \geq 0$,

and hence ρ_t exists, will be proved in Theorem 7. Since the solution is absolutely continuous, μ_t and $[\mu_t]^\theta$ in (6.7)-(6.8) can be replaced by the corresponding PDFs ρ_t and $[\rho_t]^\theta$, respectively.

6.2.3 The linear case

When $F(\mu) = \langle f, \mu \rangle$ for some function $f : \Theta \mapsto \mathbb{R}$, we have $F'(\mu) = f$ and there is no particle interaction in (6.7). While this situation is irrelevant in the neural network context, it is instructive to review results in the linear setting. In this case, (6.8) becomes a *linear* Fokker-Planck PDE, which can be shown (e.g., Proposition 6.1 in [180]) to admit a unique stationary solution with PDF $\rho_\infty(\theta, r) = \exp\left(-\beta\left(f(\theta) + \frac{1}{2}|r|^2\right)\right)/Z$, where Z is a normalizing constant. Under additional assumptions on the confining potential f, one can also study the rate of convergence of ρ_t to ρ_∞, see e.g. [13, 106, 180, 214]. Our situation corresponds to a *nonlinear* kinetic Fokker-Planck equation, which is not well-understood in the general setting. Some special cases have been studied in the literature, such as when the interaction potential $F'(\mu)$ is a convolution [31, 46, 200, 214]. The convolution structure in these references is motivated from physical dynamics–the electrostatic Coulomb interaction in plasma and semiconductor dynamics [1, 78, 80], and the gravitational Newton interaction in stellar dynamics [18, 51, 175]–and leads to the Valsov-Poisson-Fokker-Planck equations [117, 183, 209]. Unfortunately, this is not the case in our neural network setting, and these results do not directly apply. However, we will use similar techniques, and will prove that the stationary solutions have a similar characterization.

6.3 Variations of a Lyapunov functional

Hereafter, we work with the PDF trajectory $\rho_t \in C\left([0,\infty), \mathcal{P}(\mathcal{T}\Theta)\right)$ associated with the measure-valued solution trajectory $\mu_t \in C\left([0,\infty), \mathcal{M}_{\mathrm{ac}}(\mathcal{T}\Theta)\right)$ for (6.8). To characterize the stationary solutions, we will study the variations of the following Lyapunov functional, defined for $\rho \in \mathcal{P}(\mathcal{T}\Theta)$,

$$\mathcal{E}(\rho) := F([\rho]^\theta) + \left\langle \frac{1}{2}|r|^2, \rho \right\rangle + \beta^{-1} H(\rho), \tag{6.9}$$

where $H(\rho) := \langle \log \rho, \rho \rangle$ is the negative entropy. The functional $\mathcal{E}$ is often referred to as the *free energy*. In this section, we show that along the trajectory ρ_t, the functional $\mathcal{E}$ is non-increasing.

Lemma 15. *Let v_t be a vector field over $\mathcal{T}\Theta$, and let ρ_t be a solution of the continuity equation $\partial_t \rho_t = -\nabla.(\rho_t v_t)$ with initial condition $\rho_0 \in \mathcal{P}(\mathcal{T}\Theta)$. Let $\mathcal{V} : \mathcal{P}(\mathcal{T}\Theta) \mapsto \mathbb{R}$, and suppose that along ρ_t, $\mathcal{V}$ is Fréchet differentiable and $\mathcal{V}'(\rho_t) : \Theta \mapsto \mathbb{R}$ is differentiable. Then for all $t \geq 0$,*

$$\partial_t \mathcal{V}(\rho_t) = \langle \nabla \mathcal{V}'(\rho_t), \rho_t v_t \rangle_* .$$

This gives us a simple criterion for a functional $\mathcal{V}$ to be non-increasing along solution trajectories with vector field $v(\rho)$: it suffices that for all ρ, the inequality $\langle \nabla \mathcal{V}'(\rho), v(\rho) \rangle \leq 0$ holds ρ-a.e. In the case with no interaction, i.e., $\mathcal{V}(\rho) \equiv \langle V, \rho \rangle$ is linear with $V : \Theta \mapsto \mathbb{R}$, and the vector field v is independent of ρ, the condition reduces to $\langle \nabla V, v \rangle \leq 0$, which defines Lyapunov functions for single particle dynamics. With this observation, the free energy $\mathcal{E}$ can be viewed as a mean field generalization of $E(\theta, r) := f(\theta) + \frac{1}{2}|r|^2$, which is known to be a Lyapunov function for the heavy ball dynamics, see Appendix A.14.

To apply Lemma 15 to the free energy $\mathcal{E}$ in (6.9), we use the identity $\Delta_r \rho = \nabla_r \cdot \nabla_r \rho$ to formally rewrite (6.8) as follows.

$$\partial_t \rho_t = -\nabla.(\rho_t v(\rho_t)), \quad v(\rho)(\theta, r) := \begin{pmatrix} r \\ -\nabla F'([\rho]^\theta)(\theta) - \gamma r - \gamma \beta^{-1} \nabla_r \log \rho(\theta, r) \end{pmatrix}. \tag{6.10}$$

Proposition 16. *Consider the Lyapunov functional $\mathcal{E}$ in (6.9). Let ρ_t be the solution to* (6.10)*. Then*

$$\partial_t \mathcal{E}(\rho_t) = -\gamma \left\langle |r + \beta^{-1} \nabla_r \log \rho_t|^2, \rho_t \right\rangle.$$

Proof. From (6.9), we obtain

$$\mathcal{E}'(\rho)(\theta, r) = F'([\rho]^\theta)(\theta) + \frac{1}{2}|r|^2 + \beta^{-1}(1 + \log \rho(\theta, r)), \tag{6.11}$$

and, using the shorthand $\ell_\theta := \beta^{-1} \nabla_\theta \log \rho_t$, $\ell_r := \beta^{-1} \nabla_r \log \rho_t$, we compute

$$\begin{aligned} \partial_t \mathcal{E}(\rho_t) &= \langle \nabla \mathcal{E}'(\rho_t), \rho_t v(\rho_t) \rangle_* \\ &= \left\langle \begin{pmatrix} \nabla_\theta F'([\rho_t]^\theta) + \ell_\theta \\ r + \ell_r \end{pmatrix}, \rho_t \begin{pmatrix} r \\ -\nabla_\theta F'([\rho_t]^\theta) - \gamma r - \gamma \ell_r \end{pmatrix} \right\rangle_* \\ &= \left\langle -\gamma \langle r, r \rangle - \gamma \langle \ell_r, \ell_r \rangle - 2\gamma \langle r, \ell_r \rangle + \langle \ell_\theta, r \rangle - \left\langle \ell_r, \nabla_\theta F'([\rho_t]^\theta) \right\rangle, \rho_t \right\rangle. \end{aligned}$$

We conclude by showing that the last two terms, $\langle \langle \ell_\theta, r \rangle, \rho_t \rangle$ and $\left\langle \left\langle \ell_r, \nabla_\theta F'([\rho_t]^\theta) \right\rangle, \rho_t \right\rangle$ are equal to zero, using duality of the $\nabla \cdot$ and ∇ operators. The details are provided in Appendix A.17. ■

The proposition states that the free energy $\mathcal{E}$ is non-increasing along solution trajectories. This fact, together with additional bounds derived in Appendix A.14,

are the primary ingredients used to prove our main results in the next section.

6.4 Stationary solutions and global convergence

We say $\rho^\star \in \mathcal{P}$ is a stationary solution of (6.10) if the solution $(\rho_t)_{t\geq 0}$ obtained with the initial condition $\rho_0 \equiv \rho^\star$, satisfies $\rho_t \equiv \rho^\star$ for all $t \geq 0$.

In this section, we state our main results (proved in Appendix A.19), by characterizing stationary solutions (Theorem 6), proving their existence and uniqueness (Proposition 17), and establishing convergence of ρ_t, as $t \to \infty$, to the unique stationary point (Theorem 7). Furthermore, we show that the limit can be made arbitrarily close to the global infimum (Theorem 8).

Theorem 6. *Suppose $\rho^\star \in \mathcal{P}(\mathcal{T}\Theta)$ is a stationary solution of* (6.10)*. Then,*

$$\rho^\star(\theta, r) = \frac{\exp\left(-\frac{\beta}{2}|r|^2\right)}{Z_1}[\rho^\star]^\theta(\theta). \tag{6.12}$$

where Z_1 is the normalizing constant $Z_1 := \int \exp(-\frac{\beta}{2}|r|^2)\mathrm{d}r$ and $[\rho^\star]^\theta$ is the θ marginal. Furthermore, $[\rho^\star]^\theta$ solves the following fixed point equation:

$$\rho(\theta) = \frac{\exp\left(-\beta F'(\rho)(\theta)\right)}{Z_2(\rho)}, \rho \in \mathcal{P}(\Theta) \tag{6.13}$$

where $Z_2(\rho) := \int \exp\left(-\beta F'(\rho)(\theta)\right) \mathrm{d}\theta$.

The proof crucially relies on the variation of the free energy given in Proposition 16. The theorem states that a stationary solution, if it exists, must be a product distribution, where the r marginal is a Gaussian, and the θ marginal satisfies a fixed point equation. This product structure is familiar from the linear case (Section 6.2.3), where $F'(\rho) \equiv f$ and the RHS of (6.13) becomes independent

of ρ, and simply describes a Gibbs distribution. In the nonlinear case, it is not guaranteed, a priori, that (6.13) admits a solution. This is proved in the next proposition; our existence proof invokes Schauder's fixed point theorem [88, p. 286, Theorem 11.6], and this is where assumption (A4) comes into play.

Proposition 17. *Suppose assumption (A4) holds, and let* $T : \mathcal{P}(\Theta) \mapsto \mathcal{P}(\Theta)$ *be defined as follows:*

$$T(\rho)(\theta) = \frac{\exp\left(-\beta F'(\rho)(\theta)\right)}{Z_2(\rho)},$$

where $Z_2(\rho) = \int \exp\left(-\beta F'(\rho)\right)$*. Then* T *has a unique fixed point.*

We next show that the solution trajectory ρ_t converges to the unique stationary solution $\rho^\star$, under mild assumptions on the initial condition.

Theorem 7. *Consider a measure* $\mu_0 \in \mathcal{M}_{\mathrm{ac}}\left(\mathcal{T}\Theta\right)$ *satisfying the assumption (A5). Starting from such an initial condition* μ_0*, the solution* $(\mu_t)_{t\geq 0}$ *of (6.10) satisfies the following.*

(i) For each $t \geq 0$*, the measure* $\mu_t \in \mathcal{M}_{\mathrm{ac}}\left(\mathcal{T}\Theta\right)$*, i.e., the associated joint PDF* $\rho_t \in \mathcal{P}\left(\mathcal{T}\Theta\right)$ *exists.*

(ii) The trajectory $(\rho_t)_{t\geq 0}$ *converges strongly in* L^1 *to the unique stationary solution* $\rho^\star$ *of (6.10) as* $t \to \infty$*.*

Now that we have established the existence and uniqueness of a stationary solution $\rho^\star \in \mathcal{P}(\mathcal{T}\Theta)$, and convergence to $\rho^\star$, we will relate, in the next theorem, $F([\rho^\star]^\theta)$ to $\inf_{\rho\in\mathcal{P}(\Theta)} F(\rho)$. Some intuition can be gained from the linear case: when $F(\rho) = \langle f, \rho\rangle$, the stationary solution is simply given by the Gibbs distribution, $[\rho^\star]^\theta(\cdot) \propto \exp(-\beta f(\cdot))$, which concentrates around the minimizers of f as $\beta \to \infty$, thus $F([\rho^\star]^\theta)$ approaches $\inf_{\rho\in\mathcal{P}(\Theta)} F(\rho)$ as $\beta \to \infty$. The same holds in our non-linear setting, as stated in the next theorem.

For $\lambda \in [0,1]$, let $F_\lambda(\rho) := R(\langle \Psi, \rho \rangle) + \lambda \langle g, \rho \rangle$, so that $F_1 \equiv F$.

Theorem 8. *Let $\rho^\star$ be the stationary solution of (6.10), and let $[\rho^\star]^\theta$ be its marginal. Then there exists a constant C that depends on F and d, such that for all $\beta \geq 1$,*

$$F_{1-1/\beta}([\rho^\star]^\theta) \leq \inf_{\rho \in \mathcal{P}(\Theta)} F_1(\rho) + \frac{C + d \log \beta}{\beta}.$$

The proof of the above theorem has two components: the first is the observation that $\rho^\star$ is a minimizer of the free energy $\mathcal{E}$ (this follows from the characterization in Theorem 6), the second is the bounds on the difference between $\mathcal{E}$ and F derived in Appendix A.14.

As a consequence of the theorem, the objective value at the stationary point can be made arbitrarily close to the global infimum of F by taking β large enough. It is worth emphasizing that the presence of noise, i.e., the diffusion term in (6.8), is essential in guaranteeing existence and uniqueness of the stationary distribution. In the noiseless case, there may exist multiple stationary points that are not global minimizers. The addition of noise can be thought of as an *entropic regularization* of the functional $F(\rho)$, and Theorem 8 says that one can approach the infimum of the unregularized problem in the small noise limit.

6.5 Numerical simulations

To illustrate our results, we run synthetic numerical experiments following the setup used in [65]. The model ψ is a two-layer neural network, as described in Section 6.1.1, with sigmoid activation function $s(\cdot)$ and width n, i.e., $\psi(x) = \frac{1}{n} \sum_{i=1}^n b_i s(\langle a_i, x \rangle)$. The features x are normally distributed in $\mathbb{R}^{d-1}$, and the ground truth labels are generated using a similar neural network $\psi^\star$ with width

n_0, i.e., $y = \psi^\star(x)$. The risk functional is quadratic, i.e., $R(\psi) = \frac{1}{2}\|\psi(x) - y\|^2_{\mathcal{F}} = \frac{1}{2}\mathbb{E}_x[(\psi(x) - \psi^\star(x))^2]$, where the expectation is over the empirical distribution. We implement the stochastic heavy ball method using a simple Euler-Maruyama discretization of (6.7), this will be referred to as (SHB) in the figures. We also implement noiseless, second order dynamics: the heavy ball method, referred to as (HB), and Nesterov's accelerated gradient descent, referred to as (AGD).

6.5.1 Convergence to the global infimum

In a first set of experiments, we set the dimension to $d = 100$, and vary the width n of the model, while keeping the width of the ground truth network fixed to $n_0 = 20$. No regularization is used in this experiment, so that the model can theoretically achieve zero loss whenever $n \geq n_0$. The results are reported in Figure 6.1. In the left subplot, each method is run for 10^5 iterations, and we measure the loss at the last iteration. We repeat the experiment 20 times and plot the average (represented by the lines) and the individual numbers (scatter plot). The right subplot shows the full trajectory for one realization, for the width $n = 100$. The results suggest that the dynamics converge to the global infimum even with a reasonably small width n. The results also highlight the effect of noise: the stochastic heavy ball method converges closer to the global minimum when β is larger, consistent with Theorem 8. Finally, the results for the noiseless heavy ball method and Nesterov's method suggest that convergence may occur for a broader class of second-order dynamics than the setting of our analysis.

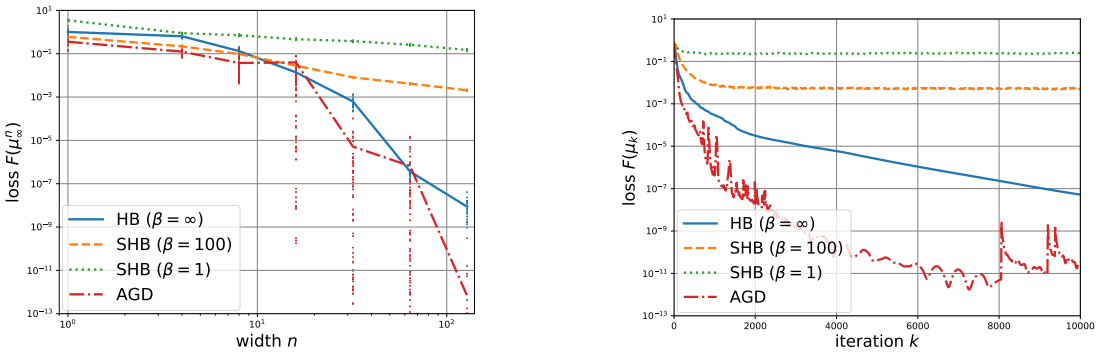

Figure 6.1: Final loss value as the width n of the network increases for several second-order dynamics (left), and sample trajectories for $n = 100$ (right).

6.5.2 Stationary distribution

In a second experiment, we illustrate the characterization of the limiting distribution, which according to Theorem 6, is the product of its marginals, where the r marginal is a Gaussian $\propto \exp(-\beta|r|^2)$, and the θ marginal is $\propto \exp(-\beta F'([\rho_\infty]^\theta))$. Recall from (6.3) that $F(\mu) = R(\langle \Psi, \mu \rangle) + \langle g, \mu \rangle$, thus

$$F'(\mu)(\cdot) = \langle R'(\langle \Psi, \mu \rangle), \Psi(\cdot) \rangle_{\mathcal{F}} + g(\cdot),$$

where g is the regularizer, which we set to $g(\theta) = 0.01|\theta|$ in this experiment. The risk functional is $R(\psi) = \frac{1}{2}\|\psi - \psi^\star\|^2_{\mathcal{F}}$, thus $R'(\psi) = \psi - \psi^\star$, and

$$F'([\mu]^\theta)(\cdot) = \mathbb{E}_x[(\left\langle \Psi, [\mu]^\theta \right\rangle (x) - \psi^\star(x))\Psi(\cdot)(x)] + g(\cdot). \tag{6.14}$$

In particular, if we apply this expression to the empirical distribution of the particles $[\mu^n]^\theta = \frac{1}{n}\sum_{i=1}^n \delta_{\theta_i}$, and use m independent samples $x_j \sim D$ to approximate the expectation, we obtain

$$F'([\mu^n]^\theta)(\cdot) \approx g(\cdot) + \frac{1}{m}\sum_{j=1}^m \Big(\frac{1}{n}\sum_{i=1}^n \Psi(\theta_i)(x_j) - \psi^\star(x_j)\Big)\Psi(\cdot)(x_j). \tag{6.15}$$

This gives us an expression of the Boltzmann distribution that we can approximate numerically in the finite particle case, by using $[\mu_k^n]^\theta$ for large k, in place of $[\rho_\infty]^\theta$ in $\exp(-\beta F'([\rho_\infty]^\theta))$.

We rerun the same experiment described above, with $n = 200$, $n_0 = 20$, and in lower dimension $d = 2$, so that we can visualize the distributions, and compare the empirical and theoretical marginals at the end of training. The result is shown in Figure 6.2, where the empirical marginals (scatter plot) appear to be consistent with the numerical approximation of the Boltzmann distribution (heat map).

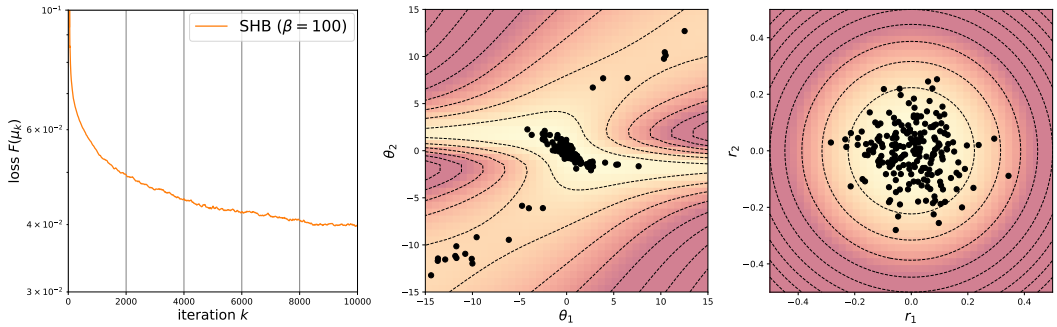

(a) Loss $F(\mu_k^n)$ **(b)** θ marginal **(c)** r marginal

Figure 6.2: Illustration of the limiting distribution under stochastic heavy ball dynamics. The loss value as a function of iteration number is shown on the left. The middle and the right plots show the marginal distributions of position θ and velocity r, at the last iteration $k = 10^5$ (scatter plot). The heat map shows a numerical approximation of the theoretical limiting distributions according to Theorem 6. The level sets represent the log of the density of the marginals, i.e. $-\beta F'(\mu_k^n)(\theta)$, and $-\beta|r|^2/2$ respectively.

6.5.3 Illustration of the interaction potential

Finally, we illustrate the interpretation of the learning dynamics as interacting particles. One can view the dynamics of the network parameters $(\theta_i)_{i=1,\dots,n}$ as evolving in a static potential given by the loss function $f(\theta_1, \dots, \theta_n) = F\left(\frac{1}{n}\sum_{i=1}^n \delta_{\theta_i}\right)$,

defined on Θ^n. But because $\nabla_{\theta_j} f(\theta_1, \ldots, \theta_n) = \frac{1}{n} \nabla F' \left(\frac{1}{n} \sum_i \delta_{\theta_i} \right) (\theta_j)$ (see Appendix A.12), in fact each of the n particles is subject to the *same time-varying potential* $F'(\mu^n) : \Theta \to \mathbb{R}$, where $\mu^n = \frac{1}{n} \sum_{i=1}^n \delta_{\theta_i}$ is the empirical distribution. The potential at any time depends on the joint distribution of particles at that time, but as the distribution converges, the interaction potential also converges. To illustrate this, we plot in Figure 6.3 the evolution of $F'(\mu_k^n)$ as the step k increases.

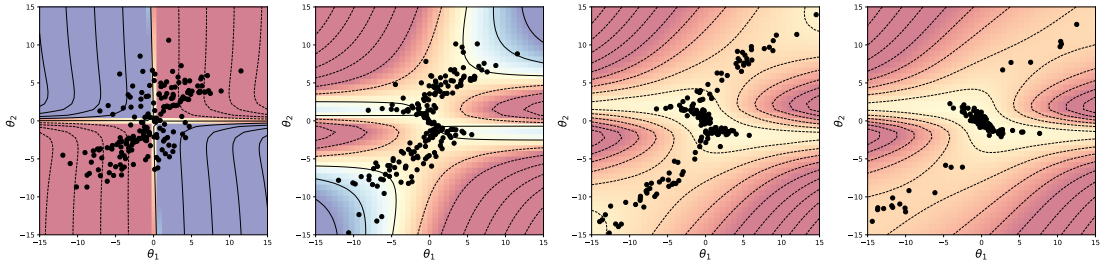

(a) $k = 10$ **(b)** $k = 100$ **(c)** $k = 1000$ **(d)** $k = 10000$

Figure 6.3: Evolution of the interaction potential $F'(\mu_k^n)$ as k increases.

6.6 Concluding remarks

We studied the stochastic heavy ball dynamics in the mean field limit, and established convergence to global minimizers. This is, to our knowledge, the first global convergence guarantee for second-order dynamics in this setting. Though the result is asymptotic, numerical experiments on synthetic problems suggest that the convergence occurs for networks of reasonably small size.

There are several possible directions to investigate quantitative results. For example, hypocoercivity [214] is concerned with the study of the rate of convergence of ρ_t to its limiting distribution, and while the theory is in its early development for the nonlinear case, we believe the techniques can be adapted under additional

assumptions on F. A second direction is the study of fluctuations of solutions, which quantifies the convergence of μ_t^n to the mean field limit μ_t as the number of particles $n \to \infty$, as was done in [188, 194] for gradient flow with quadratic loss. A third direction is to study the generalization properties of the limit. In the gradient flow case, this was investigated for instance by [66] for the logistic loss.

We believe our results can be generalized to a broader family of second-order dynamics, including Nesterov's method. One technical challenge in doing so is that the dynamics has a time-dependence due to the damping coefficient γ_t, which may require using a different Lyapunov functional.

Finally, the question of convergence for the noiseless second-order dynamics remains unsettled and requires further investigation. In general, without diffusion, there may exist stationary points that are not global minimizers (even in the mean field limit). However, one can hope to prove, under suitable assumptions on F, that such stationary points are repulsive, as was done in [65] for gradient flow.

Chapter 7

Conclusions

In this dissertation we have presented the geometric theory of generalized gradient flows. This theory provides a unifying framework that can be used to solve and compute problems in density prediction, control and learning. In Chapter 1, we gave a brief description of these subproblems and their motivation applications from different fields. We also presented a short summary of the theory of gradient flows in the space of PDFs and their connections to each of the three subproblems. Building on this geometric theory, we developed a novel algorithm to solve the density prediction in Chapter 1. This algorithm was used to compute the flow of the joint PDFs of continuous time stochastic nonlinear systems motivated by several engineering examples. We show that this proposed computational framework is convergent by establishing certain contractive properties. The efficiency of the proposed algorithm in terms of computational runtime was demonstrated by numerical examples.

In Chapter 2, we addressed the Schrödinger Bridge Problem (SBP) which is the problem of minimum energy finite horizon steering of state density between two

prescribed endpoint densities via feedback control, subject to trajectory-level dynamics with nonlinear drift. We derived the conditions for optimality and showed that two specific cases, gradient and mixed-conservative dissipative drifts, are particularly amenable for computational purposes using the algorithms we proposed in Chapter 1. This work contributes to a growing body of algorithmic advances to solving different variants of the SBP of interest. In Chapter 3, we introduced the Reflected Schrödinger Bridge Problem (RSBP) which is an SBP subject to reflecting boundary conditions on the controlled state trajectories. Combining our prior work on Wasserstein proximal recursions with some recent results on contraction mapping associated with the Schrödinger system, we provided a computational pipeline for optimal feedback synthesis. Numerical examples were given to highlight the proposed framework. In Chapter 4, we considered the SBP over input state feedback linearizable dynamics. We showed that the density steering problem can be solved in linearized feedback coordinates. We derived the state feedback controller in terms of the solutions of a pair of coupled HJB and Fokker-Planck-Kolmogorov PDEs. Furthermore, we reduced this system of coupled nonlinear PDEs to a system of boundary-coupled linear PDEs. Our results are expected to lay the foundation for developing computational algorithms solving the density steering problem.

In Chapter 6, we characterized the convergence of neural networks trained by stochastic heavy ball dynamics in the mean field limit, and established convergence to global minimizers. We show that rate of convergence to the stationary point is exponential under a weighted Wasserstein distance. This work is one of the first works to provide a proof of global guarantees for second-order dynamics.

We outline some of the possible research directions that will be pursued based on the results of this work.

7.1 Density Prediction

The classical JKO scheme was originally proposed for systems with gradient drift and the scaled isotropic noise, and later was extended to account for degenerate diffusions with mixed conservative-dissipative drifts, many engineering systems in practice do not admit such canonical forms. For example, the following cases are of significant interest:

- Euler's rigid body equations whose drift is a sum of a linear vector field and a divergence free and non-gradient vector field
- A system of coupled second order nonlinear oscillators with mixed conservative-dissipative drift with non-isotropic diffusion
- Generic nonlinear stochastic systems whose structure can be exploited to yield an appropriate distance and free energy functional

The first is a prototypical system in systems-control theory, while the second appears in dynamical models of kinematic models and interconnected power grids as well as in models in mathematical biology. Designing a JKO-like scheme for these systems amounts to amending the distance functional d and the free energy functional in the proximal operator. Finding a suitable distance function and free energy functional remains a non-trivial task. Ideally, we would like to develop a theory to design a scheme that encompasses any choice drift $\boldsymbol{f}$ and diffusion $\boldsymbol{g}$.

Next, we would like to explore the computational complexity of these algorithms and how they compare to traditional methods. The main computational bottleneck in the proximal recursion algorithm is in the multiplication of the exponentiated cost matrix $\mathbf{\Gamma}$ of size $m \times n$ with an n dimensional vector where m is the

number of particles $\boldsymbol{\varrho}_{k-1}$ and n is the number of particles $\boldsymbol{\varrho}_k$. When the dimension of the state space increases n, m also tend to increase to account for the sparsity. We would like to investigate techniques that would speed up the computational performance of this algorithm for a large number of particles.

7.2 Density Control

The classes of SBP problems investigated in this dissertation were limited to quadratic costs. A natural question is how to extend these results to generic Lagrangians. It is expected that the conditions for optimality will be more challenging. Given its contractive properties, it would be ideal to transform this system into Schrödinger System; however it is not guaranteed that the Hopf-Cole Transform will yield such a system except for a few cases. We would like to investigate whether such a generalized transform exists.

One other direction is to consider generic prior linear dynamics. The cases that we considered were limited to gradient drift, conservative-dissipative drift, and feedback linearizable,dynamics but the dynamics in practical systems have no exploitable structures. We envision that the future directions in Density Prediction will go hand in hand with solving more general trajectory dynamics.

7.3 Machine Learning

We believe our results can be generalized to a broader family of second-order dynamics, including Nesterov's method. Due to the time-dependent damping coefficient γ_t, we need to use a different Lyapunov functional. There are several

possible directions to investigate quantitative results. We are concerned with the study of the rate of convergence of ρ_t and we have some initial results on proving exponential convergence towards the equilibrium in weighted Wasserstein Distance.

Finally, the question of convergence for the noiseless second-order dynamics remains unsettled and requires further investigation. In general, without diffusion, there may exist stationary points that are not global minimizers (even in the mean field limit). However, one can hope to prove, under suitable assumptions on F, that such stationary points are repulsive, as was done in [65] for gradient flow.

Appendix A

A.1 Derivation of Fokker-Planck Equation

Consider the stochastic differential equation

$$\mathrm{d}\boldsymbol{x}_t = \boldsymbol{f}(\boldsymbol{x}_t, t)\,\mathrm{d}t + \boldsymbol{g}(\boldsymbol{x}_t, t)\,\mathrm{d}\boldsymbol{w}_t \tag{A.1}$$

and let $\phi : \mathcal{X} \mapsto \mathbb{R}$ be a smooth function with compact support. Consider the time derivative

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[\phi(\boldsymbol{x}_t)] = \frac{d}{dt}\int_{\mathcal{X}} \phi(\boldsymbol{x})\rho(x,t)\,\mathrm{d}\boldsymbol{x} = \int_{\mathcal{X}} \phi(\boldsymbol{x})\frac{\partial u}{\partial t}\,\mathrm{d}\boldsymbol{x} \tag{A.2}$$

On the other hand, by Ito's formula we have

$$\begin{aligned}
\mathrm{d}\mathbb{E}[\phi(\boldsymbol{x}_t)] &= \mathbb{E}[\mathrm{d}\phi(\boldsymbol{x}_t)] && \text{(A.3)}\\
&= \mathbb{E}\left[\left(\langle \boldsymbol{f}, \phi\rangle + \frac{1}{2}\langle \boldsymbol{g}\boldsymbol{g}^\top, \mathrm{Hess}(\phi)\rangle\right)\mathrm{d}t + \langle \nabla\phi, \boldsymbol{g}\,\mathrm{d}\boldsymbol{w}\rangle\right] && \text{(A.4)}\\
&= \left(\mathbb{E}\left[(\langle \boldsymbol{f}, \phi\rangle\right] + \frac{1}{2}\mathbb{E}\left[\langle \boldsymbol{g}\boldsymbol{g}^\top, \mathrm{Hess}(\phi)\rangle\right]\right)\,\mathrm{d}t + \mathbb{E}\left[\langle \nabla\phi, \boldsymbol{g}\,\mathrm{d}\boldsymbol{w}\rangle\right]. && \text{(A.5)}
\end{aligned}$$

Observe that the last term equals zero since $\mathrm{d}\boldsymbol{w}$ is independent of of $\boldsymbol{x}_t$ so the expectation becomes

$$\mathbb{E}\left[\langle \nabla\phi, \boldsymbol{g}\,\mathrm{d}\boldsymbol{w}\rangle\right] = \mathbb{E}\left[\nabla\phi(\boldsymbol{x}_t)^\top \boldsymbol{g}\right]\mathbb{E}[\mathrm{d}\boldsymbol{w}] = 0 \tag{A.6}$$

where $\mathbb{E}[\mathrm{d}\boldsymbol{w}] = 0$ by Gaussian property of the Wiener process. "Dividing by $\mathrm{d}t$", on both sides, we get

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[\phi(\boldsymbol{x}_t)] = \mathbb{E}\left[\langle \boldsymbol{f}, \phi\rangle\right] + \frac{1}{2}\mathbb{E}\left[\langle \boldsymbol{g}\boldsymbol{g}^\top, \mathrm{Hess}(\phi)\rangle\right]. \tag{A.7}$$

Now using the integration by parts formula

$$\int u(\nabla \cdot v)\,\mathrm{d}\boldsymbol{x} = -\int \langle \nabla u, v\rangle\,\mathrm{d}\boldsymbol{x}. \tag{A.8}$$

we obtain

$$\mathbb{E}\left[\langle \boldsymbol{f}, \phi\rangle\right] = \int_{\mathcal{X}} \langle \boldsymbol{f}, \phi\rangle \rho(x,t)\,\mathrm{d}\boldsymbol{x} = -\int_{\mathcal{X}} \nabla\cdot(\rho\boldsymbol{f})\phi(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}. \tag{A.9}$$

Applying the integration parts formula twice yields

$$\mathbb{E}\left[\langle \boldsymbol{g}\boldsymbol{g}^\top, \mathrm{Hess}(\phi)\rangle\right] = -\int_{\mathcal{X}} \sum_{i,j=1}^{n} (\rho\boldsymbol{g}\boldsymbol{g}^\top)_{ij} \frac{\partial^2\phi}{\partial x_i \partial x_j}\,\mathrm{d}\boldsymbol{x} \tag{A.10}$$

$$= -\sum_{i,j=1}^{n} \int_{\mathcal{X}} \frac{\partial}{\partial x_i}(\rho\boldsymbol{g}\boldsymbol{g}^\top)_{ij} \frac{\partial u}{\partial x_j}\,\mathrm{d}\boldsymbol{x} \tag{A.11}$$

$$= \sum_{i,j=1}^{n} \int_{\mathcal{X}} \frac{\partial}{\partial x_i \partial x_j}(\rho\boldsymbol{g}\boldsymbol{g}^\top)_{ij}\phi(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}. \tag{A.12}$$

By equation the two expressions for the $\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[\phi(\boldsymbol{x}_t)]$, we obtain

$$\int_{\mathcal{X}} \left(\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{f}) + \frac{1}{2} \sum_{i,j=1}^{n} \frac{\partial}{\partial x_i \partial x_j} (\rho \boldsymbol{g}\boldsymbol{g}^\top)_{ij} \right) \phi(\boldsymbol{x}) \, \mathrm{d}\boldsymbol{x} = 0 \tag{A.13}$$

and since ϕ is arbitrary, the evolution equation for the pdf ρ is given by

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \boldsymbol{f}) + \frac{1}{2} \sum_{i,j=1}^{n} \frac{\partial}{\partial x_i \partial x_j} (\rho \boldsymbol{g}\boldsymbol{g}^\top)_{ij} \tag{A.14}$$

A.2 Derivation of (2.28):

Noting that scaling and translation by constants do not alter the outer argmin in (2.27), we rewrite the same as

$$\begin{aligned} \boldsymbol{\varrho}_k = h \, arg\,\underset{\boldsymbol{\varrho}}{min} \Big\{ \underset{\boldsymbol{M} \in \Pi(\boldsymbol{\varrho}_{k-1}, \boldsymbol{\varrho})}{min} & \frac{1}{2h} \langle \boldsymbol{C}_k, \boldsymbol{M} \rangle + \frac{\epsilon}{h} H(\boldsymbol{M}) \\ & - \frac{\epsilon}{h} \mathbf{1}^\top \boldsymbol{M} \mathbf{1} + F(\boldsymbol{\varrho}) \Big\}, \end{aligned} \tag{A.15}$$

since $\mathbf{1}^\top \boldsymbol{M} \mathbf{1} = 1$. The Lagrangian $\mathcal{L}$ associated with the inner minimization in (A.15) is given by

$$\begin{aligned} \mathcal{L} = & \frac{1}{2h} \langle \boldsymbol{C}_k, \boldsymbol{M} \rangle + \frac{\epsilon}{h} H(\boldsymbol{M}) - \frac{\epsilon}{h} \mathbf{1}^\top \boldsymbol{M} \mathbf{1} + F(\boldsymbol{\varrho}) \\ & + \langle \boldsymbol{\lambda}_0, \boldsymbol{M}\mathbf{1} - \boldsymbol{\varrho}_{k-1} \rangle + \langle \boldsymbol{\lambda}_1, \boldsymbol{M}^\top \mathbf{1} - \boldsymbol{\varrho} \rangle. \end{aligned} \tag{A.16}$$

Setting the derivative of (A.16) w.r.t. the (i,j)-th element of $\boldsymbol{M}$ equal to zero, followed by algebraic simplification yields (2.28).

A.3 Proof of Proposition 5

We start by rewriting the Lagrangian (3.17) as

$$\begin{aligned}\mathcal{L}(\rho,\boldsymbol{u},\psi) &= \int_0^1 \int_{\mathbb{R}^n} \frac{1}{2}\|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2\,\rho(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}t \\ &\quad \underbrace{+\int_0^1 \int_{\mathbb{R}^n} \psi(\boldsymbol{x},t)\frac{\partial \rho}{\partial t}\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}t}_{\text{term 1}} \\ &\underbrace{+\int_0^1\int_{\mathbb{R}^n}\psi(\boldsymbol{x},t)\Big\{\nabla\cdot(\rho\,(\boldsymbol{f}+\boldsymbol{B}(t)\boldsymbol{u})) - \epsilon\mathbf{1}^\top(\boldsymbol{D}(t)\odot\mathbf{Hess}\,(\rho))\,\mathbf{1}\Big\}\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}t}_{\text{term 2}}.\end{aligned} \tag{A.17}$$

Next, we switch the order of integration, perform integration by parts w.r.t. t in term 1, and integration by parts w.r.t. $\boldsymbol{x}$ in term 2. Assuming the limits for $\|\,\boldsymbol{x}\,\|_2 \to \infty$ are zero, the Lagrangian (A.17) then simplifies to

$$\begin{aligned}\mathcal{L}(\rho,\boldsymbol{u},\psi) = \int_0^1\int_{\mathbb{R}^n}\Big(&\frac{1}{2}\|\boldsymbol{u}(\boldsymbol{x},t)\|_2^2 - \frac{\partial\psi}{\partial t} \\ &- \langle\nabla\psi, \boldsymbol{f}+\boldsymbol{B}(t)\boldsymbol{u}\rangle - \epsilon\langle\boldsymbol{D}(t),\mathbf{Hess}\,(\psi)\rangle\Big)\rho(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}t,\end{aligned} \tag{A.18}$$

where we have used the (two-fold) integration by parts:

$$\begin{aligned}&\int_{\mathbb{R}^n}\Big\{\mathbf{1}^\top(\boldsymbol{D}(t)\odot\mathbf{Hess}\,(\rho))\,\mathbf{1}\Big\}\,\psi(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x} \\ &= \sum_{i,j=1}^n \int_{\mathbb{R}^n} \frac{\partial^2\left(\boldsymbol{D}(t)\rho(\boldsymbol{x},t)\right)_{ij}}{\partial x_i\partial x_j}\,\psi(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x} \\ &= -\sum_{i,j=1}^n \int_{\mathbb{R}^n} \frac{\partial\left(\boldsymbol{D}(t)\rho\right)_{ij}}{\partial x_j}\frac{\partial\psi}{\partial x_i}\,\mathrm{d}\boldsymbol{x} \\ &= \sum_{i,j=1}^n \int_{\mathbb{R}^n} \left(\boldsymbol{D}(t)\rho\right)_{ij}\frac{\partial^2\psi}{\partial x_j\partial x_i}\,\mathrm{d}\boldsymbol{x} \\ &= \int_{\mathbb{R}^n}\sum_{i,j=1}^n \left(\boldsymbol{D}(t)\rho\right)_{ij}\frac{\partial^2\psi}{\partial x_j\partial x_i}\,\mathrm{d}\boldsymbol{x} \\ &= \int_{\mathbb{R}^n}\langle\boldsymbol{D}(t),\mathbf{Hess}\,(\psi)\rangle\rho(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x}.\end{aligned}$$

Pointwise minimization of (A.18) w.r.t. $\boldsymbol{u}$ for a fixed PDF ρ, gives

$$\boldsymbol{u}^{\text{opt}}(\boldsymbol{x},t) = \boldsymbol{B}(t)^{\top}\nabla\psi(\boldsymbol{x},t). \tag{A.19}$$

Substituting (A.19) back into (A.18), and equating the resulting expression to zero, we get the dynamic programming equation

$$\begin{aligned}\int_0^1 \int_{\mathbb{R}^n} \Big(-\frac{\partial\psi}{\partial t} - \frac{1}{2}\|\nabla\boldsymbol{B}(t)^{\top}\psi\|_2^2 - \langle\nabla\psi, \boldsymbol{f}\rangle \\ - \ \epsilon\langle\boldsymbol{D}(t), \mathbf{Hess}\,(\psi)\rangle\Big)\rho(\boldsymbol{x},t)\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}t = 0.\end{aligned} \tag{A.20}$$

For (A.20) to hold for arbitrary ρ, we must have

$$\frac{\partial\psi}{\partial t} + \frac{1}{2}\|\boldsymbol{B}(t)^{\top}\nabla\psi\|_2^2 + \langle\nabla\psi, \boldsymbol{f}\rangle = -\epsilon\langle\boldsymbol{D}(t), \mathbf{Hess}\,(\psi)\rangle, \tag{A.21}$$

which is the HJB PDE (3.19a). The associated FPK PDE (3.19b) results from substituting (A.19) into (3.15b). The boundary conditions (3.20) follows from (3.15c). This completes the proof. ■

A.4 Proof of Theorem 6

In (3.4a), taking the gradient of φ w.r.t. $\boldsymbol{x}$, we get

$$\nabla\varphi = \frac{1}{2\epsilon}\exp\left(\frac{\psi}{2\epsilon}\right)\nabla\psi. \tag{A.22}$$

Furthermore,

$$
\begin{aligned}
&-\epsilon\langle \boldsymbol{D}(t), \mathbf{Hess}\,(\varphi)\rangle = -\epsilon \sum_{i,j=1}^{n} (\boldsymbol{D}(t))_{ij} \frac{\partial^2}{\partial x_i \partial x_j} \exp\left(\psi/2\epsilon\right) \\
&= -\epsilon \frac{\exp\left(\psi/2\epsilon\right)}{2\epsilon} \Big\{ \sum_{i,j=1}^{n} (\boldsymbol{D}(t))_{ij} \left(\frac{\partial^2 \psi}{\partial x_i \partial x_j} + \frac{1}{2\epsilon} \frac{\partial \psi}{\partial x_i} \frac{\partial \psi}{\partial x_j} \right) \Big\} \\
&= \frac{1}{2\epsilon} \exp\left(\frac{\psi}{2\epsilon}\right) \Big\{ -\epsilon \langle \boldsymbol{D}(t), \mathbf{Hess}\,(\psi)\rangle - \frac{1}{2} \|\boldsymbol{B}(t)^\top \nabla \psi\|_2^2 \Big\}. \qquad \text{(A.23)}
\end{aligned}
$$

From (3.4a), taking partial derivative of φ w.r.t. t gives

$$
\begin{aligned}
&\frac{\partial \varphi}{\partial t} = \frac{1}{2\epsilon} \exp\left(\frac{\psi}{2\epsilon}\right) \frac{\partial \psi}{\partial t} \\
&\overset{\text{(3.19a)}}{=} \frac{1}{2\epsilon} \exp\left(\frac{\psi}{2\epsilon}\right) \Big\{ -\frac{1}{2} \|\boldsymbol{B}(t)^\top \nabla \psi\|_2^2 - \langle \nabla \psi, \boldsymbol{f}\rangle - \epsilon \langle \boldsymbol{D}(t), \mathbf{Hess}\,(\psi)\rangle \Big\} \\
&\overset{\text{(A.22),(A.23)}}{=} -\langle \nabla \varphi, \boldsymbol{f}\rangle - \epsilon \langle \boldsymbol{D}(t), \mathbf{Hess}\,(\varphi)\rangle =: L_{\mathrm{BK}} \varphi,
\end{aligned}
$$

i.e., $\varphi(\boldsymbol{x}, t)$ satisfies the backward Kolmogorov equation (3.22a).

To demonstrate that $\hat{\varphi}(\boldsymbol{x}, t)$ satisfies the forward Kolmogorov equation (4.4a), from (3.4b) we compute

$$
\nabla \hat{\varphi} = \exp\left(-\frac{\psi}{2\epsilon}\right) \left(-\frac{\rho^{\mathrm{opt}}}{2\epsilon} \nabla \psi + \nabla \rho^{\mathrm{opt}} \right). \qquad \text{(A.24)}
$$

Consequently,

$$
\begin{aligned}
&\mathbf{1}^\top \left(\boldsymbol{D}(t) \odot \mathbf{Hess}\left(\hat{\varphi}\right)\right)\mathbf{1} \\
&= \sum_{i,j=1}^{n} \left(\boldsymbol{D}(t)\right)_{ij} \frac{\partial^2 \hat{\varphi}}{\partial x_i x_j} \\
&\overset{(3.4b)}{=} \sum_{i,j=1}^{n} \left(\boldsymbol{D}(t)\right)_{ij} \Big\{ \exp\left(-\frac{\psi}{2\epsilon}\right) \frac{\partial^2 \rho^{\mathrm{opt}}}{\partial x_i \partial x_j} - \frac{1}{2\epsilon} \frac{\partial \rho^{\mathrm{opt}}}{\partial x_j} \exp\left(-\frac{\psi}{2\epsilon}\right) \frac{\partial \psi}{\partial x_i} \\
&\qquad - \frac{\partial^2 \psi}{\partial x_i \partial x_j} \frac{\rho^{\mathrm{opt}}}{2\epsilon} \exp\left(-\frac{\psi}{2\epsilon}\right) - \frac{\partial \psi}{\partial x_j} \left(\frac{1}{2\epsilon} \frac{\partial \rho^{\mathrm{opt}}}{\partial x_i} \exp\left(-\frac{\psi}{2\epsilon}\right) \right. \\
&\qquad \left. + \frac{\rho^{\mathrm{opt}}}{4\epsilon^2} \exp\left(-\frac{\psi}{2\epsilon}\right) \frac{\partial \psi}{\partial x_i} \right) \Big\} \\
&= \exp\left(-\frac{\psi}{2\epsilon}\right) \Big\{ \epsilon \mathbf{1}^\top \Big(\boldsymbol{D}(t) \odot \mathbf{Hess}(\rho^{\mathrm{opt}})\Big)\mathbf{1} - \langle \nabla \rho^{\mathrm{opt}}, \boldsymbol{D}(t) \nabla \psi \rangle \\
&\qquad - \frac{\rho^{\mathrm{opt}}}{2} \langle \boldsymbol{D}(t), \mathbf{Hess}(\psi) \rangle + \frac{\rho^{\mathrm{opt}}}{4\epsilon} \langle \nabla \psi, \boldsymbol{D}(t) \nabla \psi \rangle \Big\}.
\end{aligned} \tag{A.25}
$$

Also, we have

$$
-\nabla \cdot \Big(\rho^{\mathrm{opt}} \boldsymbol{D}(t) \nabla \psi \Big) = -\langle \nabla \rho^{\mathrm{opt}}, \boldsymbol{D}(t) \nabla \psi \rangle - \rho^{\mathrm{opt}} \langle \boldsymbol{D}(t), \mathbf{Hess}(\psi) \rangle, \tag{A.26}
$$

and

$$
\begin{aligned}
\nabla \cdot (\hat{\varphi} \boldsymbol{f}) &= \langle \nabla \hat{\varphi}, \boldsymbol{f} \rangle + \hat{\varphi} \nabla \cdot \boldsymbol{f} \\
&\overset{(5.50),(A.24)}{=} \exp\left(-\frac{\psi}{2\epsilon}\right) \Big(- \frac{\rho^{\mathrm{opt}}}{2\epsilon} \langle \nabla \psi, \boldsymbol{f} \rangle + \langle \nabla \rho^{\mathrm{opt}}, \boldsymbol{f} \rangle \\
&\qquad + \rho^{\mathrm{opt}} \nabla \cdot \boldsymbol{f} \Big),
\end{aligned} \tag{A.27}
$$

wherein we used the chain rule for the divergence operator. Taking partial deriva-

tive of (3.4b) w.r.t to t, we then get

$$\begin{aligned}
\frac{\partial\hat{\varphi}}{\partial t} &= \exp\left(-\frac{\psi}{2\epsilon}\right)\left(\frac{\partial\rho^{\mathrm{opt}}}{\partial t} - \frac{\rho^{\mathrm{opt}}}{2\epsilon}\frac{\partial\psi}{\partial t}\right) \\
&\overset{(3.19)}{=} \exp\left(-\frac{\psi}{2\epsilon}\right)\Big[\Big\{-\nabla\cdot(\rho^{\mathrm{opt}}\boldsymbol{f}) - \nabla\cdot(\rho\boldsymbol{D}(t)\nabla\psi) \\
&\quad + \epsilon\mathbf{1}^\top\left(\boldsymbol{D}(t)\odot\mathbf{Hess}\left(\rho^{\mathrm{opt}}\right)\right)\mathbf{1}\Big\} - \frac{\rho^{\mathrm{opt}}}{2\epsilon}\Big\{-\frac{1}{2}\parallel\boldsymbol{B}(t)^\top\nabla\psi\parallel_2^2 \\
&\quad -\langle\nabla\psi,\boldsymbol{f}\rangle - \epsilon\langle\boldsymbol{D}(t),\mathbf{Hess}(\psi)\rangle\Big\}\Big] \\
&\overset{(A.25),(A.26),(A.27)}{=} -\nabla\cdot(\hat{\varphi}\boldsymbol{f}) + \epsilon\Delta\hat{\varphi},
\end{aligned} \tag{A.28}$$

i.e., $\hat{\varphi}(\boldsymbol{x},t)$ satisfies (4.4a).

Combining (5.50) with (3.20), we obtain the boundary condition (3.23). Finally, combining (3.4a) with (A.19), yields the optimal control $\boldsymbol{u}^{\mathrm{opt}}(\boldsymbol{x},t) = 2\epsilon\boldsymbol{B}(t)^\top\nabla\log\varphi$.

■

A.5 Proof of Theorem 7

First, notice that (3.31) follows from combining $s = 1-t$ and (3.30a). Next, using

$$p(\boldsymbol{x},s) := q(\boldsymbol{x},s)\exp\left(-V(\boldsymbol{x})/\epsilon\right), \tag{A.29}$$

we find that

$$\nabla q = \exp\left(V/\epsilon\right)\left(\nabla p + \frac{p}{\epsilon}\nabla V\right). \tag{A.30}$$

Applying divergence operator (w.r.t. $\boldsymbol{x}$) to both sides of (A.30) yields

$$\begin{aligned}\Delta q &= \left\langle \nabla p + \frac{p}{\epsilon}\nabla V, \nabla \exp\left(V/\epsilon\right) \right\rangle + \exp\left(V/\epsilon\right) \nabla \cdot \left(\nabla p + \frac{p}{\epsilon}\nabla V\right) \\ &= \exp\left(V/\epsilon\right) \left(\frac{p}{\epsilon^2}\|\nabla V\|_2^2 + \Delta p + \frac{2}{\epsilon}\langle \nabla p, \nabla V\rangle + \frac{p}{\epsilon}\Delta V\right). \end{aligned} \tag{A.31}$$

From (A.29), we have

$$\begin{aligned}\frac{\partial p}{\partial s} &= \exp\left(-V/\epsilon\right) \frac{\partial q}{\partial s} \\ &\overset{(3.31)}{=} \exp\left(-V/\epsilon\right) \left(-\langle \nabla q, \nabla V\rangle + \epsilon \Delta q\right) \\ &\overset{(A.30),(A.31)}{=} -\left\langle \nabla p + \frac{p}{\epsilon}\nabla V, \nabla V \right\rangle + \epsilon \left(\frac{p}{\epsilon^2}\|\nabla V\|_2^2 + \Delta p \right. \\ &\qquad \left. + \frac{2}{\epsilon}\langle \nabla p, \nabla V\rangle + \frac{p}{\epsilon}\Delta V\right) \\ &= \langle \nabla p, \nabla V\rangle + \epsilon \Delta p + p \Delta V \\ &= \nabla \cdot (p \nabla V) + \epsilon \Delta p, \end{aligned} \tag{A.32}$$

which is indeed the PDE in (3.32). Setting $s = 0$ in (A.29) recovers the initial condition in (3.32). This completes the proof. ■

A.6 Proof of Theorem 8

Notice that (3.38) follows by combining $s := 1 - t$, and (3.37a). To derive (3.39), we start by taking the gradient of

$$\widetilde{p}(\boldsymbol{\xi}, -\boldsymbol{\eta}, s) = q(\boldsymbol{\xi}, \boldsymbol{\eta}, s) \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right) \tag{A.33}$$

w.r.t. $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$, respectively, to obtain

$$\nabla_{\boldsymbol{\xi}} q = \exp\left(\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right)\left(\nabla_{\boldsymbol{\xi}}\widetilde{p} + \frac{\widetilde{p}}{\epsilon}\nabla_{\boldsymbol{\xi}} V\right), \tag{A.34a}$$

$$\nabla_{\boldsymbol{\eta}} q = \exp\left(\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right)\left(-\nabla_{-\boldsymbol{\eta}}\widetilde{p} + \frac{1}{\epsilon}\widetilde{p}\,\boldsymbol{\eta}\right). \tag{A.34b}$$

Applying divergence operator w.r.t. $\boldsymbol{\eta}$ on both sides of (A.34b) yields

$$\begin{aligned}\Delta_{\boldsymbol{\eta}} q = \exp\left(\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right)\Big(&\Delta_{-\boldsymbol{\eta}}\widetilde{p} - \frac{2}{\epsilon}\langle\nabla_{-\boldsymbol{\eta}}\widetilde{p}, \boldsymbol{\eta}\rangle \\ &+ \frac{\widetilde{p}}{\epsilon} + \frac{\widetilde{p}}{\epsilon}\|\boldsymbol{\eta}\|_2^2\Big).\end{aligned} \tag{A.35}$$

Thus, we have

$$\begin{aligned}\frac{\partial p}{\partial s} = \frac{\partial \widetilde{p}}{\partial s} &\overset{\text{(A.33)}}{=} \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right)\frac{\partial q}{\partial s} \\ &\overset{\text{(3.38),(A.34),(A.35)}}{=} \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right)\Big\{\langle\boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}}\widetilde{p} + \frac{\widetilde{p}}{\epsilon}\nabla_{\boldsymbol{\xi}} V\rangle \\ &\quad - \langle\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) - \kappa\boldsymbol{\eta}, -\nabla_{-\boldsymbol{\eta}}\widetilde{p} + \frac{1}{\epsilon}\widetilde{p}\boldsymbol{\eta}\rangle \\ &\quad + \Delta_{-\boldsymbol{\eta}}\widetilde{p} - \frac{2}{\epsilon}\langle\nabla_{-\boldsymbol{\eta}}\widetilde{p}, \boldsymbol{\eta}\rangle\frac{\widetilde{p}}{\epsilon} + \frac{\widetilde{p}}{\epsilon}\|\boldsymbol{\eta}\|_2^2\Big\} \\ &= \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\eta}\|_2^2 + V(\boldsymbol{\xi})\right)\right)\Big\{\langle\boldsymbol{\eta}, \nabla_{\boldsymbol{\xi}}\widetilde{p}\rangle - \kappa\langle\nabla_{-\boldsymbol{\eta}}\widetilde{p}, \eta\rangle \\ &\quad + \langle\nabla_{-\boldsymbol{\eta}}\widetilde{p}, \nabla_{\boldsymbol{\xi}} V\rangle + \kappa\widetilde{p} + \epsilon\kappa\Delta_{-\boldsymbol{\eta}}\widetilde{p}\Big\} \\ &= \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\vartheta}\|_2^2 + V(\boldsymbol{\xi})\right)\right)\Big\{\langle-\boldsymbol{\vartheta}, \nabla_{\boldsymbol{\xi}} p\rangle + \kappa\langle\nabla_{\boldsymbol{\vartheta}} p, \vartheta\rangle \\ &\quad + \langle\nabla_{\boldsymbol{\vartheta}} p, \nabla_{\boldsymbol{\xi}} V\rangle + \kappa p + \epsilon\kappa\Delta_{\boldsymbol{\vartheta}} p\Big\} \\ &= \exp\left(-\frac{1}{\epsilon}\left(\frac{1}{2}\|\boldsymbol{\vartheta}\|_2^2 + V(\boldsymbol{\xi})\right)\right)\Big\{ - \langle\boldsymbol{\vartheta}, \nabla_{\boldsymbol{\xi}} p\rangle \\ &\quad + \nabla_{\boldsymbol{\vartheta}}\cdot\left(p\left(\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\xi}) + \kappa\boldsymbol{\vartheta}\right)\right) + \epsilon\kappa\Delta_{\boldsymbol{\vartheta}} p\Big\},\end{aligned}$$

which is the PDE in (3.39). Setting $s = 0$ in (A.33) recovers the initial condition in (3.39) . ■

A.7 Regularity of the Transition Densities for (3.33) and (3.40)

In this Appendix, we point out that the transition probability densities $K(s, \boldsymbol{y}, t, \boldsymbol{x})$ for (3.33) and (3.40), are indeed positive and continuous in $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ for all $s < t$. First, recall that the transition densities themselves solve the same PDEs as in (3.33) and (3.40), with initial condition $\lim_{t \downarrow s} K(s, \boldsymbol{y}, t, \boldsymbol{x}) = \delta(\boldsymbol{x} - \boldsymbol{y})$. From the maximum principle for parabolic PDEs, it follows that the transient solutions of (3.33) and (3.40) are positive as long as the initial conditions are positive. The continuity, in the gradient drift case (3.33), is standard assuming $V \in C^2(\mathbb{R}^n)$; see e.g., [202, Ch. 1.2]. For the degenerate diffusion case (3.40), the situation is more subtle: a result from Villani [213, Theorem 7] ensures the continuity (w.r.t. both state and time) of the transient solutions, under the assumptions on $V(\cdot)$ stated in Section 3.4.2, viz. $V \in C^2(\mathbb{R}^m)$, $\inf V > -\infty$, and $\mathbf{Hess}(V)$ uniformly lower bounded.

A.8 Proof of Theorem 1

The necessary conditions for optimality (4.4) can be deduced using the Lagrange multiplier theorem in Banach spaces; see [223, Ch.4.14, Proposition 1]. This theorem allows us set up an augmented Lagrangian associated with (4.3) and perform pointwise minimization to derive (4.4).

To apply this in our context, define the function spaces

$$\begin{aligned}
\mathcal{P}_{01} &:= \{\rho(t,\cdot) \in \mathcal{P}_2\left(\overline{\mathcal{X}}\right) \mid \rho(0,\cdot) = \rho_0, \rho(1,\cdot) = \rho_1\},\\
\widetilde{\mathcal{P}}_{01} &:= \mathcal{P}_{01} \cap L^2(H^1([0,1];\overline{\mathcal{X}})) \cap \dot{H}^1\left(\left(H^1([0,1];\overline{\mathcal{X}})\right)^*\right),\\
X &:= \widetilde{\mathcal{P}}_{01} \times L^2([0,1]\times\overline{\mathcal{X}}), \quad Y := L^2(H^{-1}([0,1];\overline{\mathcal{X}})),
\end{aligned}$$

where $[0,1]$ denotes the time interval, and $L^2(\cdot)$ denotes the space of square integrable functions. The notation $L^2([0,1];H^1(\overline{\mathcal{X}}))$ stands for the Sobolev space of functions having first order weak derivatives w.r.t. $\boldsymbol{x}_t^{\boldsymbol{u}} \in \overline{\mathcal{X}}$, and finite L^2 norms w.r.t. $t \in [0,1]$. Furthermore, $\dot{H}^1\left([0,1];\left(H^1(\overline{\mathcal{X}})\right)^*\right) := \{\phi(t,\cdot) \in L^2([0,1]) \mid \frac{\partial\phi}{\partial t} \in L^2([0,1]), \phi \in \left(H^1(\overline{\mathcal{X}})\right)^*\}$, wherein $\left(H^1(\overline{\mathcal{X}})\right)^*$ denotes the dual space of the Sobolev space $H^1(\overline{\mathcal{X}})$. We denote the dual space of $\widetilde{\mathcal{P}}_{01}$ as $\widetilde{\mathcal{P}}_{01}^*$. In the definition of Y, the notation $H^{-1}\left(\overline{\mathcal{X}}\right)$ stands for the space of all linear functionals on $H_0^1(\mathcal{X}) := \{\phi \in H^1(\mathcal{X}),$ and vanishes on $\partial\mathcal{X}\}$. Then, in (4.3a), the objective functional $F : X \mapsto \mathbb{R}$, and is given by

$$F(\rho,\boldsymbol{u}) := \int_{\overline{\mathcal{X}}}\int_0^1 \frac{1}{2}\|\boldsymbol{u}(t,\boldsymbol{x}_t^{\boldsymbol{u}})\|_2^2\,\rho(t,\boldsymbol{x}_t^{\boldsymbol{u}})\,\mathrm{d}t\,\mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}}. \tag{A.36}$$

The constraint is a mapping $G : X \mapsto Y$ given by

$$\begin{aligned}
G(\rho,\boldsymbol{u})(\psi) :=& \int_{\overline{\mathcal{X}}}\psi(1,\boldsymbol{x}_t^{\boldsymbol{u}})\rho(1,\boldsymbol{x}_t^{\boldsymbol{u}})\mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}} - \int_{\overline{\mathcal{X}}}\psi(0,\boldsymbol{x}_t^{\boldsymbol{u}})\rho(0,\boldsymbol{x}_t^{\boldsymbol{u}})\mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}}\\
&-\int_{\overline{\mathcal{X}}}\int_0^1 \frac{\partial\psi}{\partial t}\rho\,\mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}}\,\mathrm{d}t + \int_{\overline{\mathcal{X}}}\int_0^1 \psi(\nabla\cdot(\rho\boldsymbol{u}+\boldsymbol{f}) - \theta\Delta\rho)\,\mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}}\,\mathrm{d}t,
\end{aligned} \tag{A.37}$$

where we used (4.3c) so that the boundary terms vanish in the integration by parts. Following [3, p. 112-114], one can show that $G'_\rho(\rho,\boldsymbol{u})$ and $G'_{\boldsymbol{u}}(\rho,\boldsymbol{u})$ (where $'$ denotes derivative w.r.t. the subscripted variable) are surjective, and hence by [223, Ch.4.14, Proposition 1], there exists $\psi \in Y^* = L^2([0,1];H_0^1(\overline{\mathcal{X}}))$. This

result allows us to perform pointwise minimization of the augmented Lagrangian

$$\mathscr{L}(\rho, \boldsymbol{u}, \psi) := \int_0^1 \int_{\mathcal{X}} \frac{1}{2}\|\boldsymbol{u}(t, \boldsymbol{x}_t^{\boldsymbol{u}})\|_2^2 \, \rho(t, \boldsymbol{x}_t^{\boldsymbol{u}}) \, \mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}} \mathrm{d}t + \underbrace{\int_{\mathcal{X}} \int_0^1 \psi \frac{\partial \rho}{\partial t} \mathrm{d}t \mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}}}_{\text{term1}} + \underbrace{\int_0^1 \int_{\mathcal{X}} \psi \left(\nabla \cdot (\rho \boldsymbol{u} + \boldsymbol{f}) - \theta \Delta \rho\right) \, \mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}} \mathrm{d}t}_{\text{term2}}.$$

By performing integration by parts in t, term 1 becomes

$$\int_{\overline{\mathcal{X}}} \left(\underbrace{\psi(1, \boldsymbol{x}_t^{\boldsymbol{u}}) \rho(1, \boldsymbol{x}_t^{\boldsymbol{u}}) - \psi(0, \boldsymbol{x}_t^{\boldsymbol{u}}) \rho(0, \boldsymbol{x}_t^{\boldsymbol{u}})}_{\text{constant w.r.t. } (\rho, \boldsymbol{u})} - \int_0^1 \frac{\partial \psi}{\partial t} \rho \mathrm{d}t \right) \mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}}.$$

For term 2, we perform integration by parts w.r.t $\boldsymbol{x}_t^{\boldsymbol{u}}$, impose the boundary condition (4.3c), and thereby deduce that $\mathscr{L}$ (up to an additive constant) equals

$$\int_0^1 \int_{\overline{\mathcal{X}}} \left(\frac{1}{2}\|\boldsymbol{u}\|_2^2 - \frac{\partial \psi}{\partial t} - \langle \nabla \psi, \boldsymbol{u} + \boldsymbol{f} \rangle - \theta \Delta \psi \right) \rho \, \mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}} \, \mathrm{d}t. \tag{A.38}$$

Pointwise minimization of (A.38) w.r.t $\boldsymbol{u}$ while fixing ρ, gives the optimal control (5.32). Substituting (5.32) back into (A.38) and equating the resulting expression to zero results in the dynamic programming equation

$$\int_0^1 \int_{\overline{\mathcal{X}}} \left(-\frac{\partial \psi}{\partial t} - \frac{1}{2}\|\nabla \psi\|^2 - \langle \nabla \psi, \boldsymbol{f} \rangle - \theta \Delta \psi \right) \rho \, \mathrm{d}\boldsymbol{x}_t^{\boldsymbol{u}} \, \mathrm{d}t = 0.$$

Since the above holds for arbitrary ρ, we must have

$$\frac{\partial \psi}{\partial t} + \frac{1}{2}\|\nabla \psi\|^2 + \langle \nabla \psi, \boldsymbol{f} \rangle + \theta \Delta \psi = 0,$$

which is indeed the HJB PDE (4.4b). Substituting (5.32) in (4.3a) yields the FPK PDE (4.4a).

The Neumann boundary condition (4.6a) follows directly (see [152]). The endpoint

conditions (4.6b) follow from (4.3d). The Robin boundary condition (4.6c) is obtained by combining (5.32) with (4.3c). ■

A.9 Proof of Theorem 2

The system of linear PDEs (4.8) are obtained via straightforward but tedious computation detailed in [44, Appendix B]. The boundary conditions (4.9a) follow by setting $t = 0, 1$ in (4.7). To derive (4.9b), evaluate (4.7a) at a boundary point $\boldsymbol{x}_{\mathrm{bdy}} \in \partial\mathcal{X}$. In the resulting expression, take the natural log to both sides and then take the gradient w.r.t. $\boldsymbol{x}_{\mathrm{bdy}}$, to get

$$\nabla\psi(t, \boldsymbol{x}_{\mathrm{bdy}}) = 2\theta \frac{\nabla\varphi(t, \boldsymbol{x}_{\mathrm{bdy}})}{\varphi(t, \boldsymbol{x}_{\mathrm{bdy}})}. \tag{A.39}$$

In both sides of (A.39), we take the inner product with the normal vector $\boldsymbol{\nu}(\boldsymbol{x}_{\mathrm{bdy}})$, and use (4.6a), to obtain $\langle \nabla\varphi, \boldsymbol{\nu}\rangle|_{\partial\mathcal{X}} = 0$, as in (4.9b). To deduce the second equality in (4.9b), we evaluate (4.7b) at $\boldsymbol{x}_{\mathrm{bdy}}$, and then as before, take the natural log followed by the gradient w.r.t. $\boldsymbol{x}_{\mathrm{bdy}}$, and invoke (4.6a) to arrive at

$$\frac{\langle \nabla\hat{\varphi}(t, \boldsymbol{x}_{\mathrm{bdy}}), \boldsymbol{\nu}(\boldsymbol{x}_{\mathrm{bdy}})\rangle}{\hat{\varphi}(t, \boldsymbol{x}_{\mathrm{bdy}})} = \frac{\langle \rho^{\mathrm{opt}}(t, \boldsymbol{x}_{\mathrm{bdy}}), \boldsymbol{\nu}(\boldsymbol{x}_{\mathrm{bdy}})\rangle}{\rho^{\mathrm{opt}}(t, \boldsymbol{x}_{\mathrm{bdy}})}. \tag{A.40}$$

Using (4.6a) again in (4.6c), the RHS of (A.40) simplifies to $\langle \boldsymbol{f}(t, \boldsymbol{x}_{\mathrm{bdy}}), \boldsymbol{\nu}(\boldsymbol{x}_{\mathrm{bdy}})\rangle/\theta$, and thus yields the second equality in (4.9b). Finally, (4.10) follows from (4.7) and (5.32). ■

A.10 Proof of Lemma 1

Proof of (i): To demonstrate continuity, it suffices to show that the infinite sum in (4.14) for $K_\theta(x, y, t = 1)$ converges uniformly on $[a, b] \times [a, b]$. For $k \in \mathbb{N}$, let

$$f_k(x, y) := \exp\left(-\frac{\theta\pi^2 k^2}{(b-a)^2}\right) \cos\left(\frac{k\pi(x-a)}{b-a}\right) \times \cos\left(\frac{k\pi(y-a)}{b-a}\right), \tag{A.41}$$

and notice that

$$|f_k(x, y)| \leq M_k \quad \text{for all} \quad (x, y) \in [a, b] \times [a, b], \tag{A.42}$$

where $M_k := \exp\left(-\theta\pi^2 k^2/(b-a)^2\right)$. Furthermore,

$$\lim_{k\to\infty} \left|\frac{M_{k+1}}{M_k}\right| = \lim_{k\to\infty} \exp\left(-\frac{\theta\pi^2(2k+1)}{(b-a)^2}\right) = 0. \tag{A.43}$$

By the ratio test [189, Ch. 3, Theorem 3.34], we then have

$$\sum_{k=1}^{\infty} M_k < \infty. \tag{A.44}$$

From (A.42) and (A.44), the Weierstrass M-test [189, Ch. 7, Theorem 7.10] implies that $\sum_{k=1}^{\infty} f_k(x, y)$ is uniformly convergent for all $(x, y) \in [a, b] \times [a, b]$, and the resulting sum must converge to a continuous function. Therefore, $K_\theta(x, y, t = 1)$ is continuous for all $(x, y) \in [a, b] \times [a, b]$.

Proof of (ii): To establish positivity, set $r := b - a, \tilde{x} := x - a, \tilde{y} := y - a$. Using

basic trigonometry and the Euler's identity, we find that

$$\begin{aligned}
&\frac{1}{r}+\frac{2}{r}\sum_{m=1}^{\infty}\exp\left(-\frac{\theta\pi^2m^2}{r^2}\right)\left[\cos\left(\frac{m\pi\tilde{x}}{r}\right)\cos\left(\frac{m\pi\tilde{y}}{r}\right)\right]\\
&=\sum_{m=-\infty}^{\infty}\frac{1}{2r}\exp\left(-\frac{\theta\pi^2m^2}{r^2}\right)\exp\left(\frac{im\pi(\tilde{x}+\tilde{y})}{r}\right)+\\
&\quad\sum_{m=-\infty}^{\infty}\frac{1}{2r}\exp\left(-\frac{\theta\pi^2m^2}{r^2}\right)\exp\left(\frac{im\pi(\tilde{x}-\tilde{y})}{r}\right).
\end{aligned}\tag{A.45}$$

Let

$$g(m):=\underbrace{\frac{1}{2r}\exp\left(-\frac{\theta\pi^2m^2}{r^2}\right)}_{=:g_1(m)}\underbrace{\exp\left(\frac{im\pi(\tilde{x}+\tilde{y})}{r}\right)}_{=:g_2(m)},\tag{A.46}$$

and denote the Fourier transforms of $g_1(m), g_2(m)$ as $\widehat{g}_1(\widehat{m}), \widehat{g}_2(\widehat{m})$, respectively. Notice that

$$\widehat{g}_1(\widehat{m})=\frac{1}{\sqrt{4\pi\theta}}\exp\left(-\frac{r^2\widehat{m}^2}{\theta}\right),\tag{A.47}$$

$$\widehat{g}_2(\widehat{m})=\delta\left(\widehat{m}-\frac{\tilde{x}+\tilde{y}}{2r}\right),\tag{A.48}$$

where $\delta(\cdot)$ denotes the Dirac delta, and hence by the convolution theorem, the Fourier transform of g is

$$\begin{aligned}
\widehat{g}(\widehat{m})&=\frac{1}{\sqrt{4\pi\theta}}\exp\left(-\frac{r^2}{\theta}\left(\widehat{m}-\frac{(\tilde{x}+\tilde{y})}{2r}\right)^2\right)\\
&=\frac{1}{\sqrt{4\pi\theta}}\exp\left(-\frac{(2\widehat{m}r-\tilde{x}-\tilde{y})^2}{4\theta}\right).
\end{aligned}\tag{A.49}$$

Invoking the Poisson summation formula [201, Ch. 4, Theorem 2.4], we deduce

$$\sum_{m=-\infty}^{\infty}g(m)=\sum_{\widehat{m}=-\infty}^{\infty}\widehat{g}(\widehat{m}),\tag{A.50}$$

implying that the infinite sum in (A.45) is equal to

$$\frac{1}{\sqrt{4\pi\theta}} \sum_{\widehat{m}=-\infty}^{\infty} \left[\exp\left(-\frac{(2\widehat{m}r - \tilde{x} - \tilde{y})^2}{4\theta} \right) + \exp\left(-\frac{(2\widehat{m}r - \tilde{x} + \tilde{y})^2}{4\theta} \right) \right], \tag{A.51}$$

which is obviously positive. Therefore, $K_\theta(x, y, t = 1)$ is positive for all $(x, y) \in [a, b] \times [a, b]$. ■

A.11 Proof of Proposition 1

Using Lemma 9 and that $[a, b]$ is a compact metric space, the hypotheses of [53, Proposition 4 and Theorem 8] are satisfied. Therefore, the solution pair (φ_1, φ_0) exists and is unique in the projective sense. Furthermore, the fixed point recursion being contractive in Hilbert's projective metric, converges (by contraction mapping theorem) to this pair. ■

A.12 Gradient in parameter space and differential in distribution space

Recall that the risk functional is $R : \mathcal{F} \to \mathbb{R}_+$. Let us define the following functions:

- The objective function in parameter space:

$$f(\theta_1, \ldots, \theta_n) := F\left(\frac{1}{n} \sum_{i=1}^{n} \delta_{\theta_i} \right) = R\left(\frac{1}{n} \sum_{i=1}^{n} \Psi(\theta_i) \right) + \frac{1}{n} \sum_{i=1}^{n} g(\theta_i).$$

- The objective functional in distribution space: $F(\mu) := R(\langle \Psi, \mu \rangle) + \langle g, \mu \rangle$.

By the chain rule, the gradient of f is

$$\nabla_{\theta_j} f(\theta_1, \ldots, \theta_n) = \frac{1}{n} \left\langle R'(\tfrac{1}{n} \textstyle\sum_{i=1}^n \Psi(\theta_i)), \Psi'(\theta_j) \right\rangle_{\mathcal{F}} + \frac{1}{n} \nabla g(\theta_j), \tag{A.52}$$

and the differential of F is

$$F'(\mu)(\cdot) = \langle R'(\langle \Psi, \mu \rangle), \Psi(\cdot) \rangle_{\mathcal{F}} + g(\cdot). \tag{A.53}$$

Identifying (A.52) and (A.53), we see that $n \nabla_{\theta_j} f(\theta_1, \ldots, \theta_n) = \nabla F'(\mu^n)(\theta_j)$, where $\mu^n = \frac{1}{n} \sum_{i=1}^n \delta_{\theta_i}$.

Observe that the gradient is scaled by n, which leads to nonlinear dynamics in the mean field limit. A different scaling can lead to simpler, linearized dynamics that are referred to as lazy training 67 or the kernel regime 123, 220. Our analysis is concerned with the fully non-linear regime.

The stochastic heavy ball dynamics in the parameter space is given by

$$\begin{cases} \dot{\theta}_i = r_i, \\ \dot{r}_i = -n \nabla_{\theta_i} f(\theta_1, \ldots, \theta_n) - \gamma r_i + \sqrt{2\gamma\beta^{-1}}\, \mathrm{d}W_t^i, \end{cases}$$

where $i = 1, \ldots, n$. Using $n \nabla_{\theta_i} f(\theta_1, \ldots, \theta_n) = \nabla F'(\mu^n)$ in the previous equation yields (6.7).

A.13 Consistency of the mean field limit

Proof of Theorem 5. The proof of consistency follows a standard martingale argument, which we briefly sketch here. Additional details can be found in 168, 169.

For $i = 1, \ldots, n$, let

$$X_t^i := \begin{pmatrix} \theta_t^i \\ \\ r_t^i \end{pmatrix}, \; b(X_t^i, \mu_t^n) := \begin{pmatrix} r_t^i \\ \\ F'([\mu_t^n]^{\theta_t^i})(\theta_t^i) - \gamma r_t^i \end{pmatrix}, \; \sigma(X_t^i, \mu_t^n) := \sqrt{2\beta^{-1}\gamma} \begin{pmatrix} 0_{d\times d} \\ \\ I_{d\times d} \end{pmatrix},$$

and consider the system of Itô stochastic differential equations:

$$\mathrm{d}X_t^i = b(X_t^i, \mu_t^n)\,\mathrm{d}t + \sigma(X_t^i, \mu_t^n)\,\mathrm{d}W_t^i, \tag{A.54}$$

where $\mathrm{d}W_t^i$, for each $i = 1, \ldots, n$, is the standard Wiener process in $\mathcal{T}\Theta$.

For any compactly supported test function $\varphi \in C_b^2(\mathcal{T}\Theta)$, i.e., the space of all bounded continuous functions $\varphi : \mathcal{T}\Theta \mapsto \mathbb{R}$ with bounded continuous partial derivatives of first and second order, we want to describe the time evolution of the quantity

$$\langle \varphi, \mu_t^n \rangle = \frac{1}{n} \sum_{i=1}^{n} \varphi(X_t^i). \tag{A.55}$$

Using Itô's rule, we have

$$\mathrm{d}\varphi(X_t^i) = L_{\mu_t^n}\varphi(X_t^i)\,\mathrm{d}t + \nabla\varphi^\top(X_t^i)\sigma(X_t^i, \mu_t^n)\,\mathrm{d}W_t^i,$$

wherein the infinitesimal generator L is defined as

$$L_\mu \varphi(x) := \langle b(x,\mu), \nabla_x \varphi(x) \rangle + \frac{1}{2} \text{trace}\left(\sigma\sigma^\top(x,\mu) \text{Hess}(\varphi)\right). \tag{A.56}$$

Therefore,

$$\begin{aligned} \mathrm{d}\langle \varphi, \mu_t^n \rangle &= \frac{1}{n} \sum_{i=1}^n \mathrm{d}\varphi(X_t^i) \\ &= \langle L_{\mu_t^n} \varphi, \mu_t^n \rangle \, \mathrm{d}t + \frac{1}{n} \sum_{i=1}^n \nabla \varphi^\top(X_t^i) \sigma(X_t^i, \mu_t^n) \, \mathrm{d}W_t^i \\ &:= \langle L_{\mu_t^n} \varphi, \mu_t^n \rangle \, \mathrm{d}t + \mathrm{d}M_t^n, \end{aligned} \tag{A.57}$$

where M_t^n is a local martingale. Since $\varphi \in C_b^2(\mathcal{T}\Theta)$, we have $|\nabla \varphi^\top \sigma| \leq \sqrt{2\beta^{-1}\gamma} |\nabla_r \varphi| \leq C$ uniformly for some $C > 0$. Notice that the quadratic variation of the noise term in (A.57) is

$$[M_t^n] = \frac{1}{n^2} \sum_{i=1}^n \int_0^t |\nabla \varphi^\top(X_s^i) \sigma(X_s^i, \mu_s^n)|^2 \, \mathrm{d}s \leq \frac{tC^2}{n},$$

and by Doob's martingale inequality, we deduce that

$$\mathbb{E}\left(\sup_{t \leq T} M_t^n\right)^2 \leq \mathbb{E}\left(\sup_{t \leq T} (M_t^n)^2\right) \leq 4\mathbb{E}\left((M_t^n)^2\right) \leq 4\mathbb{E}([M_t^n]) \leq \frac{4tC^2}{n}. \tag{A.58}$$

So as $n \to \infty$, the noise term in (A.57) converges to zero in probability, and we get a deterministic evolution equation.

Next, we argue that sequence $\{(\mu_t^n)_{t>0}\}_{n=1}^\infty$ of measure-valued stochastic processes converges to some probability measure-valued limiting process $(\mu_t)_{t>0}$ as $n \to \infty$. To this end, we take $\{(\mu_t^n)_{t>0}\}_{n=1}^\infty$ to be the (random) elements of $\Omega = C([0,\infty), \mathcal{M}(\mathcal{T}\Theta))$, the set of continuous functions from $[0,\infty)$ into $\mathcal{M}(\mathcal{T}\Theta)$

endowed with the topology of weak convergence. Following 168, 169, it can be shown that the sequence $\mathbb{P}_n$ of probability measures on Ω induced by the processes $\{(\mu_t^n)_{t>0}\}_{n=1}^{\infty}$ weakly converges (along a subsequence) to some $\mathbb{P}$, where $\mathbb{P}$ is the measure induced by the limiting process $(\mu_t)_{t>0}$. By Skorohod's representation theorem **?**, the sequence $\{(\mu_t^n)_{t>0}\}_{n=1}^{\infty}$ converges $\mathbb{P}$-almost surely to $(\mu_t)_{t>0}$. Since the martingale term in (A.57) vanishes as $n \to \infty$, we obtain

$$\mathrm{d}\langle \varphi, \mu_t \rangle = \langle L_\mu \varphi, \mu_t \rangle \, \mathrm{d}t = \langle \varphi, L_\mu^* \mu_t \rangle \, \mathrm{d}t, \tag{A.59}$$

which is valid almost everywhere for any test function $\varphi \in C_b^2(\mathcal{T}\Theta)$. In (A.59), L^* is the adjoint operator of L given by (A.56), and is defined as

$$L_m^* \mu(x) := -\nabla \cdot (\mu b(x, m)) + \frac{1}{2} \sum_{i,j=1}^{n} \frac{\partial^2}{\partial x_i x_j} (\mu \sigma \sigma^\top(x, m))_{ij}. \tag{A.60}$$

This shows that μ_t is almost surely a weak solution to the nonlinear Fokker-Planck PDE (6.8). ■

A.14 Variations and bounds on the free energy

This section provides the details of the proofs in Section 6.3, and additional bounds that are used in the proofs of the main results.

A.15 Proof of Lemma 15

Proof. By assumption, ρ_t satisfies the continuity equation $\partial \rho_t = -\nabla \cdot (\rho_t v_t)$, thus

$$\begin{aligned} \partial_t \mathcal{V}(\rho_t) &= \langle \mathcal{V}'(\rho_t), \partial_t \rho_t \rangle && \text{by the chain rule} \\ &= \langle \mathcal{V}'(\rho_t), -\nabla \cdot (\rho_t v_t) \rangle && \text{by the continuity equation} \\ &= \langle \nabla \mathcal{V}'(\rho_t), \rho_t v_t \rangle_* , \end{aligned}$$

where the last equality follows by duality of the gradient and divergence operators ∇ and $\nabla\cdot$, in the following sense: if $f : \mathcal{T}\Theta \to \mathbb{R}$ is a differentiable scalar function and $G : \mathcal{T}\Theta \to \mathbb{R}^d$ is a vector field, then

$$\langle \nabla f, G \rangle_* + \langle f, \nabla \cdot G \rangle = 0. \tag{A.61}$$

The inner product in the first summand above is for vector fields whereas the same in the second summand is for scalar-valued functions. ■

A.16 Lyapunov function in the single particle case

This section highlights a connection between Lyapunov functions for the single particle case, and Lyapunov functionals for the mean-field dynamics. To simplify the notation, let $\xi = (\theta, r)$ denote a position-velocity pair. Lemma 15 states that if $\rho_t \in \mathcal{P}(\mathcal{T}\Theta)$ solves the continuity equation $\partial_t \rho_t = -\nabla \cdot (\rho_t v(\rho_t))$, then the time

derivative of a functional $\mathcal{V}(\rho)$ along the solution trajectory ρ_t is given by

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{V}(\rho_t) = \langle \nabla\mathcal{V}'(\rho_t), \rho_t v(\rho_t)\rangle_* = \int_{\mathcal{T}\Theta} \langle \nabla\mathcal{V}'(\rho_t)(\xi), v(\rho_t)(\xi)\rangle \, \rho_t(\xi)\mathrm{d}\xi. \tag{A.62}$$

In the single particle case, if $\xi_t \in \mathcal{T}\Theta$ solves the differential equation $\dot{\xi}_t = v(\xi_t)$ for a vector field v, then the time derivative of a function $V(\xi)$ along ξ_t is, by the chain rule,

$$\frac{\mathrm{d}}{\mathrm{d}t}V(\xi_t) = \langle \nabla V(\xi_t), v(\xi_t)\rangle \,. \tag{A.63}$$

Comparing the two expressions, we see that (A.62) can be viewed as an integral version of (A.63). This connection is particularly simple in the linear case with no interaction: suppose $\mathcal{V}(\rho) = \langle V, \rho\rangle$ for a differentiable function $V : \mathcal{T}\Theta \to \mathbb{R}$, and the vector field v does not depend on ρ. Then the following holds:

If $\langle \nabla V(\xi), v(\xi)\rangle \leq 0$ for all ξ, then $\mathcal{V}$ is non-increasing along ρ_t, and V is non-increasing along ξ_t.

In other words, the same sufficient condition describes Lyapunov functions for ξ_t and Lyapunov functionals for ρ_t. In the nonlinear case, the condition becomes:

If $\langle \nabla\mathcal{V}'(\rho)(\xi), v(\rho)(\xi)\rangle \leq 0$ for all ρ and all ξ, then $\mathcal{V}$ is non-increasing along ρ_t, and $\mathcal{V}'(\rho)$ is non-increasing along the solution to $\dot{\xi}_t = v(\rho)(\xi_t)$.

In this case, the condition describes a *family* of single-particle dynamics $v(\rho)$ and corresponding Lyapunov functions $\mathcal{V}'(\rho)$, where the family is indexed by ρ.

We examine the case of the noiseless heavy ball dynamics as an example. In this

case, we have

$$v_{\mathrm{HB}}(\rho)(\theta, r) = \begin{pmatrix} r \\ -\nabla F'(\rho)(\theta) - \gamma r \end{pmatrix}, \tag{A.64}$$

$$\mathcal{V}_{\mathrm{HB}}(\rho)(\theta, r) = F(\rho) + \left\langle \frac{1}{2}|r|^2, \rho \right\rangle, \tag{A.65}$$

$$\mathcal{V}'_{\mathrm{HB}}(\rho)(\theta, r) = F'(\rho)(\theta) + \frac{1}{2}|r|^2, \tag{A.66}$$

corresponding to equations (6.8),(6.9) without diffusion.

Viewed as a single-particle dynamics, $v_{\mathrm{HB}}(\rho)(\cdot)$ describes the damped nonlinear oscillator with potential $F'(\rho)(\cdot)$. It is well-known from the optimization literature that (A.66) is a Lyapunov function for the dynamics (A.64), see, e.g., 87. This fact can be easily verified: for all θ, r,

$$\begin{aligned} \langle \nabla \mathcal{V}'_{\mathrm{HB}}(\rho)(\theta, r), v_{\mathrm{HB}}(\rho)(\theta, r) \rangle &= \left\langle \begin{pmatrix} \nabla F'(\rho)(\theta) \\ r \end{pmatrix}, \begin{pmatrix} r \\ -\nabla F'(\rho)(\theta) - \gamma r \end{pmatrix} \right\rangle \\ &= -\gamma |r|^2 \leq 0. \end{aligned}$$

Therefore, that $\mathcal{V}_{\mathrm{HB}}$ is a Lyapunov functional for the mean-field dynamics is a simple consequence of the single particle case. Proposition 16 is an extension of this fact to the case with diffusion.

A.17 Time-derivative of the free energy

Proof of Proposition 16. From the expression of the free energy $\mathcal{E}(\rho) = F([\rho]^\theta) + \left\langle \frac{1}{2}|r|^2, \rho \right\rangle + \frac{1}{\beta}\langle \log \rho, \rho \rangle$, we obtain,

$$\mathcal{E}'(\rho)(\theta, r) = F'([\rho]^\theta)(\theta) + \frac{1}{2}|r|^2 + \beta^{-1}(1 + \log \rho(\theta, r)), \tag{A.67}$$

and, using the shorthand $\ell_\theta := \beta^{-1}\nabla_\theta \log \rho_t$, $\ell_r := \beta^{-1}\nabla_r \log \rho_t$, we compute

$$\begin{aligned}
\partial_t \mathcal{E}(\rho_t) &= \langle \nabla \mathcal{E}'(\rho_t), \rho_t v(\rho_t) \rangle_* \\
&= \left\langle \begin{pmatrix} \nabla_\theta F'([\rho_t]^\theta) + \ell_\theta \\ r + \ell_r \end{pmatrix}, \rho_t \begin{pmatrix} r \\ -\nabla_\theta F'([\rho_t]^\theta) - \gamma r - \gamma \ell_r \end{pmatrix} \right\rangle_* \\
&= \left\langle -\gamma \langle r, r \rangle - \gamma \langle \ell_r, \ell_r \rangle - 2\gamma \langle r, \ell_r \rangle + \langle \ell_\theta, r \rangle - \left\langle \ell_r, \nabla_\theta F'([\rho_t]^\theta) \right\rangle, \rho_t \right\rangle.
\end{aligned}$$

We conclude by showing that $\langle \langle \ell_\theta, r \rangle, \rho_t \rangle$, $\left\langle \left\langle \ell_r, \nabla_\theta F'([\rho_t]^\theta) \right\rangle, \rho_t \right\rangle$ are equal to zero. Indeed,

$$\begin{aligned}
&\int_{\mathcal{T}\Theta} \langle \ell_\theta(\theta, r), r \rangle \, \rho_t(\theta, r) \, \mathrm{d}\theta \mathrm{d}r \\
&\quad = \beta^{-1} \int_{\mathcal{T}\Theta} \langle \nabla_\theta \log \rho_t(\theta, r), r \rangle \, \rho_t(\theta, r) \, \mathrm{d}\theta \mathrm{d}r \qquad \text{by definition of } \ell_\theta \\
&\quad = \beta^{-1} \int_{\mathcal{T}\Theta} \langle \nabla_\theta \rho_t(\theta, r), r \rangle \, \mathrm{d}\theta \mathrm{d}r \\
&\quad = -\beta^{-1} \int_{\mathcal{T}\Theta} \langle \rho_t(\theta, r), \nabla_\theta \cdot r \rangle \, \mathrm{d}\theta \mathrm{d}r \qquad \text{by duality (A.61)} \\
&\quad = 0,
\end{aligned}$$

and similarly,

$$\begin{aligned}
&\int_{\mathcal{T}\Theta} \left\langle \ell_r(\theta, r), \nabla_\theta F'([\rho_t]^\theta(\theta)) \right\rangle \rho_t(\theta, r)\, \mathrm{d}\theta \mathrm{d}r \\
&\quad = \beta^{-1} \int_{\mathcal{T}\Theta} \left\langle \nabla_r \log \rho_t(\theta, r), \nabla_\theta F'([\rho_t]^\theta(\theta)) \right\rangle \rho_t(\theta, r)\, \mathrm{d}\theta \mathrm{d}r \quad \text{by definition of } \ell_r \\
&\quad = \beta^{-1} \int_{\mathcal{T}\Theta} \left\langle \nabla_r \rho_t(\theta, r), \nabla_\theta F'([\rho_t]^\theta)(\theta) \right\rangle \mathrm{d}\theta \mathrm{d}r \\
&\quad = -\beta^{-1} \int_{\mathcal{T}\Theta} \left\langle \rho_t(\theta, r), \nabla_r \cdot \nabla_\theta F'([\rho_t]^\theta)(\theta) \right\rangle \mathrm{d}\theta \mathrm{d}r \qquad \text{by duality (A.61)} \\
&\quad = 0,
\end{aligned}$$

where the last equality is due to the fact $F'([\rho_t]^\theta)$ does not depend on r. ■

A.18 Additional bounds on the entropy and free energy

We recall the expression of the free energy:

$$\begin{aligned}
\mathcal{E}(\rho) &= F([\rho]^\theta) + \left\langle \frac{1}{2}|r|^2, \rho \right\rangle + H(\rho) \\
&= F_0([\rho]^\theta) + \left\langle g(\theta) + \frac{1}{2}|r|^2, \rho \right\rangle + H(\rho),
\end{aligned}$$

where $H(\rho) := \langle \log \rho, \rho \rangle$ is the negative entropy, $F_0(\rho) = R(\langle \Psi, \rho \rangle)$ is the unregularized risk, and $g : \Theta \to \mathbb{R}_+$ is the regularization function.

Let $\mathcal{K}$ be the set

$$\mathcal{K} := \{\rho \in \mathcal{P}(\mathcal{T}\Theta) : \left\langle g(\theta) + |r|^2/2, \rho \right\rangle < \infty\}. \tag{A.68}$$

First, we provide the following lower-bound on the free energy. For $\rho \in \mathcal{P}(\mathcal{T}\Theta)$,

we write $\log \rho = \log^+ \rho - \log^- \rho$, where $\log^+ \rho := \max\{\log \rho, 0\}$ and $\log^- \rho := \max\{-\log \rho, 0\}$.

Proposition 18. *Suppose that assumptions (A1)-(A4) hold. Then there exists a positive function $C(\alpha)$ such that for all $\rho \in \mathcal{K}$, and all $\alpha \leq \beta$*

$$\mathcal{E}(\rho) \geq F_0([\rho]^\theta) + (1 - \alpha/\beta) \langle g, \rho \rangle - \frac{C(\alpha)}{\beta}. \tag{A.69}$$

Proof. We can decompose $\mathcal{E}$ into

$$\mathcal{E}(\rho) = F_0([\rho]^\theta) + \left\langle g(\theta) + |r^2|/2, \rho \right\rangle + \frac{1}{\beta}(\left\langle \log^+ \rho, \rho \right\rangle - \left\langle \log^- \rho, \rho \right\rangle). \tag{A.70}$$

We focus on bounding the last term. First, following **?**, observe that for any constant $c \geq 1$, we have

$$x \log^- x \leq c(x + e^{-c}) \quad \text{for all } x \geq 0.$$

The inequality is trivial for $x \geq 1$ since the LHS is 0, by definition. For $x \in [0, 1]$, this can be verified by noting that the difference $d(x) := x \log^- x - c(x + e^{-c})$ attains its maximum at $x = e^{-c-1}$, and $d(e^{-c-1}) \leq 0$. Applying the previous inequality with a function $c : \mathcal{T}\Theta \to [1, +\infty)$, we have

$$\left\langle \log^- \rho, \rho \right\rangle \leq \left\langle c(\theta, r), \rho(\theta, r) + e^{-c(\theta,r)} \right\rangle.$$

Let $\alpha > 0$ and take $c(\theta, r) := 1 + \alpha(g(\theta) + \frac{|r|^2}{2})$, which is ≥ 1 since the regularizer g is non-negative by assumption. Then

$$\left\langle \log^- \rho, \rho \right\rangle \leq 1 + \alpha \left\langle g(\theta) + |r|^2/2, \rho \right\rangle + \int_{\mathcal{T}\Theta} (1 + \alpha g(\theta) + \alpha \frac{|r|^2}{2}) e^{-1 - \alpha g(\theta) - \alpha \frac{|r|^2}{2}} \mathrm{d}\theta \mathrm{d}r. \tag{A.71}$$

We shall prove that the last term, which we denote by $C(\alpha) := \int_{\mathcal{T}\Theta}(1 + \alpha g + \alpha\frac{|r|^2}{2})e^{-1-\alpha g(\theta)-\alpha\frac{|r|^2}{2}}$, is finite by virtue of assumption (A4). Indeed, the assumption guarantees that $e^{-\alpha g}$ is integrable. It also follows that $ge^{-\alpha g}$ is integrable: indeed, for any $\epsilon \in (0, \alpha)$, using the inequality $1 + \epsilon g \leq e^{\epsilon g}$, we can write that $ge^{-\alpha g} \leq \frac{e^{-\alpha g+\epsilon g}-e^{-\alpha g}}{\epsilon}$, and the upper-bound is integrable by assumption (A4).

To summarize, we obtain

$$\left\langle \log^- \rho, \rho \right\rangle \leq \alpha \left\langle g(\theta) + |r|^2/2, \rho \right\rangle + C(\alpha), \tag{A.72}$$

for a finite, positive function $C(\alpha)$. Using the last inequality in (A.70), and the fact $\left\langle \log^+ \rho, \rho \right\rangle \geq 0$, we obtain

$$\mathcal{E}(\rho) \geq F_0([\rho]^\theta) + \left\langle g(\theta) + |r^2|/2, \rho \right\rangle (1 - \alpha/\beta) - \frac{C(\alpha)}{\beta}.$$

Finally, taking $\alpha \leq \beta$ guarantees that the term $\langle |r|^2, \rho \rangle (1 - \alpha/\beta)$ is non-negative, and proves the claim (A.69). ■

Proposition 19. *Let $\rho_t \in C([0, \infty), \mathcal{P}(\mathcal{T}\Theta))$ be a solution to* (6.10) *with initial condition $\rho_0 \in \mathcal{P}(\mathcal{T}\Theta)$, and suppose that ρ_0 satisfies assumption (A5). Then for all $t \geq 0$, the quantities $\mathcal{E}(\rho_t)$, $F_0([\rho_t]^\theta)$, $\langle g(\theta) + |r|^2/2, \rho_t \rangle$, $\langle \log^+ \rho_t, \rho_t \rangle$, are bounded independently of t.*

Proof. From (A.70), we have

$$F_0([\rho_t]^\theta) + \left\langle g(\theta) + |r|^2/2, \rho_t \right\rangle + \beta^{-1}\langle \log^+ \rho_t, \rho_t \rangle = \mathcal{E}(\rho_t) + \beta^{-1}\langle \log^- \rho_t, \rho_t \rangle.$$

The terms on the left-hand-side are non-negative. We upper bound the right-

hand-side using (A.72), to obtain

$$F_0\left([\rho_t]^\theta\right) + \left\langle g(\theta) + |r|^2/2, \rho_t \right\rangle (1 - \alpha/\beta) + \beta^{-1}\langle \log^+ \rho_t, \rho_t \rangle \leq \mathcal{E}(\rho_t) + \beta^{-1}C(\alpha). \tag{A.73}$$

Choosing $\alpha < \beta$, as in the proof of Proposition 18, and using the fact $\mathcal{E}(\rho_t)$ is a decreasing function of t (Proposition 16), we have

$$\begin{aligned} 0 &\leq F_0\left([\rho_t]^\theta\right) + \left\langle g(\theta) + |r|^2/2, \rho_t \right\rangle (1 - \alpha/\beta) + \beta^{-1}\langle \log^+ \rho_t, \rho_t \rangle \\ &\leq \mathcal{E}(\rho_t) + \beta^{-1}C(\alpha) \\ &\leq \mathcal{E}(\rho_0) + \beta^{-1}C(\alpha) < \infty \end{aligned}$$

where the $\mathcal{E}(\rho_0)$ is finite by virtue of assumption (A5). The statement follows. ■

The following is a consequence of Propositions 16 and 19.

Theorem 9. *Consider the set up in Propositions 16 and 19. Then the solution trajectory $(\rho_t)_{t\geq 0}$ for (6.10) satisfies*

$$\lim_{t\to\infty} \int_{\mathcal{T}\Theta} |r + \beta^{-1}\nabla_r \log \rho_t|^2 \rho_t \, \mathrm{d}\theta \mathrm{d}r = 0. \tag{A.74}$$

Proof. Propositions 16 and 19 allow us to deduce that the functional $\mathcal{E}(\rho_t)$ given by (6.9) has a finite limit as $t \to \infty$. Now our strategy is to prove that $G := \partial_t \mathcal{E}$ is uniformly continuous in t. Then, by Barbalat's lemma **?**, the claim (A.74) follows.

To prove the uniform continuity of G in t, it suffices to show that $\partial_t G$ is upper bounded for all $t \geq 0$. First notice that

$$|\partial_t G| = |\langle \nabla G'(\rho_t), \rho_t v_t \rangle| = |\mathbb{E}_{\rho_t}\left[\nabla G', v_t\right]| \leq \sqrt{\mathbb{E}_{\rho_t}\left[|\nabla G'(\rho_t)|^2\right]} \sqrt{\mathbb{E}_{\rho_t}\left[|v_t|^2\right]}, \tag{A.75}$$

where the last inequality is due to Cauchy-Schwarz. By applying Cauchy-Schwarz again,

$$\begin{aligned}\mathbb{E}_{\rho_t}\Big[|v_t|^2\Big] &= \int_{\mathcal{T}\Theta}\Big(|r|^2+|\nabla_\theta F'([\rho_t]^\theta)+\gamma r+\gamma\ell_r|^2\Big)\,\rho_t \mathrm{d}\theta\mathrm{d}r \\ &\leq \int_{\mathcal{T}\Theta}\Big((1+3\gamma^2)|r|^2+3\gamma^2|\ell_r|^2+3|\nabla_\theta F'([\rho_t]^\theta)|^2\Big)\,\rho_t \mathrm{d}\theta\mathrm{d}r. \end{aligned} \tag{A.76}$$

Per Assumption (A3), $\nabla_\theta F'(\rho_t^\theta)(\theta) \in L^\infty(\Theta)$, and hence $\int|\nabla_\theta F'(\rho_t^\theta)|^2\rho_t \mathrm{d}\theta\mathrm{d}r < \infty$. From Proposition 19, we know that $\int|r|^2\rho_t \mathrm{d}\theta\mathrm{d}r < \infty$. Noting that $\int|\ell_r|^2\rho_t \mathrm{d}\theta\mathrm{d}r = \beta^{-2}\int\frac{|\nabla_r\rho_t|^2}{\rho_t}\mathrm{d}\theta\mathrm{d}r = 4\beta^{-2}\int|\nabla_r\sqrt{\rho_t}|^2\mathrm{d}\theta\mathrm{d}r$, and that $\nabla_r\sqrt{\rho} \in L^2([0,T],\mathcal{T}\Theta)$ for any $T > 0$ (see e.g., **?**), we have $\int|\ell_r|^2\rho_t \mathrm{d}\theta\mathrm{d}r < \infty$. Putting these together, we find that (A.76) is finite for all $t \geq 0$. We also note that the finiteness of $\int|\ell_r|^2\rho_t \mathrm{d}\theta\mathrm{d}r = \beta^{-2}\int\frac{|\nabla_r\rho_t|^2}{\rho_t}\mathrm{d}\theta\mathrm{d}r$ implies that ρ_t is positive almost everywhere, and that the Fisher information $\int\frac{|\nabla\rho_t|^2}{\rho_t}\mathrm{d}\theta\mathrm{d}r < \infty$.

To show that the other factor in the right-hand-side of (A.75) is finite, let $G_1 := -\gamma\langle|r|^2,\rho_t\rangle$, $G_2 := -2\gamma\int\langle r,\ell_r\rangle\rho_t \mathrm{d}\theta\mathrm{d}r$, $G_3 := -\gamma\langle|\ell_r|^2,\rho_t\rangle$, and notice that

$$G = -\gamma\int_{\mathcal{T}\Theta}|r+\ell_r|^2\rho_t\,\mathrm{d}\theta\mathrm{d}r = G_1+G_2+G_3. \tag{A.77}$$

Direct calculation of the functional derivatives yield

$$G_1' = -\gamma|r|^2, \tag{A.78a}$$

$$G_2' = 2\beta^{-1}\gamma\nabla_{\binom{\theta}{r}} \cdot \frac{\partial}{\partial\nabla_{\binom{\theta}{r}}\rho_t}\left\langle \begin{pmatrix}0\\ r\end{pmatrix}, \nabla_{\binom{\theta}{r}}\rho_t \right\rangle = 2d\beta^{-1}\gamma, \tag{A.78b}$$

$$\begin{aligned} G_3' &= -\beta^{-2}\gamma\left(-\frac{|\nabla_r\rho_t|^2}{\rho_t^2} - \nabla_{\binom{\theta}{r}} \cdot \rho_t^{-1}\frac{\partial}{\partial\nabla_{\binom{\theta}{r}}\rho_t}\left\langle \begin{pmatrix}0\\ \nabla_r\rho_t\end{pmatrix}, \nabla_{\binom{\theta}{r}}\rho_t \right\rangle\right) \\ &= -\beta^{-2}\gamma\left(-\frac{|\nabla_r\rho_t|^2}{\rho_t^2} - \frac{2}{\rho_t}\Delta_r\rho_t + \frac{2}{\rho_t^2}|\nabla_r\rho_t|^2\right) = -\beta^{-2}\gamma\left(\frac{|\nabla_r\rho_t|^2}{\rho_t^2} - \frac{2}{\rho_t}\Delta_r\rho_t\right). \end{aligned} \tag{A.78c}$$

Combining (A.77) and (A.78), we get

$$G' = -\gamma|r|^2 + 2d\beta^{-1}\gamma - \beta^{-2}\gamma\left(\rho_t^{-2}|\nabla_r\rho_t|^2 - 2\rho_t^{-1}\Delta_r\rho_t\right).$$

Therefore,

$$\nabla G' = \begin{pmatrix}\nabla_\theta G'\\ \nabla_r G'\end{pmatrix} = -\gamma\begin{pmatrix}\beta^{-2}\nabla_\theta\left(\rho_t^{-2}|\nabla_r\rho_t|^2 - 2\rho_t^{-1}\Delta_r\rho_t\right)\\ 2r + \beta^{-2}\nabla_r\left(\rho_t^{-2}|\nabla_r\rho_t|^2 - 2\rho_t^{-1}\Delta_r\rho_t\right)\end{pmatrix}. \tag{A.79}$$

Recalling that $\beta^{-2}\rho_t^{-2}|\nabla_r\rho_t|^2 = |\ell_r|^2$, we get

$$|\nabla G'|^2 \le 3\gamma^2\left(4|r|^2 + \left|\nabla_{\binom{\theta}{r}}|\ell_r|^2\right|^2 + 4\beta^{-4}\left|\nabla_{\binom{\theta}{r}}\rho_t^{-1}\Delta_r\rho_t\right|^2\right), \tag{A.80}$$

and hence $\mathbb{E}_{\rho_t}\left[|\nabla G'|^2\right]$ (the other factor in the RHS of (A.75)) is less than or equal

to

$$12\gamma^2 \underbrace{\int |r|^2 \rho_t \mathrm{d}\theta\mathrm{d}r}_{\text{term 1}} + 3\gamma^2 \underbrace{\int \left|\nabla_{\binom{\theta}{r}} |\ell_r|^2\right|^2 \rho_t \mathrm{d}\theta\mathrm{d}r}_{\text{term 2}} + 12\beta^{-4}\gamma^2 \underbrace{\int \left|\nabla_{\binom{\theta}{r}} \rho_t^{-1}\Delta_r \rho_t\right|^2 \rho_t \mathrm{d}\theta\mathrm{d}r}_{\text{term 3}}. \tag{A.81}$$

By Proposition 19, the term 1 in (A.81) is finite. Showing the finiteness of the terms 2 and 3 in (A.81) requires somewhat tedious estimates. We only sketch the main ideas for the same.

Letting $u := |\ell_r|^2$, term 2 equals $\int |\nabla u|^2 \rho_t \mathrm{d}\theta\mathrm{d}r = \int \langle \nabla u, \rho_t \nabla u\rangle \mathrm{d}\theta\mathrm{d}r$, which upon integration-by-parts and setting the boundary term to zero becomes:

$$-\int u\nabla\cdot(\rho_t\nabla u)\mathrm{d}\theta\mathrm{d}r = -\int u\left(\langle\nabla\rho_t, \nabla u\rangle + \rho_t\Delta u\right)\mathrm{d}\theta\mathrm{d}r \tag{A.82}$$

$$= -\mathbb{E}_{\rho_t}\left[\langle \nabla \log\rho_t, u\nabla u\rangle\right] + \mathbb{E}_{\rho_t}\left[u\Delta u\right]. \tag{A.83}$$

Thus, term 2 in (A.81) can be written as

$$\begin{aligned}\int |\nabla u|^2 \rho_t \mathrm{d}\theta\mathrm{d}r &= \left|-\mathbb{E}_{\rho_t}\left[\langle \nabla \log\rho_t, u\nabla u\rangle\right] + \mathbb{E}_{\rho_t}\left[u\Delta u\right]\right| \\ &\leq \mathbb{E}_{\rho_t}\left[\left|-\langle \nabla \log\rho_t, u\nabla u\rangle + u\Delta u\right|\right] \\ &\leq \mathbb{E}_{\rho_t}\left[\left|-\langle \nabla \log\rho_t, u\nabla u\rangle\right|\right] + \mathbb{E}_{\rho_t}\left[\left|u\Delta u\right|\right] \\ &\leq |\nabla\log\rho_t|_{L^2(\rho_t)} |u\nabla u|_{L^2(\rho_t)} + |u|_{L^1(\rho_t)} |\Delta u|_{L^\infty(\rho_t)},\end{aligned} \tag{A.84}$$

wherein we used the Jensen's, triangle and Hölder's inequalities, respectively. Finiteness for two of the four terms in (A.84) have been pointed out before: $|\nabla\log\rho_t|_{L^2(\rho_t)} = \int \frac{|\nabla\rho_t|^2}{\rho_t}\mathrm{d}\theta\mathrm{d}r$ (the Fisher information) $< \infty$, and

$$|u|_{L^1(\rho_t)} = \int |\ell_r|^2 \rho_t \mathrm{d}\theta\mathrm{d}r < \infty.$$

Following some calculation, the same estimates can be used to bound the remaining two terms.

For term 3 in (A.81), notice that $\rho_t^{-1}\Delta_r\rho_t = \beta^2|\ell_r|^2+\beta\nabla_r\cdot\ell_r$, and hence term 3 equals

$$\mathbb{E}_{\rho_t}\left|2\beta^2\left(\frac{\partial\ell_r}{\partial\theta}\right)^\top\ell_r+\beta\nabla_\theta\left(\nabla_r\cdot\ell_r\right)\right|^2+\mathbb{E}_{\rho_t}\left|2\beta^2\left(\frac{\partial\ell_r}{\partial r}\right)^\top\ell_r+\beta\nabla_r\left(\nabla_r\cdot\ell_r\right)\right|^2.$$

Similar estimates as before show the finiteness of the above. We summarize: since each of the two factors in the RHS of (A.75) are finite, $\partial_t G$ is upper bounded for all $t\geq 0$, which suffices to conclude that G is uniformly continuous in t. Then by Barbalat's lemma **?**, (A.74) follows. ■

A.19 Stationary solutions and convergence

A.20 Proof of Theorem 6

We seek to prove that any stationary solution $\rho^\star$ decomposes into the product of marginals $\rho^\star = \frac{\exp(-\beta\frac{|r|^2}{2})}{Z_1}[\rho^\star]^\theta$, where $[\rho^\star]^\theta$ is a solution to the Boltzmann equation (6.13).

Proof. Let ρ_t be the solution initialized at $\rho^\star$. Since $\rho^\star$ is stationary, we must have $\partial_t\mathcal{E}(\rho_t)=0$, i.e.,

$$\int|r+\beta^{-1}\nabla_r\log\rho^\star|^2\mathrm{d}\rho^\star=0, \tag{A.85}$$

by Proposition 16. Let $\rho^\star(\theta,r)=\frac{\exp\left(-\frac{\beta}{2}|r|^2\right)}{Z_1}\eta(\theta,r)$. To prove the first part of the claim, we seek to show that $\eta(\theta,\cdot)$ is a constant for a.e. θ. We have $r+$

$\beta^{-1}\nabla_r \log \rho^\star = \beta^{-1}\nabla_r \log \eta$, thus (A.85) yields

$$\nabla_r \log \eta = 0 \quad \eta\text{-a.e.}$$

This implies that for a.e. θ, the function $\eta(\theta, \cdot)$ is constant on its support; but it is also constant (equal to 0) outside its support, thus by continuity it is constant on its entire domain. This proves the first part of the claim.

So far, we have shown that there exists $\eta \in \mathcal{P}(\Theta)$ such that

$$\rho^\star(\theta, r) = \frac{\exp\left(-\frac{\beta}{2}|r|^2\right)}{Z_1}\eta(\theta),$$

and we seek to characterize η. By the continuity equation (6.10), since $\rho^\star$ is stationary, we must have $\nabla \cdot (\rho^\star v(\rho^\star)) = 0$, i.e.,

$$\begin{aligned}
0 &= \nabla_\theta \cdot (\rho^\star r) + \nabla_r \cdot [\rho^\star(-\nabla F'(\eta) - \gamma r - \gamma\beta^{-1}\nabla_r \log \rho^\star)] \\
&= \nabla_\theta \cdot (\rho^\star r) - \nabla_r \cdot [\rho^\star \nabla F'(\eta)] \quad \text{since } \nabla_r \log \rho^\star = -\beta r, \\
&= \left\langle \frac{\exp\left(-\frac{\beta}{2}|r|^2\right)}{Z_1} r, \nabla_\theta \eta \right\rangle - \left\langle \eta \nabla_\theta F'(\eta), \nabla_r \frac{\exp\left(-\frac{\beta}{2}|r|^2\right)}{Z_1} \right\rangle \\
&= \left\langle \frac{\exp\left(-\frac{\beta}{2}|r|^2\right)}{Z_1} r, \nabla_\theta \eta + \beta\eta\nabla_\theta F'(\eta) \right\rangle,
\end{aligned}$$

where the equality is for a.e. r, θ. Therefore, $\nabla_\theta \eta + \beta\eta\nabla_\theta F'(\eta) = 0$. This is equivalent to

$$\nabla_\theta \log \eta + \beta \nabla_\theta F'(\eta) = 0, \ \eta\text{-a.e.},$$

and integrating, we obtain: $\log \eta = -\beta \nabla_\theta F'(\eta) +$ a constant. This is equivalent to the Boltzmann fixed point equation (6.13), as desired. ■

A.21 Proof of Proposition 17

We seek to prove that the operator

$$T : \rho \mapsto T(\rho) = \frac{\exp(-\beta F'(\rho))}{Z_2(\rho)}$$

admits a unique point on $\mathcal{P}(\Theta)$. Note that T is well-defined for all ρ by virtue of assumption (A4). Indeed, $F'(\rho) = F_0'(\rho) + g$, and assumption (A4) states that F_0' is uniformly bounded on $\mathcal{P}$, and $\exp(-\beta g)$ is integrable, thus if M is an upper bound on $\|F_0'(\rho)\|_\infty$, we have $\exp(-\beta F'(\rho)) \leq \exp \beta M \exp(-\beta g)$, which is integrable.

Proof of existence: We will use Schauder's fixed point theorem **?**,**?**, stated below. Recall that a subset of a metric space is precompact if any sequence in that subset has a converging subsequence.

Theorem (Schauder's fixed point theorem)**.** *Let X be a Banach space and $M \subset X$ be non-empty, convex and closed. If $T : M \mapsto M$ is a continuous operator such that $T(M)$ is precompact, then T has a fixed point.*

Let $B = \{\rho \in L^1(\Theta) : \rho \geq 0, \|\rho\|_1 \leq 1\}$. Note that $T(B) \subset \mathcal{P}(\Theta)$, since $T(\rho)$ is normalized. Thus, to prove that T has a fixed point on $\mathcal{P}$, it suffices to prove that T has a fixed point on B. To this end, we apply Schauder's theorem with $X = L^1(\Theta)$ and $M = B$.

Let (ρ_n) be a sequence of elements in B. We shall prove that $(T(\rho_n))$ has a converging subsequence. We have that F' decomposes into $F'(\rho_n) = F_0'(\rho_n) + g$. By assumption (A4), $(F_0'(\rho_n))_n$ is uniformly equicontinuous and uniformly bounded, thus by the Arzela-Ascoli theorem, there exists a subsequence $(F_0'(\rho_{k_n}))$

that converges uniformly to some continuous, bounded function ℓ. We will show that $T(\rho_{k_n})$ converges in L^1 to $\rho := \exp(-\beta(\ell + g))/\|\exp(-\beta(\ell + g))\|_1$, which is well-defined since ℓ is bounded and g is confining. Observing that for all n, $\|T(\rho_{k_n})\|_1 = \|\rho\|_1 = 1$, we have by Scheffé's lemma **?** that poitwise convergence of $T(\rho_{k_n})$ to ρ implies convergence in L^1. Thus it suffices to prove pointwise convergence.

By continuity of the exponential function, we have

$$\exp(-\beta F'(\rho_{k_n})) \to \exp(-\beta(\ell + g)), \tag{A.86}$$

where the convergence is pointwise. By assumption (A4), $\{F_0'(\rho), \rho \in B\}$ is uniformly bounded and $\exp(-\beta g)$ is integrable, thus, by the dominated convergence theorem,

$$\|\exp(-\beta F'(\rho_{k_n}))\|_1 \to \|\exp(-\beta(\ell + g))\|_1. \tag{A.87}$$

By (A.86) and (A.87), we have $T(\rho_{k_n})$ converges pointwise to ρ, which concludes the proof.

■

Proof of uniqueness: Suppose $\rho_1, \rho_2 \in \mathcal{P}(\Theta)$ are two fixed points of T. Then we have for $i \in \{1, 2\}$,

$$\log Z(\rho_i) = -\log \rho_i(\theta) - \beta F'(\rho_i)(\theta) \text{ for a.e. } \theta. \tag{A.88}$$

We have

$$
\begin{aligned}
0 &= \langle \log Z(\rho_1) - \log Z(\rho_2), \rho_1 - \rho_2 \rangle \quad \text{since } \log Z(\rho_i) \text{ are constants} \\
&= -\langle \log \rho_1 + \beta F'(\rho_1) - \log \rho_2 - \beta F'(\rho_2), \rho_1 - \rho_2 \rangle \quad \text{by (A.88)} \\
&= -\langle F'(\rho_1) - F'(\rho_2), \rho_1 - \rho_2 \rangle - D_{\mathrm{KL}}(\rho_1 \| \rho_2) - D_{\mathrm{KL}}(\rho_2 \| \rho_1) \\
&\leq -D_{\mathrm{KL}}(\rho_1 \| \rho_2) - D_{\mathrm{KL}}(\rho_2 \| \rho_1) \quad \text{by convexity of } F.
\end{aligned}
$$

where $D_{\mathrm{KL}}(\rho_1 \| \rho_2) = \left\langle \log \frac{\rho_1}{\rho_2}, \rho_1 \right\rangle$. Note that ρ_1, ρ_2 are both normalized by assumption, so both KL divergences are non-negative, with equality if and only if $\rho_1 = \rho_2$ a.e. This concludes the proof. ■

A.22 Proof of Theorem 7

Proof. (i) Recall from Section 2.2 that under the stated conditions on the initial measure μ_0, the equation (6.10) admits a unique solution $(\mu_t)_{t\geq 0}$ satisfying $\mu_t \in C\left([0,\infty), \mathcal{M}(\mathcal{T}\Theta)\right)$, that is, $(\mu_t)_{t\geq 0}$ is a continuous measure-valued trajectory satisfying $\int \mathrm{d}\mu_t < \infty$ for all $t \geq 0$.

From Proposition 19, we know that the quantities $F_0([\mu_t]^\theta) < \infty$, $\int_{\mathcal{T}\Theta}(g(\theta) + |r|^2/2)\mathrm{d}\mu_t < \infty$, $\int_{\mathcal{T}\Theta} \log^+ \mu_t \mathrm{d}\mu_t < \infty$ for all $t \geq 0$, with their upper bounds being independent of t. Hence by the Dunford-Pettis theorem **?**, the solutions $(\mu_t)_{t\geq 0}$ are weakly compact in $L^1(\mathcal{T}\Theta)$. Thus, there exists $\mu^\star$ and a subsequence $(\mu_{t_k})_{k\geq 1}$ such that $(\mu_{t_k})_{k\geq 1}$ converges weakly to $\mu^\star$.

To prove μ_t is absolutely continuous (w.r.t. the Lebesgue measure) for each $t \geq 0$, we now show that the sequence of random vectors $(X_k)_{k\geq 1} := (\theta_{t_k}, r_{t_k})_{k\geq 1}$ are uniformly integrable. By de la Vallée-Poussin's criterion **?**, the latter holds if and

only if there is an increasing function $\Phi : \mathbb{R}_{>0} \mapsto \mathbb{R}_{>0}$ satisfying $\lim_{x\to\infty} \frac{\Phi(x)}{x} = +\infty$, such that $\sup_{k\geq 1} \mathbb{E}_{\mu_{t_k}}[\Phi(X_k)] < \infty$. To apply this in our context, we set $\Phi(x) \equiv x^2$, and use the result from Proposition (19) that $\int_{\mathcal{T}\Theta} |r_{t_k}|^2 \mathrm{d}\mu_{t_k}$ is uniformly upper bounded for all $k \geq 1$. Therefore, $(X_k)_{k\geq 1}$ are uniformly integrable, and equivalently, the measures μ_{t_k} are absolutely continuous, and the corresponding joint PDFs exist for all $k \geq 1$. Taking $\{t_k\}_{k\geq 1}$ to be an arbitrary sequence, we deduce that μ_t is absolutely continuous for each $t \geq 0$. Taking $\{t_k\}_{k\geq 1}$ to be the sequence corresponding to the weakly convergent subsequence $(\mu_{t_k})_{k\geq 1}$ mentioned in the previous paragraph, we deduce that $\mu^\star$ is absolutely continuous.

(ii) Let us consider the joint PDF trajectory $(\rho_t)_{t\geq 0}$ corresponding to the measure-valued trajectory $(\mu_t)_{t\geq 0}$ that solves (6.10). From part (i), we know that $(\rho_t)_{t\geq 0}$ exists and is weakly compact in $L^1(\mathcal{T}\Theta)$. Letting

$$\zeta_t(s, \theta, r) := \rho_{t+s}(\theta, r),$$

we now prove that $(\zeta_t)_{t\geq 0}$ is strongly compact in $C([0, T], L^1(\mathcal{T}\Theta))$ for any $T > 0$. From Theorem 9, we can write

$$\lim_{t\to\infty} \int_0^T \partial_t \mathcal{E}(t + s)\, \mathrm{d}s = 0, \tag{A.89}$$

which combined with Proposition 16 yields

$$\lim_{t\to\infty} \| r\sqrt{\zeta_t} + 2\beta^{-1}\nabla_r \sqrt{\zeta_t}\|_{L^2([0,T]\times\mathcal{T}\Theta)} = 0. \tag{A.90}$$

The remaining proof follows the same line of arguments as in **?**. Specifically, for any given sequence $\{t_k\}_{k\geq 1}$ with $\lim_{k\to\infty} t_k = \infty$, letting $\zeta_k := \zeta_{t_k}$, one shows that the sequence of $\{\zeta_k\}_{k\geq 1}$ is relatively compact in $C([0, T], L^1(\mathcal{T}\Theta))$ for any $T > 0$.

Consequently, $\zeta_k \to \zeta_\infty$ strongly in $L^1(\mathcal{T}\Theta)$, which is to say $\rho_t \to \rho^\star$ (equivalently, $\mu_t \to \mu^\star$) strongly in $L^1(\mathcal{T}\Theta)$, as desired. ■

A.23 Proof of Theorem 8

We start by showing that $\rho^\star$ is a minimizer of $\mathcal{E}(\rho)$ over $\mathcal{K}$, by adapting the argument from the first-order case **?**. We omit some details and emphasize the differences.

Recall that

$$\mathcal{K} = \{\rho \in \mathcal{P}(\mathcal{T}\Theta) : \left\langle g(\theta) + |r|^2/2, \rho \right\rangle < \infty\}.$$

Lemma 20. *Let $\rho^\star$ be the unique solution of the Boltzmann fixed point equation* (6.13)*. Then for all $\rho \in \mathcal{K}$, $\mathcal{E}(\rho) \geq \mathcal{E}(\rho^\star)$.*

Proof. First, we argue that $\mathcal{E}$ has a minimizer over $\mathcal{K}$. Note that $\mathcal{E}(\rho)$ is lower-bounded on $\mathcal{K}$ by Proposition 18. Thus, $\inf_{\rho\in\mathcal{K}} \mathcal{E}(\rho)$ is finite and there exists a sequence $\rho_k \in \mathcal{K}$ such that $\lim_{k\to\infty} \mathcal{E}(\rho_k) = \inf_{\rho\in\mathcal{K}} \mathcal{E}(\rho)$. Furthermore, by the same argument as the proof of Proposition 19, the quantities

$$F_0([\rho_k]^\theta), \left\langle g(\theta) + |r|^2/2, \rho_k \right\rangle, \left\langle \log^+ \rho_k, \rho_k \right\rangle \tag{A.91}$$

are bounded uniformly in k. Thus, by de la Vallée-Poussin's criterion **?**, there exists $\rho_\infty \in \mathcal{K}$ such that a subsequence of ρ_k converges weakly to ρ_∞. By lower semi-continuity of $\mathcal{E}$, we have $\mathcal{E}(\rho_\infty) = \inf_{\rho\in\mathcal{K}} \mathcal{E}(\rho)$.

Second, we show that any minimizer of $\mathcal{E}$ on $\mathcal{K}$, must, in fact, be equal to $\rho^\star$. Let $\bar{\rho}$ be such a minimizer. Then $\bar{\rho}$ must be positive a.e., otherwise, a perturbation of $\bar{\rho}$

can decrease the value of $\mathcal{E}$. Indeed, suppose that there exists a bounded subset S of positive Lebesgue measure, such that $\bar{\rho} \equiv 0$ on S, and define $\bar{\rho}_\epsilon = (1-\epsilon)\bar{\rho} + \epsilon u_S$, where $u_S = 1_S / \|1_S\|_1$ is the uniform distribution over S. Then $\bar{\rho}_\epsilon$ is in $\mathcal{K}$ (since S is bounded), and there exist constants A_0 and B_0 such that

$$\begin{aligned}
F([\bar{\rho}_\epsilon]^\theta) &\leq (1-\epsilon) F([\bar{\rho}]^\theta) + \epsilon A_0 && \text{by convexity of } F \\
\left\langle |r|^2/2, \bar{\rho}_\epsilon \right\rangle &\leq (1-\epsilon) \left\langle |r|^2/2, \rho \right\rangle + \epsilon B_0 && \text{by boundedness of } S \\
H(\bar{\rho}_\epsilon) &= \langle \log((1-\epsilon)\bar{\rho} + \epsilon u_S), (1-\epsilon)\bar{\rho} + \epsilon u_S \rangle \\
&\leq (1-\epsilon) H(\bar{\rho}) + \epsilon \log \frac{\epsilon}{\|1_S\|_1}.
\end{aligned}$$

Summing the previous inequalities, we see that there exists a constant C such that $\mathcal{E}(\bar{\rho}_\epsilon) \leq (1-\epsilon)\mathcal{E}(\bar{\rho}) + \epsilon C + \frac{1}{\beta}\epsilon \log(\epsilon)$, which is strictly less than $\mathcal{E}(\bar{\rho})$ for $\epsilon < e^{-\beta C}$, a contradiction. Therefore $\bar{\rho}$ must be positive a.e.

Once we have established that $\bar{\rho}$ is positive a.e., we can show that $\bar{\rho}$ satisfies the Boltzmann fixed point equation (6.13). Indeed, consider the set $\Gamma_k := \{(\theta, r) : \frac{1}{k} \leq \bar{\rho}(\theta, r) \leq k\}$, and let $\mathcal{T}_k = \{f \in C^\infty(\mathcal{T}\Theta) : \text{support}(f) \subseteq \Gamma_k, \|f\|_\infty \leq 1, \int f = 0\}$. In other words, $\mathcal{T}_k$ is a set of tangent vectors such that $\bar{\rho} + \frac{1}{k}\mathcal{T}_k \subset \mathcal{K}$. The directional derivative of $\mathcal{E}$ in the direction $f \in \mathcal{T}_k$ is well-defined and given by

$$\lim_{\epsilon \to 0} \frac{\mathcal{E}(\bar{\rho} + \epsilon f) - \mathcal{E}(\bar{\rho})}{\epsilon} = \left\langle F'([\bar{\rho}]^\theta) + \frac{1}{2}|r|^2 + (1 + \log \bar{\rho}), f \right\rangle \tag{A.92}$$

and since $\bar{\rho}$ is a minimizer of $\mathcal{E}$ on $\mathcal{K}$, (A.92) must be non-negative for all f. Therefore one must have that the integrand $F'([\bar{\rho}]^\theta) + \frac{1}{2}|r|^2 + (1 + \log \bar{\rho})$ is zero a.e. on Γ_k. But since $\mathcal{T}\Theta = \cup_{k \geq 1} \Gamma_k$, it must be zero a.e. on $\mathcal{T}\Theta$. This implies that $\bar{\rho}$ is a solution to the Boltzmann fixed point equation (6.13), which admits a unique solution $\rho^\star$ by Proposition 17. This concludes the proof. ■

Proof of Theorem 8. Let F_λ denote the regularized functional with regularization coefficient λ, i.e. $F_\lambda(\rho) = F_0(\rho) + \lambda \langle g, \rho \rangle$. We shall prove that there exists a constant C_1 such that, for all $\beta \geq 1$,

$$F_{1-1/\beta}([\rho^\star]^\theta) \leq \inf_{\eta \in \mathcal{P}(\Theta)} F(\eta) + \frac{C_1 + d \log \beta}{\beta}.$$

By Lemma 20, we have $\mathcal{E}(\rho^\star) \leq \mathcal{E}(\rho)$ for all $\rho \in \mathcal{K}$, and observing that $\rho^\star \in \mathcal{K}$, we have by Proposition 18 applied to $\rho^\star$ and $\alpha = 1$, $F_{1-1/\beta}([\rho^\star]^\theta) \leq \mathcal{E}(\rho^\star) + C(1)/\beta$. Combining the previous bounds, we have for all $\rho \in \mathcal{K}$,

$$F_{1-1/\beta}([\rho^\star]^\theta) \leq \mathcal{E}(\rho) + \frac{C(1)}{\beta}. \tag{A.93}$$

In order to conclude, we shall bound the difference between $\mathcal{E}$ and F. Note that $\mathcal{E}(\rho) - F([\rho]^\theta) = \frac{1}{2} \langle |r|^2, \rho \rangle + \frac{1}{\beta} \langle \log \rho, \rho \rangle$, which can be arbitrarily large due to the entropy term. To resolve this issue, one can take a convolution with a Gaussian to control the entropy. More precisely, let $\eta \in \mathcal{P}(\Theta)$, and define $\rho_\eta \in \mathcal{K}$ as the product:

$$\rho_\eta(\theta, r) := [g_1 * \eta](\theta) g_2(r)$$

where g_1, g_2 are two Gaussian PDFs over θ and r respectively, each with mean 0 and variance $1/\beta$, and $*$ denotes the convolution. Our goal is to bound the difference between $\mathcal{E}(\rho_\eta)$ and $F(\eta)$. Following the same line of argument as in **?**,

there exists a constant K such that

$$
\begin{aligned}
&F(g_1 * \eta) \leq F(\eta) + \frac{K}{\beta}, \\
&\left\langle \frac{1}{2}|r|^2, g_2 \right\rangle = \frac{d/2}{\beta}, \\
&H(\rho_\eta) \leq (H(g_1) + H(g_2)) = -d \log(2\pi e/\beta).
\end{aligned}
$$

Summing the inequalities above, we obtain

$$
\mathcal{E}(\rho_\eta) = F(g_1 * \eta) + \frac{1}{2}\left\langle |r|^2, g_2 \right\rangle + \frac{1}{\beta} H(\rho_\eta) \leq F(\eta) + \frac{K + d/2 - d\log(2\pi e/\beta)}{\beta}. \tag{A.94}
$$

Finally, we combine the inequalities (A.93) and (A.94), to obtain, for all $\eta \in \mathcal{P}(\Theta)$,

$$
F_{1-1/\beta}([\rho^\star]^\theta) \leq F(\eta) + \frac{C_1 + d\log(\beta)}{\beta},
$$

where C_1 is a constant equal to $C(1) + K + d/2 - d\log(2\pi e)$. Taking the infimum over η yields the desired result. ■

A.24 The case of quadratic loss

In this section, we illustrate the assumptions in the quadratic loss case.

Let R be given by $R(\psi) = \frac{1}{2}\mathbb{E}_{(x,y)\sim D}(\psi(x) - y)^2 = \frac{1}{2}\|\psi - y\|_{\mathcal{F}}^2$, where (x, y) are the input feature and labels, respectively, and D is the joint data distribution. The

functional F is the sum $F(\mu) = F_0(\mu) + \langle g, \mu \rangle$, where for $\mu \in \mathcal{M}(\Theta)$,

$$\begin{aligned} F_0(\mu) &= R(\langle \Psi, \mu \rangle) \\ &= \frac{1}{2} \| \langle \Psi, \mu \rangle - y \|_{\mathcal{F}}^2 \\ &= \frac{1}{2} \| \langle \Psi, \mu \rangle \|_{\mathcal{F}}^2 - \langle \langle \Psi, \mu \rangle, y \rangle_{\mathcal{F}} + \frac{1}{2} \|y\|_{\mathcal{F}}^2. \end{aligned} \tag{A.95}$$

The first term in (A.95) can be written as

$$\begin{aligned} \frac{1}{2} \| \langle \Psi, \mu \rangle \|_{\mathcal{F}}^2 &= \frac{1}{2} \mathbb{E}_x \left(\int_\Theta \Psi(\theta)(x) \mathrm{d}\mu(\theta) \right)^2 = \frac{1}{2} \iint_\Theta \mathbb{E}_x [\Psi(\theta)(x) \Psi(\tilde{\theta})(x)] \mathrm{d}\mu(\theta) \mathrm{d}\mu(\tilde{\theta}) \\ &= \frac{1}{2} U[\mu, \mu], \end{aligned}$$

where $U(\theta, \tilde{\theta}) := \mathbb{E}_x[\Psi(\theta)(x)\Psi(\tilde{\theta})(x)]$, and the symbol $U[\mu, \nu]$ denotes the double integral $\iint U(\theta, \tilde{\theta}) \mathrm{d}\mu(\theta) \mathrm{d}\nu(\tilde{\theta})$. The second term in (A.95) can be written as

$$- \langle \langle \Psi, \mu \rangle, y \rangle_{\mathcal{F}} = -\mathbb{E}_{(x,y)} \left[y \int_\Theta \Psi(\theta)(x) \mathrm{d}\mu(\theta) \right] = \langle V, \mu \rangle,$$

where $V(\theta) := -\mathbb{E}_{x,y}[y\Psi(\theta)(x)]$. The last term in (A.95) is a constant independent of μ. To summarize, the functional F_0 can be written as

$$F_0(\mu) = \frac{1}{2} U[\mu, \mu] + \langle V, \mu \rangle + \frac{1}{2} \|y\|_{\mathcal{F}}^2. \tag{A.96}$$

We now discuss our assumptions in this quadratic case. In particular, we show that the assumptions made in 158 (for the first-order gradient flow) imply our assumption (A4).

First, it is assumed in 158 that a quadratic regularizer is used, $g(\theta) = |\theta|^2/2$, which is confining since $\lim_{|\theta| \to \infty} g(\theta) = \infty$ and $\exp(-\beta g)$ is integrable for all $\beta > 0$. This proves the second part of assumption (A4). They also make the

following assumptions on U, V.

(B1) U and V are uniformly bounded i.e., there exists $C_1, C_2 > 0$ such that

$$\|U(\cdot,\cdot)\|_\infty \leq C_1, \quad \|V(\cdot)\|_\infty \leq C_2.$$

(B2) U and V are differentiable, and have bounded gradients, i.e., there exist C_3, C_4 such that

$$\|\nabla U(\cdot,\tilde{\theta})\|_\infty \leq C_4 \text{ for all } \tilde{\theta}, \quad \|\nabla V(\cdot)\|_\infty \leq C_3.$$

To prove that (A4) is satisfied, we need to show that the family $\{F_0'(\rho), \rho \in B\}$ is uniformly equicontinuous and uniformly bounded, where $B = \{\rho \in L^1(\Theta), \|\rho\|_1 \leq 1\}$. From (A.96), the Fréchet differential of F_0 is given by

$$F_0'(\rho)(\theta) = U[\rho](\theta) + V(\theta),$$

where the symbol $U[\rho]$ denotes the function $U[\rho](\theta) = \int U(\theta,\tilde{\theta})\rho(\tilde{\theta})\mathrm{d}\theta$. Then,

- By (B1), $U[\rho] + V$ is bounded, uniformly in $\rho \in B$.
- From (B2), it also follows that for $U[\rho] + V$ is Lipschitz continuous, with a Lipschitz constant independent of the choice of $\rho \in B$ and thus the family $\{U[\rho] + V, \rho \in B\}$ is uniformly equicontinuous.

Thus assumption (A4) is satisfied.

Note that in 158, the regularization term $g(\theta) = \langle |\theta|^2/2, \mu \rangle$, together with the

boundedness assumptions (B1)-(B2), are crucial to guarantee integrability of

$$\exp(-\beta F'(\mu)),$$

so that the Boltzmann distribution (6.13) is well-defined. In the linear case described in Section 6.2.3, it is also common to assume that the potential (which in this case is the same as our regularizer g) is confining, see for example **?**.

Assumption (A4) generalizes the conditions on F_0 from the quadratic setting to the convex setting, and replaces the quadratic regularizer with a more general confining regularizer.

A.25 Proof of Theorem 3

The derivation of (4.19a)-(4.19b) follows [44, Appendix C]. By substituting the identity $\nabla q = \exp(V/\theta)(\nabla p + p\nabla V/\theta)$ in $0 = \langle \nabla\varphi, \boldsymbol{\nu}\rangle = \langle \nabla q, \boldsymbol{\nu}\rangle$, we find (4.19c). ■

Bibliography

[1] Naoufel Ben Abdallah and Pierre Degond. The Child-Langmuir law in the kinetic theory of charged particles: Semiconductors models. *Mathematical Problems in Semiconductor Physics*, 340:76, 1995.

[2] Andrei Agrachev and Paul Lee. Optimal transportation under nonholonomic constraints. *Transactions of the American Mathematical Society*, 361(11):6019–6047, 2009.

[3] Giacomo Albi, Young-Pil Choi, Massimo Fornasier, and Dante Kalise. Mean field control hierarchy. *Applied Mathematics & Optimization*, 76(1):93–135, 2017.

[4] Fernando Albiac and Nigel John Kalton. *Topics in Banach space theory*, volume 233. Springer, 2006.

[5] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252. PMLR, 2019.

[6] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.

[7] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR, 2019.

[8] Zahra Askarzadeh, Rui Fu, Abhishek Halder, Yongxin Chen, and Tryphon T Georgiou. Stability theory of stochastic models in opinion dynamics. *IEEE Transactions on Automatic Control, doi: https://doi.org/10.1109/TAC.2019.2912490*, 2019.

[9] Hedy Attouch and Felipe Alvarez. The heavy ball with friction dynamical

system for convex constrained minimization problems. In *Lecture Notes in Econom. and Math. Systems*, pages 25–35. Springer, 2000.

[10] Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.

[11] Efstathios Bakolas. Covariance control for discrete-time stochastic linear systems with incomplete state information. In *2017 American Control Conference (ACC)*, pages 432–437. IEEE, 2017.

[12] Efstathios Bakolas. Finite-horizon covariance control for discrete-time stochastic linear systems subject to input constraints. *Automatica*, 91:61–68, 2018.

[13] Dominique Bakry, Patrick Cattiaux, and Arnaud Guillin. Rate of convergence for ergodic continuous markov processes: Lyapunov versus poincaré. *J. Funct. Anal.*, 254(3):727–759, 2008.

[14] Radu Balescu. *Equilibrium and nonequilibrium statistical mechanics*. Wiley, 1975.

[15] Saptarshi Bandyopadhyay, Soon-Jo Chung, and Fred Y Hadaegh. Probabilistic and distributed control of a large-scale swarm of autonomous agents. *IEEE Transactions on Robotics*, 33(5):1103–1123, 2017.

[16] Andrew R Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.

[17] PeterL Bartlett. Thesamplecomplexityofp atternclassification withneuralnetworks: Thesizeoftheweightsismo reimportantthan thesizeofthenetwork. *IEEETrans. Inf. Theory*, 44(2), 1998.

[18] Jürgen Batt and Gerhard Rein. A rigorous stability result for the Vlasov-Poisson system in three dimensions. *Annali di matematica pura ed applicata*, 164(1):133–154, 1993.

[19] Heinz H Bauschke and Patrick L Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.

[20] Richard Ernest Bellman. *Dynamic Programming*. Courier Dover Publications, 1957.

[21] Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

[22] Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.

[23] Jean-David Benamou, Guillaume Carlier, and Maxime Laborde. An augmented Lagrangian approach to Wasserstein gradient flows and applications. *ESAIM: Proceedings and Surveys*, 54:1–17, 2016.

[24] Yoshua Bengio, Nicolas Le Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex neural networks. In *Advances in neural information processing systems*, pages 123–130, 2006.

[25] Arne Beurling. An automorphism of product measures. *Annals of Mathematics*, pages 189–200, 1960.

[26] Rabi N Bhattacharya and Edward C Waymire. *Stochastic processes with applications*, volume 61. SIAM, 2009.

[27] Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.

[28] François Bolley, Ivan Gentil, and Arnaud Guillin. Convergence to equilibrium in Wasserstein distance for Fokker–Planck equations. *Journal of Functional Analysis*, 263(8):2430–2457, 2012.

[29] François Bolley, Ivan Gentil, and Arnaud Guillin. Uniform convergence to equilibrium for granular media. *Archive for Rational Mechanics and Analysis*, 208(2):429–445, 2013.

[30] François Bolley, Arnaud Guillin, and Florent Malrieu. Trend to equilibrium and particle approximation for a weakly selfconsistent vlasov-fokker-planck equation. *ESAIM: Mathematical Modelling and Numerical Analysis*, 44(5):867–884, 2010.

[31] François Bouchut and Jean Dolbeault. On long time asymptotics of the Vlasov-Fokker-Planck equation and of the Vlasov-Poisson-Fokker-Planck system with Coulombic and Newtonian potentials. *Differential and Integral Equations*, 8(3):487–514, 1995.

[32] Roger Brockett. Notes on stochastic processes on manifolds. In *Systems and Control in the Twenty-first Century*, pages 75–100. Springer, 1997.

[33] Roger Brockett. Notes on the control of the Liouville equation. In *Control of partial differential equations*, pages 101–129. Springer, 2012.

[34] Roger W Brockett. Optimal control of the Liouville equation. *AMS IP Studies in Advanced Mathematics*, 39:23, 2007.

[35] RW Brockett and JC Willems. Stochastic control and the second law of thermodynamics. In *1978 IEEE Conference on Decision and Control including the 17th Symposium on Adaptive Processes*, pages 1007–1011. IEEE, 1979.

[36] Peter J Bushell. Hilbert's metric and positive contraction mappings in a Banach space. *Archive for Rational Mechanics and Analysis*, 52(4):330–338, 1973.

[37] Alexandre Cabot, Hans Engler, and Sébastien Gadat. On the long time behavior of second order differential equations with asymptotically small dissipation. *Transactions of the American Mathematical Society*, 361(11):5983–6017, 2009.

[38] Kenneth F Caluya and Abhishek Halder. Finite horizon density control for static state feedback linearizable systems. *arXiv preprint arXiv:1904.02272*, 2019.

[39] Kenneth F Caluya and Abhishek Halder. Finite horizon density steering for multi-input state feedback linearizable systems. *arXiv preprint arXiv:1909.12511*, 2019.

[40] Kenneth F Caluya and Abhishek Halder. Finite horizon density steering for multi-input state feedback linearizable systems. *arXiv preprint arXiv:1909.12511, to appear in 2020 American Control Conference*, 2019.

[41] Kenneth F Caluya and Abhishek Halder. Gradient flow algorithms for density propagation in stochastic systems. *IEEE Transactions on Automatic Control*, 2019.

[42] Kenneth F Caluya and Abhishek Halder. Proximal recursion for solving the Fokker-Planck equation. In *2019 Annual American Control Conference (ACC)*, 2019.

[43] Kenneth F Caluya and Abhishek Halder. Proximal recursion for solving the Fokker-Planck equation. In *2019 American Control Conference (ACC)*, pages 4098–4103. IEEE, 2019.

[44] Kenneth F Caluya and Abhishek Halder. Wasserstein proximal algorithms for the Schrödinger bridge problem: Density control with nonlinear drift. *arXiv preprint arXiv:1912.01244*, 2019.

[45] Vincent Calvez and José A Carrillo. Volume effects in the keller–segel model:

energy estimates preventing blow-up. *Journal de mathématiques pures et appliquées*, 86(2):155–175, 2006.

[46] José A Carrillo, Robert J McCann, and Cédric Villani. Kinetic equilibration rates for granular media and related equations: entropy dissipation and mass transportation estimates. *Revista Matemática Iberoamericana*, 19(3):971–1018, 2003.

[47] José A Carrillo, Robert J McCann, and Cédric Villani. Contractions in the 2-wasserstein length space and thermalization of granular media. *Archive for Rational Mechanics and Analysis*, 179(2):217–263, 2006.

[48] José A Carrillo, Giuseppe Toscani, et al. Contractive probability metrics and asymptotic behavior of dissipative kinetic equations. 2007.

[49] Patrick Cattiaux, Arnaud Guillin, and Florent Malrieu. Probabilistic approach for granular media equations in the non-uniformly convex case. *Probability theory and related fields*, 140(1):19–40, 2008.

[50] Subhash Challa and Yaakov Bar-Shalom. Nonlinear filter design using Fokker-Planck-Kolmogorov probability density evolutions. *IEEE Transactions on Aerospace and Electronic Systems*, 36(1):309–315, 2000.

[51] Subrahmanyan Chandrasekhar. Stochastic problems in physics and astronomy. *Reviews of Modern Physics*, 15(1):1–89, 1943.

[52] Loïc Chaumont and Marc Yor. *Exercises in Probability: a guided tour from measure theory to random processes, via conditioning*. Number 35 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2012.

[53] Yongxin Chen, Tryphon Georgiou, and Michele Pavon. Entropic and displacement interpolation: a computational approach using the Hilbert metric. *SIAM Journal on Applied Mathematics*, 76(6):2375–2396, 2016.

[54] Yongxin Chen and Tryphon T Georgiou. Stochastic bridges of linear systems. *IEEE Transactions on Automatic Control*, 61(2):526–531, 2015.

[55] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Fast cooling for a system of stochastic oscillators. *Journal of Mathematical Physics*, 56(11):113302, 2015.

[56] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal steering of a linear stochastic system to a final probability distribution, Part I. *IEEE Transactions on Automatic Control*, 61(5):1158–1169, 2015.

[57] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal steering of a linear stochastic system to a final probability distribution, part ii. *IEEE Transactions on Automatic Control*, 61(5):1170–1180, 2015.

[58] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal steering of a linear stochastic system to a final probability distribution, Part II. *IEEE Transactions on Automatic Control*, 61(5):1170–1180, 2015.

[59] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. On the relation between optimal transport and schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169(2):671–691, 2016.

[60] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal transport over a linear dynamical system. *IEEE Transactions on Automatic Control*, 62(5):2137–2152, 2016.

[61] Xiang Cheng, Niladri S Chatterji, Peter L Bartlett, and Michael I Jordan. Underdamped langevin mcmc: A non-asymptotic analysis. In *Conference on Learning Theory*, pages 300–323. PMLR, 2018.

[62] Michael Chertkov and Vladimir Chernyak. Ensemble of thermostatically controlled loads: Statistical physics approach. *Scientific reports (Nature)*, 7(1):8673, 2017.

[63] Cheng-Chih Chien, Youping Zhang, and Petros A Ioannou. Traffic density control for automated highway systems. *Automatica*, 33(7):1273–1285, 1997.

[64] Gregory S Chirikjian and Alexander B Kyatkin. *Harmonic Analysis for Engineers and Applied Scientists: Updated and Expanded Edition*. Courier Dover Publications, 2016.

[65] Lenaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, pages 3036–3046, 2018.

[66] Lenaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. *arXiv preprint arXiv:2002.04486*, 2020.

[67] Lenaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pages 2933–2943, 2019.

[68] Julian D Cole. On a quasi-linear parabolic equation occurring in aerodynamics. *Quarterly of Applied Mathematics*, 9(3):225–236, 1951.

[69] Rinaldo M Colombo and Magali Lécureux-Mercier. Nonlocal crowd dynamics models for several populations. *Acta Mathematica Scientia*, 32(1):177–196, 2012.

[70] Giovanni Conforti and Michele Pavon. Extremal curves in Wasserstein space. In *International Conference on Geometric Science of Information*, pages 91–99. Springer, 2017.

[71] Dario Cordero-Erausquin, Wilfrid Gangbo, and Christian Houdré. Inequalities for generalized entropy and optimal transportation. *Contemporary Mathematics*, 353:73–94, 2004.

[72] John C Cox, Jonathan E Ingersoll Jr, and Stephen A Ross. A theory of the term structure of interest rates. *Econometrica*, 53(2):385–408, 1985.

[73] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013.

[74] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[75] Paolo Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied mathematics and Optimization*, 23(1):313–329, 1991.

[76] Paolo Dai Pra and Michele Pavon. On the Markov processes of Schrödinger, the Feynman-Kac formula and stochastic control. In *Realization and Modelling in System Theory*, pages 497–504. Springer, 1990.

[77] Fred Daum. Nonlinear filters: beyond the Kalman filter. *IEEE Aerospace and Electronic Systems Magazine*, 20(8):57–69, 2005.

[78] Pierre Degond. Mathematical modelling of microelectronics semiconductor devices. *AMS/IP Studies in Advanced Mathematics*, 15:77–110, 2000.

[79] Vaibhav Deshmukh, Karthik Elamvazhuthi, Shiba Biswal, Zahi Kakish, and Spring Berman. Mean-field stabilization of markov chain models for robotic swarms: Computational approaches and experimental results. *IEEE Robotics and Automation Letters*, 3(3):1985–1992, 2018.

[80] Jean Dolbeault. Stationary states in plasma physics: Maxwellian solutions of the Vlasov-Poisson system. *Mathematical Models and Methods in Applied Sciences*, 1(2):183–208, 1991.

[81] Manh Hong Duong, Mark A Peletier, and Johannes Zimmer. Conservative-

dissipative approximation schemes for a generalized Kramers equation. *Mathematical Methods in the Applied Sciences*, 37(16):2517–2540, 2014.

[82] Andreas Eberle, Arnaud Guillin, Raphael Zimmer, et al. Couplings and quantitative contraction rates for langevin dynamics. *Annals of Probability*, 47(4):1982–2010, 2019.

[83] Martin Ehrendorfer. The Liouville equation and its potential usefulness for the prediction of forecast skill. part I: Theory. *Monthly Weather Review*, 122(4):703–713, 1994.

[84] Karthik Elamvazhuthi, Piyush Grover, and Spring Berman. Optimal transport over deterministic discrete-time nonlinear systems using stochastic feedback laws. *IEEE control systems letters*, 3(1):168–173, 2018.

[85] Wendell H Fleming. Logarithmic transformations and stochastic control. In *Advances in Filtering and Optimal Stochastic Control*, pages 131–141. Springer, 1982.

[86] Robert Fortet. Résolution d'un systeme d'équations de m. Schrödinger. *J. Math. Pure Appl. IX*, 1:83–105, 1940.

[87] Sébastien Gadat, Fabien Panloup, and Sofiane Saadane. Stochastic heavy ball. *Electronic Journal of Statistics*, 12(1):461–529, 2018.

[88] David Gilbarg and Neil S Trudinger. *Elliptic partial differential equations of second order*. Springer, 2015.

[89] Peter W Glynn and Rob J Wang. On the rate of convergence to equilibrium for reflected Brownian motion. *Queueing Systems*, 89(1-2):165–197, 2018.

[90] Javier Gomez-Serrano, Carl Graham, and Jean-Yves Le Boudec. The bounded confidence model of opinion dynamics. *Mathematical Models and Methods in Applied Sciences*, 22(02):1150007, 2012.

[91] Xavier Goudou and Julien Munier. The gradient and heavy ball with friction dynamical systems: the quasiconvex case. *Mathematical Programming*, 116(1-2):173–191, 2009.

[92] Karolos M Grigoriadis and Robert E Skelton. Minimum-energy covariance controllers. *Automatica*, 33(4):569–578, 1997.

[93] Arnaud Guillin, Pierre Le Bris, et al. Convergence rates for the vlasov-fokker-planck equation and uniform in time propagation of chaos in non convex cases. *arXiv preprint arXiv:2105.09070*, 2021.

[94] Abhishek Halder and Raktim Bhattacharya. Beyond Monte Carlo: A computational framework for uncertainty propagation in planetary entry, descent and landing. In *AIAA Guidance, Navigation, and Control Conference*, page 8029, 2010.

[95] Abhishek Halder and Raktim Bhattacharya. Dispersion analysis in hypersonic flight during planetary entry using stochastic Liouville equation. *Journal of Guidance, Control, and Dynamics*, 34(2):459–474, 2011.

[96] Abhishek Halder and Raktim Bhattacharya. Model validation: A probabilistic formulation. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 1692–1697. IEEE, 2011.

[97] Abhishek Halder and Raktim Bhattacharya. Further results on probabilistic model validation in Wasserstein metric. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 5542–5547. IEEE, 2012.

[98] Abhishek Halder and Raktim Bhattacharya. Probabilistic model validation for uncertain nonlinear systems. *Automatica*, 50(8):2038–2050, 2014.

[99] Abhishek Halder, Xinbo Geng, PR Kumar, and Le Xie. Architecture and algorithms for privacy preserving thermal inertial load management by a load serving entity. *IEEE Transactions on Power Systems*, 32(4):3275–3286, 2016.

[100] Abhishek Halder and Tryphon T Georgiou. Gradient flows in uncertainty propagation and filtering of linear Gaussian systems. *2017 IEEE Conference on Decision and Control, arXiv preprint arXiv:1704.00102*, 2017.

[101] Abhishek Halder and Tryphon T Georgiou. Gradient flows in filtering and Fisher-Rao geometry. In *2018 Annual American Control Conference (ACC)*, pages 4281–4286. IEEE, 2018.

[102] Abhishek Halder and Tryphon T Georgiou. Proximal recursion for the Wonham filter. *2019 IEEE 58th Annual Conference on Decision and Control (CDC), arXiv:1909.05827*, 2019.

[103] Abhishek Halder, Kooktae Lee, and Raktim Bhattacharya. Optimal transport approach for probabilistic robustness analysis of F-16 controllers. *Journal of Guidance, Control, and Dynamics*, 38(10):1935–1946, 2015.

[104] Abhishek Halder and Eric DB Wendel. Finite horizon linear quadratic Gaussian density regulator with Wasserstein terminal cost. In *2016 American Control Conference (ACC)*, pages 7249–7254. IEEE, 2016.

[105] Heiko Hamann and Heinz Wörn. A framework of space–time continuous models for algorithm design in swarm robotics. *Swarm Intelligence*, 2(2-4):209–239, 2008.

[106] Alain Haraux. *Systèmes dynamiques dissipatifs et applications*, volume 17. Masson, 1991.

[107] Rolland L Hardy. Multiquadric equations of topography and other irregular surfaces. *Journal of Geophysical Research*, 76(8):1905–1915, 1971.

[108] Rolland Lee Hardy. Research results in the application of multiquadratic equations to surveying and mapping problems. *Surveying and Mapping*, 1975.

[109] J Michael Harrison. *Brownian models of performance and control*. Cambridge University Press, 2013.

[110] J Michael Harrison and Ruth J Williams. Multidimensional reflected Brownian motions having exponential stationary distributions. *The Annals of Probability*, pages 115–137, 1987.

[111] Siegfried Hess. Fokker-Planck-equation approach to flow alignment in liquid crystals. *Zeitschrift für Naturforschung A*, 31(9):1034–1037, 1976.

[112] Desmond J Higham, Xuerong Mao, and Andrew M Stuart. Strong convergence of Euler-type methods for nonlinear stochastic differential equations. *SIAM Journal on Numerical Analysis*, 40(3):1041–1063, 2002.

[113] David Hilbert. Über die gerade linie als kürzeste verbindung zweier punkte. *Mathematische Annalen*, 46(1):91–96, 1895.

[114] Eberhard Hopf. The partial differential equation $u_t + uu_x = \mu_{xx}$. *Communications on Pure and Applied mathematics*, 3(3):201–230, 1950.

[115] Anthony Hotz and Robert E Skelton. Covariance control theory. *International Journal of Control*, 46(1):13–32, 1987.

[116] Kaitong Hu, Zhenjie Ren, David Siska, and Lukasz Szpruch. Mean-field langevin dynamics and energy landscape of neural networks. *arXiv preprint arXiv:1905.07769*, 2019.

[117] Chaocheng Huang and Richard Jordan. Variational formulations for Vlasov–Poisson–Fokker–Planck systems. *Mathematical methods in the applied sciences*, 23(9):803–843, 2000.

[118] Martin Hutzenthaler and Arnulf Jentzen. *Numerical approximations of*

stochastic differential equations with non-globally Lipschitz continuous coefficients, volume 236. American Mathematical Society, 2015.

[119] Martin Hutzenthaler, Arnulf Jentzen, and Peter E Kloeden. Strong and weak divergence in finite time of Euler's method for stochastic differential equations with non-globally Lipschitz continuous coefficients. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2130):1563–1576, 2010.

[120] Martin Hutzenthaler, Arnulf Jentzen, Peter E Kloeden, et al. Strong convergence of an explicit numerical method for SDEs with nonglobally Lipschitz continuous coefficients. *The Annals of Applied Probability*, 22(4):1611–1641, 2012.

[121] Nobuyuki Ikeda. On the construction of two-dimensional diffusion processes satisfying Wentzell's boundary conditions and its application to boundary value problems. *Memoirs of the College of Science, University of Kyoto. Series A: Mathematics*, 33(3):367–427, 1961.

[122] Alberto Isidori. *Nonlinear control systems: an introduction*. Springer, 1985.

[123] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.

[124] Benton Jamison. The Markov processes of Schrödinger. *Probability Theory and Related Fields*, 32(4):323–331, 1975.

[125] Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker–Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.

[126] Yuri P Kalmykov and William T Coffey. Analytical solutions for rotational diffusion in the mean field potential: application to the theory of dielectric relaxation in nematic liquid crystals. *Liquid crystals*, 25(3):329–339, 1998.

[127] Edward J Kansa. Multiquadrics—a scattered data approximation scheme with applications to computational fluid-dynamics—ii solutions to parabolic, hyperbolic and elliptic partial differential equations. *Computers & Mathematics with Applications*, 19(8-9):147–161, 1990.

[128] EJ Kansa. Application of Hardy's multiquadric interpolation to hydrodynamics. Technical report, Lawrence Livermore National Lab., CA (USA), 1985.

[129] Samuel Karlin and Howard E Taylor. *A second course in stochastic processes*. Elsevier, 1981.

[130] Johan Karlsson and Axel Ringh. Generalized Sinkhorn iterations for regularizing inverse problems using optimal mass transport. *SIAM Journal on Imaging Sciences*, 10(4):1935–1962, 2017.

[131] Anna Kazeykina, Zhenjie Ren, Xiaolu Tan, and Junjian Yang. Ergodicity of the underdamped mean-field langevin dynamics. *arXiv preprint arXiv:2007.14660*, 2020.

[132] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

[133] Peter E Kloeden and Eckhard Platen. *Numerical solution of stochastic differential equations*, volume 23. Springer Science & Business Media, 2013.

[134] Walid Krichene and Peter L Bartlett. Acceleration and averaging in stochastic descent dynamics. In *Advances in Neural Information Processing Systems*, pages 6796–6806, 2017.

[135] Walid Krichene, Alexandre Bayen, and Peter L Bartlett. Accelerated mirror descent in continuous and discrete time. In *Advances in neural information processing systems*, pages 2845–2853, 2015.

[136] Walid Krichene, Kenneth F Caluya, and Abhishek Halder. Global convergence of second-order dynamics in two-layer neural networks. *arXiv preprint arXiv:2007.06852*, 2020.

[137] Lukasz Kruk, John Lehoczky, Kavita Ramanan, Steven Shreve, et al. An explicit formula for the Skorokhod map on [0, a]. *The Annals of Probability*, 35(5):1740–1768, 2007.

[138] Maxime Laborde. On some nonlinear evolution systems which are perturbations of Wasserstein gradient flows. *Topological Optimization and Optimal Transport: In the Applied Sciences*, 17:304, 2017.

[139] Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.

[140] Flavien Léger. A geometric perspective on regularized optimal transport. *Journal of Dynamics and Differential Equations*, pages 1–15, 2018.

[141] Flavien Léger and Wuchen Li. Hopf-Cole transformation via generalized Schrödinger bridge problem. *arXiv preprint arXiv:1901.09051*, 2019.

[142] Bas Lemmens and Roger Nussbaum. *Nonlinear Perron-Frobenius Theory*, volume 189. Cambridge University Press, 2012.

[143] Christian Léonard. From the Schrödinger problem to the Monge–Kantorovich problem. *Journal of Functional Analysis*, 262:1879–1920, 2012.

[144] Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.

[145] Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete & Continuous Dynamical Systems-A*, 34(4):1533–1574, 2014.

[146] Jr-Shin Li and Navin Khaneja. Ensemble control of bloch equations. *IEEE Transactions on Automatic Control*, 54(3):528–536, 2009.

[147] Wuchen Li, Penghang Yin, and Stanley Osher. Computations of optimal transport distance with Fisher information regularization. *Journal of Scientific Computing*, 75(3):1581–1595, 2018.

[148] Daniel Liberzon and Roger W Brockett. Nonlinear feedback systems perturbed by noise: Steady-state probability distributions and optimal control. *IEEE Transactions on Automatic Control*, 45(6):1116–1130, 2000.

[149] Yongdo Lim. Nonlinear equations based on jointly homogeneous mappings. *Linear Algebra and Its Applications*, 430(1):279–285, 2009.

[150] Vadim Linetsky. On the transition densities for reflected diffusions. *Advances in Applied Probability*, 37(2):435–460, 2005.

[151] Pierre-Louis Lions and Alain-Sol Sznitman. Stochastic differential equations with reflecting boundary conditions. *Communications on Pure and Applied Mathematics*, 37(4):511–537, 1984.

[152] PL Lions. Optimal control of reflected diffusion processes. In *Filtering and Control of Random Processes*, pages 157–163. Springer, 1984.

[153] Harald Luschgy and Gilles Pagès. Functional quantization of a class of Brownian diffusions: a constructive approach. *Stochastic Processes and their Applications*, 116(2):310–336, 2006.

[154] Florent Malrieu. Uniform propagation of chaos for McKean-Vlasov equations. `https://tinyurl.com/ycxk5g9t`, 2011. [Online; accessed 18-November-2018].

[155] Florent Malrieu et al. Convergence to equilibrium for granular media equations and their euler schemes. *Annals of applied Probability*, 13(2):540–560, 2003.

[156] Michael B Marcus and Jay Rosen. *Markov processes, Gaussian processes, and local times*, volume 100. Cambridge University Press, 2006.

[157] SS Mazurenko. The dynamic programming method in systems with states in the form of distributions. *Moscow University Computational Mathematics and Cybernetics*, 35(3):133, 2011.

[158] Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.

[159] Toshio Mikami. Variational processes from the weak forward equation. *Communications in Mathematical Physics*, 135(1):19–40, 1990.

[160] Toshio Mikami. Monge's problem with a quadratic cost by the zero-noise limit of h-path processes. *Probability Theory and Related Fields*, 129(2):245–260, 2004.

[161] Toshio Mikami and Michèle Thieullen. Optimal transportation problem by stochastic optimal control. *SIAM Journal on Control and Optimization*, 47(3):1127–1139, 2008.

[162] Jan Kloppenborg Møller and Henrik Madsen. *From state dependent diffusion to constant diffusion in stochastic differential equations by the Lamperti transform*. DTU Informatics, 2010.

[163] Bharat Monga, Gary Froyland, and Jeff Moehlis. Synchronizing and desynchronizing neural populations through phase distribution control. In *2018 Annual American Control Conference (ACC)*, pages 2808–2813. IEEE, 2018.

[164] W Muschik and B Su. Mesoscopic interpretation of Fokker-Planck equation describing time behavior of liquid crystal orientation. *The Journal of Chemical Physics*, 107(2):580–584, 1997.

[165] Yurii Nesterov. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. In *Sov. Math. Dokl*, volume 27, pages 372–376, 1983.

[166] Atsushi Nitanda, Denny Wu, and Taiji Suzuki. Particle dual averaging: Optimization of mean field neural networks with global convergence rate analysis. *arXiv preprint arXiv:2012.15477*, 2020.

[167] Roger D Nussbaum and David Hilbert. *Hilbert's projective metric and iterated nonlinear maps*, volume 391. American Mathematical Soc., 1988.

[168] Karl Oelschläger. A martingale approach to the law of large numbers for weakly interacting stochastic processes. *The Annals of Probability*, pages 458–479, 1984.

[169] Karl Oelschläger. A law of large numbers for moderately interacting diffusion processes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 69(2):279–322, 1985.

[170] Kazuhide Okamoto, Maxim Goldshtein, and Panagiotis Tsiotras. Optimal covariance control for stochastic systems under chance constraints. *IEEE Control Systems Letters*, 2(2):266–271, 2018.

[171] Kazuhide Okamoto and Panagiotis Tsiotras. Input hard constrained optimal covariance steering. *arXiv preprint arXiv:1903.10964, IEEE Conference on Decision and Control, Nice, France*, 2019.

[172] Kazuhide Okamoto and Panagiotis Tsiotras. Optimal stochastic vehicle path planning using covariance steering. *IEEE Robotics and Automation Letters*, 4(3):2276–2281, 2019.

[173] Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

[174] Felix Otto. The geometry of dissipative evolution equations: the porous medium equation. 2001.

[175] Thanu Padmanabhan. Statistical mechanics of gravitating systems. *Physics Reports*, 188(5):285–362, 1990.

[176] Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.

[177] Wooram Park, Jin Seob Kim, Yu Zhou, Noah J Cowan, Allison M Okamura, and Gregory S Chirikjian. Diffusion-based motion planning for a nonholonomic flexible needle model. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 4600–4605. IEEE, 2005.

[178] Wooram Park, Yan Liu, Yu Zhou, Matthew Moses, and Gregory S Chirikjian. Kinematic state estimation and motion planning for stochastic nonholonomic systems using the exponential map. *Robotica*, 26(4):419–434, 2008.

[179] Grigorios A Pavliotis. Stochastic processes and applications, 2011.

[180] Grigorios A Pavliotis. *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*, volume 60. Springer, 2014.

[181] Grigorios A Pavliotis. Stochastic processes and applications, 2015.

[182] Michele Pavon, Esteban G Tabak, and Giulio Trigila. The data-driven Schroedinger bridge. *arXiv preprint arXiv:1806.01364*, 2018.

[183] Benoît Perthame. Mathematical tools for kinetic equations. *Bulletin of the American Mathematical Society*, 41(2):205–244, 2004.

[184] Andrey Pilipenko. *An introduction to stochastic differential equations with reflection*, volume 1. Universitätsverlag Potsdam, 2014.

[185] Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.

[186] Hannes Risken. *Fokker-Planck equation: Methods of solution and applications*. Springer, 1996.

[187] Grant Rotskoff, Samy Jelassi, Joan Bruna, and Eric Vanden-Eijnden. Neuron birth-death dynamics accelerates gradient descent and converges asymptotically. In *International Conference on Machine Learning*, pages 5508–5517, 2019.

[188] Grant Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. In *Advances in neural information processing systems*, pages 7146–7155, 2018.

[189] Walter Rudin. *Principles of Mathematical Analysis*. Third Edition, McGraw-Hill, USA, 1976.

[190] Filippo Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7(1):87–154, 2017.

[191] Erwin Schrödinger. Über die umkehrung der naturgesetze. volume 10, pages 144–153. Sitzungsberichte der Preuss. Phys. Math. Klasse., 1931.

[192] Erwin Schrödinger. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. In *Annales de l'institut Henri Poincaré*, volume 2, pages 269–310, 1932.

[193] Bin Shi, Simon S Du, Weijie Su, and Michael I Jordan. Acceleration via symplectic discretization of high-resolution differential equations. In *Advances in Neural Information Processing Systems*, pages 5745–5753, 2019.

[194] Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *arXiv preprint arXiv:1805.01053*, 2018.

[195] Robert E Skelton and Masao Ikeda. Covariance controllers for linear continuous-time systems. *International Journal of Control*, 49(5):1773–1785, 1989.

[196] Robert E Skelton, T Iwasaki, and KM Grigoriadis. A unified algebraic approach to linear control design, 2013.

[197] Anatoliy V Skorokhod. Stochastic equations for diffusion processes in a bounded region. *Theory of Probability & Its Applications*, 6(3):264–274, 1961.

[198] Anatoliy V Skorokhod. Stochastic equations for diffusion processes in a bounded region. II. *Theory of Probability & Its Applications*, 7(1):3–23, 1962.

[199] Jean-Jacques E Slotine and Weiping Li. *Applied nonlinear control*, volume 199. Prentice hall Englewood Cliffs, NJ, 1991.

[200] Juan Soler, José A Carrillo, and Luis L Bonilla. Asymptotic behavior of an initial-boundary value problem for the Vlasov–Poisson–Fokker–Planck system. *SIAM Journal on Applied Mathematics*, 57(5):1343–1372, 1997.

[201] Elias M Stein and Rami Shakarchi. *Complex analysis*, volume 2. Princeton University Press, 2010.

[202] Daniel W Stroock. *Partial differential equations for probabilists*. Cambridge Univ. Press, 2008.

[203] Daniel W Stroock and SR Srinivasa Varadhan. Diffusion processes with boundary conditions. *Communications on Pure and Applied Mathematics*, 24(2):147–225, 1971.

[204] Weijie Su, Stephen Boyd, and Emmanuel Candès. A differential equation for modeling Nesterov's accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems*, pages 2510–2518, 2014.

[205] Yifei Sun and Mrinal Kumar. Uncertainty forecasting in the perturbed two-body problem via tensor decomposition. In *American Control Conference (ACC), 2016*, pages 5431–5436. IEEE, 2016.

[206] Ilya Sutskever, James Martens, George Dahl, and Geoffery Hinton. On the importance of momentum and initialization in deep learning. In *30th International Conference on Machine Learning*, pages 404–439, 2013.

[207] Anthony C Thompson. On certain contraction mappings in a partially ordered vector space. *Proceedings of the American Mathematical Society*, 14(3):438–443, 1963.

[208] Chad M Topaz, Andrea L Bertozzi, and Mark A Lewis. A nonlocal continuum model for biological aggregation. *Bulletin of mathematical biology*, 68(7):1601, 2006.

[209] Harold Dean Victory Jr and Brian P O'Dwyer. On classical solutions of Vlasov-Poisson Fokker-Planck systems. *Indiana University mathematics journal*, pages 105–156, 1990.

[210] Cédric Villani. *Topics in optimal transportation*. Number 58. American Mathematical Soc., 2003.

[211] Cédric Villani. Trend to equilibrium for dissipative equations, functional inequalities and mass transportation. *Contemporary Mathematics*, 353:95, 2004.

[212] Cédric Villani. Hypocoercive diffusion operators. In *International Congress of Mathematicians*, volume 3, pages 473–498, 2006.

[213] Cédric Villani. Hypocoercivity. *arXiv preprint math/0609050*, 2006.

[214] Cédric Villani. Hypocoercivity. *Memoirs of the American Mathematical Society*, 202(950), 2009.

[215] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.

[216] A Wakolbinger. Schrödinger bridges from 1931 to 1991. In *Proc. of the 4th Latin American Congress in Probability and Mathematical Statistics, Mexico City*, pages 61–79, 1990.

[217] Shinzo Watanabe. On stochastic differential equations for multi-dimensional diffusion processes with boundary conditions. *Journal of Mathematics of Kyoto University*, 11(1):169–180, 1971.

[218] Andre Wibisono, Ashia C Wilson, and Michael I Jordan. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, 113(47):E7351–E7358, 2016.

[219] David Williams. *Probability with martingales*. Cambridge university press, 1991.

[220] Blake Woodworth, Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. *arXiv preprint arXiv:1906.05827*, 2019.

[221] Jean-Claude Zambrini. Stochastic mechanics according to E. Schrödinger. *Physical Review A*, 33(3):1532, 1986.

[222] Jean-Claude Zambrini. Variational processes and stochastic versions of mechanics. *Journal of Mathematical Physics*, 27(9):2307–2330, 1986.

[223] Eberhard Zeidler. *Applied functional analysis: main principles and their applications*, volume 109. Springer Science & Business Media, 1995.

[224] Guoming Zhu, Karolos M Grigoriadis, and Robert E Skelton. Covariance control design for hubble space telescope. *Journal of Guidance, Control, and Dynamics*, 18(2):230–236, 1995.