

Neural Schrödinger Bridge With Sinkhorn Losses: Application to Data-Driven Minimum Effort Control of Colloidal Self-Assembly

Iman Nodozi^{ID}, *Graduate Student Member, IEEE*, Charlie Yan, Mira Khare^{ID},
Abhishek Halder^{ID}, *Senior Member, IEEE*, and Ali Mesbah^{ID}, *Senior Member, IEEE*

Abstract—We show that the minimum effort control of colloidal self-assembly (SA) can be naturally formulated in the order-parameter space as a generalized Schrödinger bridge problem (GSBP)—a class of fixed-horizon stochastic optimal control problems that originated in the works of Erwin Schrödinger in the early 1930s. In recent years, this class of problems has seen a resurgence of research activities in the control and machine-learning communities. Different from the existing literature on the theory and computation for such problems, the controlled drift and diffusion coefficients for colloidal SA are typically nonaffine in control and are difficult to obtain from physics-based modeling. We deduce the conditions of optimality for such generalized problems and show that the resulting system of equations is structurally very different from the existing results in a way that standard computational approaches no longer apply. Thus motivated, we propose a data-driven learning and control framework, named “neural Schrödinger bridge,” to solve such generalized Schrödinger bridge problems by innovating on recent advances in neural networks (NNs). We illustrate the effectiveness of the proposed framework using a numerical case study of colloidal SA. We learn the controlled drift and diffusion coefficients as two NNs using molecular dynamics (MD) simulation data and then use these two to train a third network with Sinkhorn losses designed for distributional endpoint constraints, specific for this class of control problems.

Index Terms—Colloidal self-assembly (SA), physics-informed neural networks (PINNs), Schrödinger bridge, Sinkhorn loss, stochastic optimal control.

I. INTRODUCTION

MOTIVATED by feedback control of colloidal self-assembly (SA), this work focuses on learning the solution of the nonlinear stochastic optimal control problems

Manuscript received 25 July 2023; revised 13 October 2023; accepted 23 November 2023. Date of publication 7 December 2023; date of current version 25 April 2024. This work was supported by NSF under Award 2112754 and Award 2112755. Recommended by Associate Editor M. Abbaszadeh. (*Corresponding author: Abhishek Halder.*)

Iman Nodozi and Charlie Yan are with the Department of Electrical and Computer Engineering, University of California at Santa Cruz, Santa Cruz, CA 95064 USA (e-mail: inodozi@ucsc.edu; cyan140@ucsc.edu).

Mira Khare and Ali Mesbah are with the Department of Chemical and Biomolecular Engineering, University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: mira_khare@berkeley.edu; mesbah@berkeley.edu).

Abhishek Halder is with the Department of Aerospace Engineering, Iowa State University, Ames, IA 50011 USA (e-mail: ahalder@iastate.edu).

Digital Object Identifier 10.1109/TCST.2023.3337588

over a given fixed time horizon $[0, T]$ of the form

$$\inf_{u \in \mathcal{U}} \mathbb{E}_{\mu^u} \left[\int_0^T \frac{1}{2} \|u(t, \mathbf{x})\|_2^2 dt \right] \quad (1a)$$

$$\text{subject to } d\mathbf{x} = \mathbf{f}(t, \mathbf{x}, \mathbf{u})dt + \sqrt{2} \mathbf{g}(t, \mathbf{x}, \mathbf{u})d\mathbf{w} \quad (1b)$$

$$\mathbf{x}(t=0) \sim \mu_0 \text{ (given)}, \mathbf{x}(t=T) \sim \mu_T \text{ (given)} \quad (1c)$$

where μ_0, μ_T denote the prescribed probability measures over the state space $\mathcal{X} \subseteq \mathbb{R}^n$ at $t=0$ and $t=T$, respectively. The constraint in (1b) is a controlled Itô stochastic differential equation (SDE) with the state vector $\mathbf{x} \in \mathcal{X}$, the control vector $\mathbf{u} \in \mathbb{R}^m$, and the standard Wiener process $\mathbf{w} \in \mathbb{R}^p$. For the solution to the SDE (1b) to be for colloidal SA systems, the state vector $\mathbf{x}$ represents suitable order parameters. The *drift coefficient* $\mathbf{f}$ is a vector field given by mapping $\mathbf{f} : [0, T] \times \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}^n$, and the *diffusion coefficient* $\mathbf{g}$ is a matrix field given by mapping $\mathbf{g} : [0, T] \times \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}^{n \times p}$. For the SDE solutions to be well-posed, we will detail suitable smoothness assumptions on $\mathbf{f}$ and $\mathbf{g}$.

Associated with the diffusion coefficient $\mathbf{g}$, is a *diffusion tensor* $\mathbf{G} := \mathbf{g}\mathbf{g}^\top \in \mathbb{S}_+^n$, which being an outer product, is a symmetric positive semidefinite matrix field $\mathbf{G} : [0, T] \times \mathcal{X} \times \mathcal{U} \mapsto \mathbb{S}_+^n$. In (1a), we suppose that the set of admissible controls $\mathcal{U}$ comprises of finite energy Markovian inputs within a prescribed time horizon, i.e.,

$$\mathcal{U} := \{\mathbf{u} : [0, T] \times \mathcal{X} \mapsto \mathbb{R}^m \mid \langle \mathbf{u}, \mathbf{u} \rangle < \infty\} \quad (2)$$

where $\langle \cdot, \cdot \rangle$ denotes the standard Euclidean inner product. The symbol $\mathbb{E}_{\mu^u}[\cdot]$ in (1a) denotes the mathematical expectation with respect to the controlled state probability measure μ^u , that is, $\mathbb{E}_{\mu^u}[\cdot] := \int (\cdot) d\mu^u$. The superscript $\mathbf{u}$ in μ^u indicates that the joint measure depends on the choice of control $\mathbf{u}$. Thus, the objective in (1a) is to minimize the control effort in steering the state statistics from μ_0 to μ_T under a prespecified time horizon and controlled stochastic dynamics constraints, over all admissible control policies $\mathbf{u}(t, \mathbf{x})$ in $\mathcal{U}$.

In feedback control of colloidal SA systems, the objective generally is to design control policies that steer the system from an initial disordered stochastic state to a desired terminal ordered crystalline stochastic state [1], [2]. These stochastic states are naturally encoded in terms of suitable order parameters. As such, formulation (1) is particularly appealing in this

context because it allows for directly shaping the multivariate distribution of order parameters via optimal control synthesis. The drift and diffusion coefficients $\mathbf{f}$, $\mathbf{g}$ in equation (1b) allow for the representation of the free energy landscape, which is crucial for circumventing kinetic traps or local minima when directing the system toward a desired end state, typically a global minimum within the solution space.

However, in practice, the drift and diffusion coefficients are difficult to model from first principles. This is because accurately capturing the interplay between various forces and interactions, such as van der Waals forces, electrostatic interactions, and solvent-mediated interactions, is challenging. As a result, empirical or semi-empirical approaches [3], as well as coarse-grained or phenomenological models [4], are often employed to approximate these coefficients based on either experimental data or molecular dynamics (MD) simulation data.

Another modeling difficulty specific to colloidal SA is that both $\mathbf{f}$, $\mathbf{g}$ are typically nonlinear in state $\mathbf{x}$, as well as nonaffine in control $\mathbf{u}$. Furthermore, $\mathbf{f}$ and $\mathbf{g}$ both have explicit time dependence in practice. To circumvent these modeling issues, in this work, we propose a learning and control framework where $\mathbf{f}$ and $\mathbf{g}$ are learned from high-fidelity MD simulation data as the outputs of neural network (NN) representations $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$, respectively. With these learned representations for $\mathbf{f}$ and $\mathbf{g}$, we propose a computational framework-based on another NN-to numerically solve (1) for control synthesis.

A. Relation to the SBP

We refer to (1) as a *generalized Schrödinger bridge problem (GSBP)* since it is related to distributional two-point boundary value problems originating in two papers of Erwin Schrödinger in 1931–32 [5], [6]. The qualifier “generalized” points to the presence of prior dynamics given by the controlled drift-diffusion coefficient pair $(\mathbf{f}, \mathbf{g})$, which generalizes the setting considered in Schrödinger’s original investigations [5], [6]. In the special case $\mathbf{f} \equiv \mathbf{u}$, $\mathbf{g} \equiv \mathbf{I}_n$, formulation (1) reduces to the classical Schrödinger Bridge Problem (SBP) à la Schrödinger. From this perspective, classical SBP is the problem of minimum effort distribution steering with zero prior drift, i.e., the problem of controlling Brownian motion with endpoint distribution constraints.

A different way to interpret classical SBP is to view it as a stochastic dynamic version of the optimal mass transport (OMT) problem. The dynamic OMT [7] is a special case of (1) with $\mathbf{f} \equiv \mathbf{u}$, $\mathbf{g} \equiv \mathbf{0}$. For details on these connections from a stochastic control perspective, we refer the readers to [8]. In recent years, SBPs and their generalizations have come to prominence in both control [8], [9], [10], [11] and machine learning [12], [13], [14], [15] communities. In particular, a data-driven maximum likelihood sampling solution of the classical SBP (i.e., with $\mathbf{f} \equiv \mathbf{u}$, $\mathbf{g} \equiv \mathbf{I}_n$) was proposed in [16] assuming availability of the samples from the endpoint measures μ_0, μ_T . Similar lines of ideas were pursued in [17] and [18].

While solution methods for the GSBP (1) in general are not available in the current literature, specialized algorithms for

particular forms of $\mathbf{f}$, $\mathbf{g}$ have appeared. For instance, Caluya and Halder [9] considered the case when the drift coefficient $\mathbf{f}$ is control affine and the diffusion coefficient $\mathbf{g}$ is $C([0, T])$ matrix that is independent of state and input, i.e.,

$$\mathbf{f}(t, \mathbf{x}, \mathbf{u}) \equiv \tilde{\mathbf{f}}(t, \mathbf{x}) + \mathbf{B}(t)\mathbf{u}, \quad \mathbf{g}(t, \mathbf{x}, \mathbf{u}) \equiv \mathbf{B}(t) \in \mathbb{R}^{n \times m}.$$

In this case, $m = p$ and the stochastic process noise enters through the input channels (e.g., modeling disturbance in forcing and/or actuation uncertainties). The results in [9] showed that if $\tilde{\mathbf{f}}$ is gradient of a potential, or if $(\tilde{\mathbf{f}}, \mathbf{B}(t))$ is of mixed conservative-dissipative form, then certain proximal recursions can be designed to numerically solve the corresponding GSBP. In [10], this result was extended for the case when additional (deterministic) state constraints are present.

GSBPs with nonlinear drifts and full-state feedback linearizable structures were considered in [19] and [20]. GSBP instances for both first- and second-order noisy nonuniform Kuramoto oscillator models were solved in [21] using Feynman-Kac path integral techniques. Closest to the GSBP (1) is the work in [22], which considered *control nonaffine* drift and diffusion coefficients and showed that the conditions of optimality involves additional coupled PDEs compared to the control-affine case. However, the developments in [22] were still model-based. Data-driven solution of control nonaffine GSBPs at the level of generality (1), as pursued in this work, is novel with respect to the existing literature.

B. Related Works on Control of Colloidal SA

Feedback control has emerged as a promising approach to enhance the reproducibility of colloidal SA systems toward desired structures. Previous studies [23], [24] demonstrated the effectiveness of proportional-integral control on simple test systems. However, applying such basic control approaches to complex colloidal SA systems with possible kinetically arrested dynamics may not yield satisfactory results. Alternative approaches like model predictive control (MPC) or dynamic programming have been suggested. For instance, Tang et al. [3] present an MPC approach based on energy landscapes estimated from MD simulations. However, this method can be computationally demanding for large-scale systems or systems with complex interactions, especially considering that the solution time for MPC might exceed the sampling time of SA, particularly for fast dynamics. This challenge becomes even more compounded as the size and complexity of the SA model grows.

On the other hand, Tang et al. [4] utilize a dynamic programming-based approach, which results in a lookup table of optimal actions for given states. Despite its theoretical elegance, dynamic programming suffers from the “curse of dimensionality,” rendering it impractical for systems of higher complexity due to the exponential growth in computational resources required. For both these methods, the accuracy of the control relies heavily on the quality of the underlying model. To this end, model-free reinforcement learning can alleviate modeling challenges in optimal control of colloidal SA systems [25], [26]. Furthermore, recent advances in NNs have provided a promising alternative for modeling the hidden physics of stochastic dynamic systems without making

assumptions about the final equations representing the physics of the system (e.g., [27], [28], [29], [30]). We employ the NNs, $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$, to represent the energy and diffusion landscapes of a colloidal SA control problem, critical to guiding the system to a desired final state. Our focus shifts to governing the temporal progression of the joint probability distribution encompassing the states involved in colloidal SA. Building upon previous work [22], we incorporate data-driven models based on physics-informed NNs (PINNs) and trained on high-fidelity MD simulation data, which allows for a more realistic representation of the colloidal SA process.

C. Contributions

This article makes the following specific contributions.

- 1) Building on our prior work [22], we propose that the controlled colloidal SA can be naturally formulated as a distribution steering problem in a suitable order-parameter space. This offers a newfound connection between the controlled colloidal SA and a nonstandard stochastic optimal control problem with hard constraints on the endpoint state statistics. The resulting stochastic optimal control problem takes the form of a control nonaffine GSBP.
- 2) To the best of the authors' knowledge, this is the first work to derive and numerically solve the conditions of optimality for control-nonaffine GSBP in multidimensional state-space. As detailed in Section II, the resulting system of equations is fundamentally different from the control-affine SBPs in that the optimal control is no longer an explicit functional of the (sub)gradient of the associated value function solving a Hamilton–Jacobi–Bellman (HJB) PDE. Instead, the m -dimensional optimal control $\mathbf{u}_{\text{opt}}$ here solves a system of m PDEs, which are coupled nonlinearly with two more PDEs and endpoint boundary conditions. As a result, existing approaches from the literature such as the Hopf–Cole transform [31], [32] followed by a contractive fixed point recursion [9], [33], or Feynman–Kac path integral techniques [21], cannot be used to numerically solve our system of equations. Leveraging recent advances in NNs, we propose a computational framework to learn the solution for this system of $m + 2$ PDEs and boundary conditions.
- 3) Our proposed computational approach, dubbed “neural Schrödinger bridge,” is “neural” in two ways: 1) a pair of NNs are trained to approximate the $\mathbf{f}$ and $\mathbf{g}$ in (1) using MD simulation data and 2) the GSBP optimality conditions, derived as functions of these NN representations, are further solved via a PINN [34], [35]. However, standard PINNs with mean squared error (MSE) losses are not appropriate to enforce distributional endpoint constraints (1c). To address this, we propose a PINN with Sinkhorn a.k.a. entropy-regularized Wasserstein losses for these constraints, and differentiate through these losses for training. The resulting architecture could be of independent interest.

We clarify here that, from a methodological viewpoint, the proposed framework is different from two recent works

[36] and [37], which also bring together SBPs and NNs. In [36], the main idea was to learn the *uncontrolled* $\mathbf{f}$, $\mathbf{g}$ as NNs, i.e., to learn an *unforced* neural SDE using the population samples via SBP. The unforced SDE was learned via a stochastic version of the principle of least action, i.e., by appealing to how SBP can be seen as a stochastic version of dynamic OMT, as we explained in Section I-A. The work in [37] proposed learning a classical SBP between unpaired images. Different from these works, our colloidal SA context requires learning the *controlled* $\mathbf{f}$, $\mathbf{g}$ as controlled neural SDEs before proceeding for optimal control synthesis—the latter is an instance of GSBP, which is then solved via a new variant of PINN that we propose herein.

D. Organization

In Section II, we define the GSBP (1) in terms of NN representations of $\mathbf{f}$ and $\mathbf{g}$. In Section III, we then discuss the solution of the GSBP conditions of optimality using a novel PINN with Sinkhorn losses. A detailed numerical case study of a colloidal SA system in an isothermal-isobaric (NPT) ensemble is then presented in Section IV. Section V concludes the article.

II. NEURAL SCHRÖDINGER BRIDGE

Our overall approach is to learn $(\mathbf{f}, \mathbf{g})$ as fully connected feed-forward NN representations, denoted by $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$, respectively. Both these NNs are designed to be functions of the current time $t \in [0, T]$, the system state $\mathbf{x}$, and the control input $\mathbf{u}$. These two networks are trained to predict the future states of the system based on the tuple $(t, \mathbf{x}, \mathbf{u})$. Training of these networks using MD simulations is detailed in Section IV. Fig. 1 gives an overview of the proposed learning and control framework. We next state the smoothness of the learned NN representations required for the control problem (i.e., the GSBP) to be well-posed.

A. Smoothness of the Learned $\mathbf{f}$ and $\mathbf{g}$

We consider both $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$ to have tangent hyperbolic, i.e., $\tanh(\cdot)$ activation functions. Tangent hyperbolic nonlinearities are known to be slope-restricted [38, Prop. 2]. As a result, the output of a fully connected feed-forward NN with $\tanh$ activation remains componentwise slope-restricted. Consequently, $\mathbf{f}$, $\mathbf{g}$ being the respective outputs of the networks $\mathcal{N}_{\text{Drift}}$, $\mathcal{N}_{\text{Diffusion}}$, are guaranteed [39, Th. 2] to be Lipschitz continuous.

Motivated by the Lipschitz continuity of the outputs of $\mathcal{N}_{\text{Drift}}$, $\mathcal{N}_{\text{Diffusion}}$ for an admissible Markovian policy $\mathbf{u}(t, \mathbf{x}) \in \mathcal{U}$, we assume that the coefficients $\mathbf{f}$ and $\mathbf{g}$ satisfy

(A1) *nonexplosion and Lipschitz conditions*: there exist constants c_1, c_2 such that

$$\|\mathbf{f}(t, \mathbf{x}, \mathbf{u}(t, \mathbf{x}))\|_2 + \|\mathbf{g}(t, \mathbf{x}, \mathbf{u}(t, \mathbf{x}))\|_2 \leq c_1(1 + \|\mathbf{x}\|_2)$$

and that

$$\|\mathbf{f}(t, \mathbf{x}, \mathbf{u}(t, \mathbf{x})) - \mathbf{f}(t, \tilde{\mathbf{x}}, \mathbf{u}(t, \tilde{\mathbf{x}}))\|_2 \leq c_2\|\mathbf{x} - \tilde{\mathbf{x}}\|_2$$

for all $\mathbf{x}, \tilde{\mathbf{x}} \in \mathcal{X}$, $t \in [0, T]$ and




Fig. 1. Overview of the proposed learning and control framework for solving the GSBP (1) for colloidal SA. Here, ρ_0 and ρ_T denote the pdf's associated with the endpoint measures μ_0 and μ_T , respectively.

(A2) *uniformly lower bounded diffusion*: there exists constant c_3 such that the diffusion tensor $\mathbf{G} = \mathbf{g}\mathbf{g}^\top$ satisfies

$$\langle \mathbf{x}, \mathbf{G}(t, \mathbf{x}, \mathbf{u}(t, \mathbf{x}))\mathbf{x} \rangle \geq c_3 \|\mathbf{x}\|_2^2$$

for all $t \in [0, T]$.

The assumption (A1) guarantees [40, p. 66] existence-uniqueness for the sample path of the SDE (1b). The assumptions (A1), (A2) together guarantee [41, Ch. 1] that the generator associated with (1b) yields absolutely continuous probability measures μ^u for all $t > 0$ provided the prescribed initial probability measure $\mu_0 := \mu^u(t=0, \mathbf{x})$ is absolutely continuous.

In this work, we assume that the given endpoint measures μ_0, μ_T are absolutely continuous, i.e., $\mu_0 = \rho_0(\mathbf{x})d\mathbf{x}$, $\mu_T = \rho_T(\mathbf{x})d\mathbf{x}$ where ρ_0, ρ_T are the corresponding endpoint joint state probability density functions (pdfs). If the solution for (1) exists, then under the stated regularity assumptions on $\mathbf{f}$ and $\mathbf{g}$, the corresponding controlled measure $\mu^u(t, \mathbf{x})$ will remain absolutely continuous with $d\mu^u(t, \mathbf{x}) = \rho^u(t, \mathbf{x})d\mathbf{x}$ for admissible $\mathbf{u} \in \mathcal{U}$. We next discuss reformulating (1) in terms of the controlled joint state pdf ρ^u .

B. PDF Steering Problem

To state the pdf steering problem, we set up some notations. We use the symbol nabla (∇) to denote the gradient with respect to its subscript vector. So for $\mathbf{x} \in \mathbb{R}^n$, we have

$$\text{gradient operator } \nabla_{\mathbf{x}} := \begin{pmatrix} \frac{\partial}{\partial x_1} \\ \vdots \\ \frac{\partial}{\partial x_n} \end{pmatrix}$$

$$\text{divergence operator } \nabla_{\mathbf{x}} \cdot := \frac{\partial}{\partial x_1} + \dots + \frac{\partial}{\partial x_n}$$

$$\text{Laplacian operator } \Delta_{\mathbf{x}} := \nabla_{\mathbf{x}} \cdot \nabla_{\mathbf{x}} = \frac{\partial^2}{\partial x_1^2} + \dots + \frac{\partial^2}{\partial x_n^2}.$$

For matrices $\mathbf{P}, \mathbf{Q}$ with commensurate dimensions and respective (i, j) th entries P_{ij}, Q_{ij} , their Frobenius a.k.a.

Hilbert-Schmidt inner product

$$\langle \mathbf{P}, \mathbf{Q} \rangle := \text{trace}(\mathbf{P}^\top \mathbf{Q}) = \sum_{i,j} P_{ij} Q_{ij}. \quad (3)$$

We use the symbol **Hess** to denote the Euclidean Hessian operator defined for any real-valued twice differentiable function $h : \mathcal{X} \subseteq \mathbb{R}^n \mapsto \mathbb{R}$, as

$$\mathbf{Hess}(h) := \begin{pmatrix} \frac{\partial^2 h}{\partial x_1^2} & \frac{\partial^2 h}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 h}{\partial x_1 \partial x_n} \\ \frac{\partial^2 h}{\partial x_2 \partial x_1} & \frac{\partial^2 h}{\partial x_2^2} & \cdots & \frac{\partial^2 h}{\partial x_2 \partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 h}{\partial x_n \partial x_1} & \frac{\partial^2 h}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 h}{\partial x_n^2} \end{pmatrix} \quad \forall \mathbf{x} \in \mathcal{X}.$$

In general, the entries of **Hess**(h) depend on $\mathbf{x}$, i.e., **Hess** returns a symmetric matrix field. The (i, j) th entry of the operator **Hess** is $(\partial^2 / \partial x_i \partial x_j)$.

Following (3), we define the Frobenius a.k.a. Hilbert-Schmidt inner product between the operator **Hess** and a matrix field $\mathbf{Q}(\mathbf{x})$ where $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$, as

$$\langle \mathbf{Hess}, \mathbf{Q}(\mathbf{x}) \rangle := \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} Q_{ij}(\mathbf{x}). \quad (4)$$

With the assumptions in Section II-A, the GSBP (1) can be rewritten as a state pdf steering problem

$$\inf_{(\rho^u, \mathbf{u})} \int_0^T \int_{\mathcal{X}} \frac{1}{2} \|\mathbf{u}(t, \mathbf{x})\|_2^2 \rho^u(t, \mathbf{x}) d\mathbf{x} dt \quad (5a)$$

$$\text{subject to } \frac{\partial \rho^u}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\rho^u \mathbf{f}) + \langle \mathbf{Hess}, \mathbf{G} \rho^u \rangle \quad (5b)$$

$$\rho^u(0, \mathbf{x}) = \rho_0, \quad \rho^u(T, \mathbf{x}) = \rho_T. \quad (5c)$$

The constraint (5b) is the controlled Fokker-Planck-Kolmogorov (FPK) PDE which governs the flow of the joint state pdf ρ^u associated with the SDE (1b). For a derivation of (5b) from (1b), see e.g., [42, Prop. 3.3].

For the term $\langle \text{Hess}, G\rho^u \rangle$ in (5b), note from (4) that

$$\langle \text{Hess}, G\rho^u \rangle = \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} (G_{ij}(t, \mathbf{x}, \mathbf{u}(t, \mathbf{x})) \rho^u(t, \mathbf{x})).$$

The boundary conditions (5c) at $t = 0$ and $t = T$ involve the prescribed initial and terminal joint state PDFs ρ_0 and ρ_T , respectively.

We note that when $\mathbf{f} \equiv \mathbf{u}$, $\mathbf{g}$ (and hence $\mathbf{G} \equiv \mathbf{I}_n$), then (5b) reduces to the controlled heat PDE, and problem (5a) reduces to the classical SBP, as mentioned in Section I-A. Furthermore, when $\mathbf{f} \equiv \mathbf{u}$, $\mathbf{g} \equiv \mathbf{0}$, then (5b) reduces to the Liouville PDE [43] for integrator dynamics $\dot{\mathbf{x}} = \mathbf{u}$, and problem (5a) reduces to the dynamic OMT, as mentioned in Section I-A.

Remark 1: To better understand the correspondence between (1) and (5a), notice that (5a) is simply a re-writing of (1a) by “opening up” the expectation operator with respect to the controlled state probability measure $d\mu^u(t, \mathbf{x}) = \rho^u(t, \mathbf{x})d\mathbf{x}$. The constraint (5b) is the PDF dynamics induced by the sample path dynamics (1b). Intuitively, the term $(1/2)\|\mathbf{u}\|_2^2 \rho^u d\mathbf{x}$ is a generalized kinetic energy, and the state-time integral (5a) encodes total control effort over the finite horizon $[0, T]$. The FPK PDE (5b) is a continuity equation expressing the conservation of probability mass under the drift coefficient $\mathbf{f}$ and the diffusion coefficient $\mathbf{g}$ (thus the diffusion tensor $\mathbf{G}$).

C. Existence and Uniqueness of Solution

Under the assumptions stated already in Section II-A, the controlled pdf ρ^u exists for $\mathbf{u} \in \mathcal{U}$. For the existence-uniqueness of solution for the variational problem (5a), we further assume that

(A3) the pdf ρ^u remains positive and continuous for all $t \in [0, T]$.

Then, following [44, Th. 3.2], [45], problem (5a) is guaranteed to admit a unique solution; see also [46, Sec. 10].

We next deduce the first-order optimality conditions for the GSBP (5a) in the form of a coupled system of $m + 2$ PDEs with boundary conditions, where m is the number of control inputs. With respect to the existing literature on the conditions of optimality for GSBPs, this system of PDEs for nonaffine control is the most general, and is a new result.

D. Conditions for Optimality

We start with the Lagrangian associated with the GSBP (5a)

$$\begin{aligned} \mathcal{L}(\rho^u, \mathbf{u}, \psi) \\ := \int_0^T \int_{\mathcal{X}} \left\{ \frac{1}{2} \|\mathbf{u}(t, \mathbf{x})\|_2^2 \rho^u(t, \mathbf{x}) + \psi(t, \mathbf{x}) \right. \\ \left. \times \left(\frac{\partial \rho^u}{\partial t} + \nabla_{\mathbf{x}} \cdot (\rho^u \mathbf{f}) - \langle \text{Hess}, G\rho^u \rangle \right) \right\} d\mathbf{x} dt \end{aligned} \quad (6)$$

where $\psi(t, \mathbf{x})$ is a $C^2([0, T]; \mathcal{X})$ Lagrange multiplier. Let

$$\begin{aligned} \mathcal{P}_{0T}(\mathcal{X}) := \left\{ \rho(t, \mathbf{x}) \geq 0 \mid \int_{\mathcal{X}} \rho d\mathbf{x} = 1 \right. \\ \left. \rho(t = 0, \mathbf{x}) = \rho_0, \quad \rho(t = T, \mathbf{x}) = \rho_T \right\}. \end{aligned} \quad (7)$$

Performing the unconstrained minimization of the Lagrangian $\mathcal{L}$ over $\mathcal{P}_{0T}(\mathcal{X}) \times \mathcal{U}$, where $\mathcal{U}$ is given in (2), we get the following result.

Theorem 1 (Optimal Control and Optimal State pdf): Let the set of feasible Markovian controls be given by (2). Then the pair $(\rho_{\text{opt}}^u(t, \mathbf{x}), \mathbf{u}_{\text{opt}}(t, \mathbf{x}))$ that solves (5a), must satisfy the following system of $m + 2$ coupled PDEs:

$$\frac{\partial \psi}{\partial t} = \frac{1}{2} \|\mathbf{u}_{\text{opt}}\|_2^2 - \langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle - \langle \mathbf{G}, \text{Hess}(\psi) \rangle \quad (8a)$$

$$\frac{\partial \rho_{\text{opt}}^u}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\rho_{\text{opt}}^u \mathbf{f}) + \langle \text{Hess}, G\rho_{\text{opt}}^u \rangle \quad (8b)$$

$$\mathbf{u}_{\text{opt}} = \nabla_{\mathbf{u}_{\text{opt}}} (\langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle + \langle \mathbf{G}, \text{Hess}(\psi) \rangle) \quad (8c)$$

with boundary conditions

$$\rho_{\text{opt}}^u(0, \mathbf{x}) = \rho_0, \quad \rho_{\text{opt}}^u(T, \mathbf{x}) = \rho_T \quad (9)$$

where $\psi(t, \mathbf{x})$ is a $C^2([0, T]; \mathcal{X})$ value function.

Proof: For $\mathcal{X} \subseteq \mathbb{R}^n$, let $r_0 \in \overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$ (two point compactification of the real line $\mathbb{R}$) be defined as $r_0 := \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$.

We write the Lagrangian (6) as the sum of three state-time integrals

$$\begin{aligned} \int_0^T \int_{\mathcal{X}} \frac{1}{2} \|\mathbf{u}\|_2^2 \rho^u d\mathbf{x} dt + \int_0^T \int_{\mathcal{X}} \psi \frac{\partial \rho^u}{\partial t} d\mathbf{x} dt \\ + \int_0^T \int_{\mathcal{X}} \left(\frac{\partial \rho^u}{\partial t} + \nabla_{\mathbf{x}} \cdot (\rho^u \mathbf{f}) - \langle \mathbf{G}, \text{Hess}(\rho^u) \rangle \right) \psi d\mathbf{x} dt. \end{aligned} \quad (10)$$

In above, for the second summand, we invoke the Fubini–Tonelli theorem to switch the order of integration and perform integration by parts with respect to t . This gives

$$\begin{aligned} \int_0^T \int_{\mathcal{X}} \psi \frac{\partial \rho^u}{\partial t} d\mathbf{x} dt \\ = \int_{\mathcal{X}} \left(\int_0^T \psi \frac{\partial \rho^u}{\partial t} dt \right) d\mathbf{x} \\ = \int_{\mathcal{X}} \left([\psi \rho^u]_{t=0}^{t=T} - \int_0^T \frac{\partial \psi}{\partial t} \rho^u dt \right) d\mathbf{x} \\ = \underbrace{\int_{\mathcal{X}} (\psi(T, \mathbf{x}) \rho_T(\mathbf{x}) - \psi(0, \mathbf{x}) \rho_0(\mathbf{x})) d\mathbf{x}}_{\text{constant over } \mathcal{P}_{0T}(\mathcal{X}) \times \mathcal{U}} - \int_0^T \int_{\mathcal{X}} \frac{\partial \psi}{\partial t} \rho^u d\mathbf{x} dt. \end{aligned} \quad (11)$$

For the third summand in (10), we perform integration by parts with respect to $\mathbf{x}$, to obtain

$$\begin{aligned} \int_0^T \int_{\mathcal{X}} \left(\frac{\partial \rho^u}{\partial t} + \nabla_{\mathbf{x}} \cdot (\rho^u \mathbf{f}) - \langle \text{Hess}, G\rho^u \rangle \right) \psi d\mathbf{x} dt \\ = \int_0^T \left(\int_{\mathcal{X}} \left(\frac{\partial \rho^u}{\partial t} + \nabla_{\mathbf{x}} \cdot (\rho^u \mathbf{f}) \right) \psi d\mathbf{x} - \langle \text{Hess}, G\rho^u \rangle \psi d\mathbf{x} \right) dt \\ = \int_0^T \left(\lim_{\|\mathbf{x}\|_2 \rightarrow r_0} \left[\psi(t, \mathbf{x}) \int \frac{\partial \rho^u}{\partial t} d\mathbf{x} \right] \right) dt - \int_0^T \int_{\mathcal{X}} \langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle \rho^u d\mathbf{x} dt \\ - \int_0^T \int_{\mathcal{X}} \langle \text{Hess}, G\rho^u \rangle \psi d\mathbf{x} dt \\ = - \int_0^T \int_{\mathcal{X}} \langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle \rho^u d\mathbf{x} dt - \int_0^T \int_{\mathcal{X}} \langle \text{Hess}, G\rho^u \rangle \psi d\mathbf{x} dt \end{aligned} \quad (12)$$

where we assumed that the limits at $\|\mathbf{x}\|_2 \rightarrow r_0$ are zero.

Now consider the second summand in (12), and perform twofold integration by parts with respect to $\mathbf{x}$ as

$$\begin{aligned}
\int_{\mathcal{X}} \langle \mathbf{Hess}, \mathbf{G} \rho^u \rangle \psi \, d\mathbf{x} &= \int_{\mathcal{X}} \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} (G_{ij} \rho^u) \psi \, d\mathbf{x} \\
&= \sum_{i,j} \int_{\mathcal{X}} \frac{\partial^2}{\partial x_i \partial x_j} (G_{ij} \rho^u) \psi \, d\mathbf{x} \\
&= - \sum_{i,j} \int_{\mathcal{X}} \frac{\partial (G_{ij} \rho^u)}{\partial x_j} \frac{\partial \psi}{\partial x_i} \, d\mathbf{x} \\
&= \sum_{i,j} \int_{\mathcal{X}} (G_{ij} \rho^u) \frac{\partial^2 \psi}{\partial x_j \partial x_i} \, d\mathbf{x} \\
&= \int_{\mathcal{X}} \sum_{i,j} (G_{ij} \rho^u) \frac{\partial^2 \psi}{\partial x_j \partial x_i} \, d\mathbf{x} \\
&= \int_{\mathcal{X}} \langle \mathbf{G}, \mathbf{Hess}(\psi) \rangle \rho^u \, d\mathbf{x} \quad (13)
\end{aligned}$$

which helps rewrite (12) as

$$- \int_0^T \int_{\mathcal{X}} \langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle \rho^u \, d\mathbf{x} \, dt - \int_0^T \int_{\mathcal{X}} \langle \mathbf{G}, \mathbf{Hess}(\psi) \rangle \rho^u \, d\mathbf{x} \, dt. \quad (14)$$

Combining (11), (12), (14), and dropping the constant term, the Lagrangian (10) simplifies to

$$\int_0^T \int_{\mathcal{X}} \left(\frac{1}{2} \|\mathbf{u}\|_2^2 - \frac{\partial \psi}{\partial t} - \langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle - \langle \mathbf{G}, \mathbf{Hess}(\psi) \rangle \right) \rho^u \, d\mathbf{x} \, dt. \quad (15)$$

Minimizing (15) with respect to $\mathbf{u}$ for a fixed pdf ρ^u yields (8c).

We then substitute (8c) back in (15), and equate the resulting expression to zero, to arrive at the dynamic programming equation

$$\begin{aligned}
&\int_0^T \int_{\mathcal{X}} \left(\frac{1}{2} \|\nabla_{\mathbf{u}_{\text{opt}}} (\langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle + \langle \mathbf{G}, \mathbf{Hess}(\psi) \rangle) \|_2^2 \right. \\
&\quad \left. - \frac{\partial \psi}{\partial t} - \langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle - \langle \mathbf{G}, \mathbf{Hess}(\psi) \rangle \right) \rho^u(t, \mathbf{x}) \, d\mathbf{x} \, dt = 0. \quad (16)
\end{aligned}$$

Since (16) should be satisfied for arbitrary ρ^u , we get

$$\begin{aligned}
\frac{\partial \psi}{\partial t} &= \frac{1}{2} \|\nabla_{\mathbf{u}} (\langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle + \langle \mathbf{G}, \mathbf{Hess}(\psi) \rangle) \|_2^2 - \langle \nabla_{\mathbf{x}} \psi, \mathbf{f} \rangle \\
&\quad - \langle \mathbf{G}, \mathbf{Hess}(\psi) \rangle
\end{aligned}$$

which is the HJB PDE (8a). The FPK PDE (8b) and the boundary conditions (9) follow from the primal feasibility conditions (5b) and (5c), respectively. ■

Remark 2: The conditions of optimality (8) relate the primal variables ($\rho_{\text{opt}}^u(t, \mathbf{x})$, $\mathbf{u}_{\text{opt}}(t, \mathbf{x})$) with the dual variable (i.e., Lagrange multiplier) $\psi(t, \mathbf{x})$. Specifically, the HJB PDE (8a) and the controlled FPK PDE (8b) express the dual and the primal feasibility, respectively. The optimal control policy equation (8c) expresses the primal-dual relation.

Structurally, the system of coupled PDEs (8) for our control nonaffine GSBP is quite different from the corresponding system for control-affine SBPs [9, eqs. (20)–(21)], [10, eq. (4)],

[11, eqs. (5.7)–(5.8)]. In the control-affine SBPs, the conditions of optimality involve two coupled PDEs: one being the HJB PDE and another being the controlled FPK PDE, as in (8a) and (8b). Once this pair of PDEs are solved for two unknowns ρ_{opt}^u , ψ using techniques such as Hopf-Cole transform followed by a contractive fixed point recursion [9], [33] or Feynman-Kac path integrals [21], the optimal control $\mathbf{u}_{\text{opt}}$ is obtained as a scaled (sub)gradient of ψ . In other words, for control-affine SBPs, $\mathbf{u}_{\text{opt}}$ is an explicit functional of ψ .

In contrast, the system (8) for our nonaffine GSBP comprises of $m+2$ coupled PDEs in three unknowns: ρ_{opt} , $\mathbf{u}_{\text{opt}}$, ψ , where m is the number of control inputs. This is because (8c) itself gives m PDEs coupled in ψ and $\mathbf{u}_{\text{opt}}$, while the equation pair (8a)–(8b) are coupled in ρ_{opt} , $\mathbf{u}_{\text{opt}}$, ψ . Existing techniques such as Hopf-Cole transform or Feynman-Kac path integrals no longer apply for this situation, and new ideas are needed to numerically solve the coupled system (8)–(9).

Table I summarizes how known results in the literature can be recovered as special cases of (8).

III. SOLVING THE CONDITIONS FOR OPTIMALITY USING PINN WITH SINKHORN LOSSES

In this Section, we propose a new variant of the PINN [34], [35] designed for numerically solving (8)–(9). To do so, we first introduce the 2-Wasserstein distance followed by its entropic regularization.

Definition 1 (2-Wasserstein Distance): The squared 2-Wasserstein distance W between a pair of probability measures μ_1, μ_2 supported, respectively, on $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^n$, is

$$W^2(\mu_1, \mu_2) := \inf_{\mu \in \mathcal{M}(\mu_1, \mu_2)} \int_{\mathcal{X} \times \mathcal{Y}} \|\mathbf{x} - \mathbf{y}\|_2^2 \, d\mu(\mathbf{x}, \mathbf{y}) \quad (17)$$

where $\mathcal{M}(\mu_1, \mu_2)$ is the set of joint probability measures or couplings over the product space $\mathcal{X} \times \mathcal{Y}$ having $\mathbf{x}$ marginal μ_1 , and $\mathbf{y}$ marginal μ_2 . Hereafter, we refer to (17) as the squared Wasserstein distance, dropping the prefix 2.

For metric properties of W , see e.g., [47, Ch. 7]. Whenever μ_1 and μ_2 are absolutely continuous, their respective PDFs ρ_1, ρ_2 exist, i.e., $d\mu_1(\mathbf{x}) = \rho_1(\mathbf{x})d\mathbf{x}$ and $d\mu_2(\mathbf{y}) = \rho_2(\mathbf{y})d\mathbf{y}$, and we use the equivalent notation $W^2(\rho_1, \rho_2)$. Note that (17) corresponds to a linear program (LP) and is in fact, the Kantorovich formulation [48] of OMT.

Definition 2 (Sinkhorn Loss): The Sinkhorn loss between a pair of probability measures μ_1, μ_2 supported, respectively, on $\mathcal{X}, \mathcal{Y} \subseteq \mathbb{R}^n$, with fixed $\varepsilon > 0$, is the entropy-regularized squared Wasserstein distance, i.e.,

$$W_\varepsilon^2(\mu_1, \mu_2) := \inf_{\mu \in \mathcal{M}(\mu_1, \mu_2)} \int_{\mathcal{X} \times \mathcal{Y}} \{ \|\mathbf{x} - \mathbf{y}\|_2^2 + \varepsilon \log \mu \} \, d\mu(\mathbf{x}, \mathbf{y}) \quad (18)$$

where $\mathcal{M}(\mu_1, \mu_2)$ is the set of joint probability measures or couplings over the product space $\mathcal{X} \times \mathcal{Y}$ having $\mathbf{x}$ marginal μ_1 , and $\mathbf{y}$ marginal μ_2 .

It is known [49] that $W_\varepsilon^2 \rightarrow W^2$ in the limit $\varepsilon \downarrow 0$. Even though the Sinkhorn loss (18) does not define a metric over $\mathcal{M}$, its computation offers several advantages over that of (17). For instance, the entropic regularization makes the objective in (18) strictly convex, and its discrete implementation was

TABLE I
SPECIAL CASES OF THE GSBP (1) [EQUIVALENTLY (5A)] AND CORRESPONDING REDUCTIONS OF THE OPTIMALITY CONDITIONS (8)

Special case	Drift $\mathbf{f}$	Diffusion $\mathbf{g}$	Form of (8a)	Form of (8b)	Form of (8c)
OMT [7]	$\mathbf{u}$	$\mathbf{0}_{n \times p}$	$\frac{\partial \psi}{\partial t} + \frac{1}{2} \ \nabla_{\mathbf{x}} \psi\ _2^2 = 0$	$\frac{\partial \rho_{\text{opt}}^{\mathbf{u}}}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\rho_{\text{opt}}^{\mathbf{u}} \nabla_{\mathbf{x}} \psi)$	$\mathbf{u}_{\text{opt}} = \nabla_{\mathbf{x}} \psi$
SBP [5], [6]	$\mathbf{u}$	$\mathbf{I}_{n \times n}$	$\frac{\partial \psi}{\partial t} + \frac{1}{2} \ \nabla_{\mathbf{x}} \psi\ _2^2 + \Delta_{\mathbf{x}} \psi = 0$	$\frac{\partial \rho_{\text{opt}}^{\mathbf{u}}}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\rho_{\text{opt}}^{\mathbf{u}} \nabla_{\mathbf{x}} \psi) + \Delta_{\mathbf{x}} \psi$	$\mathbf{u}_{\text{opt}} = \nabla_{\mathbf{x}} \psi$
Control-affine GSBP [9]	$\tilde{\mathbf{f}}(t, \mathbf{x})$ $+\mathbf{B}(t)\mathbf{u}$	$\mathbf{B}(t) \in \mathbb{R}^{n \times m}$	$\frac{\partial \psi}{\partial t} + \frac{1}{2} \ \mathbf{B}(t)^\top \nabla_{\mathbf{x}} \psi\ _2^2 + \langle \nabla_{\mathbf{x}} \psi, \tilde{\mathbf{f}} \rangle$ $+\langle \mathbf{B}(t)\mathbf{B}(t)^\top, \mathbf{Hess}(\psi) \rangle = 0$	$\frac{\partial \rho_{\text{opt}}^{\mathbf{u}}}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\rho_{\text{opt}}^{\mathbf{u}} (\tilde{\mathbf{f}} + \mathbf{B}(t)\mathbf{B}(t)^\top \nabla_{\mathbf{x}} \psi))$ $+\langle \mathbf{B}(t)\mathbf{B}(t)^\top, \mathbf{Hess}(\rho_{\text{opt}}^{\mathbf{u}}) \rangle$	$\mathbf{u}_{\text{opt}} = \mathbf{B}(t)^\top \nabla_{\mathbf{x}} \psi$

proposed [50] as a fast numerical approximant of the OMT (17). As we explain next, (18) is also better suited for automatic differentiation with respect to NN parameters for PINN training, which is what we need for our boundary conditions (1c) [or equivalently (9)].

A. Learning With Sinkhorn Losses

To better understand the advantage of learning with Sinkhorn losses, consider the squared Euclidean distance matrix $\mathbf{C} \in \mathbb{R}^{d \times d}$, and for a given pair of d -dimensional probability vectors $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2$, let $\Pi(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2)$ denote the set of all coupling matrices, i.e.,

$$\Pi(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2) := \{\mathbf{M} \in \mathbb{R}^{d \times d} \mid \mathbf{M} \geq 0 \text{ (element-wise)} \\ \mathbf{M}\mathbf{1} = \boldsymbol{\mu}_1, \mathbf{M}^\top \mathbf{1} = \boldsymbol{\mu}_2\}. \quad (19)$$

The dimension d here represents the number of samples involved, i.e., the dimensionality of the standard simplex in which $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2$ belong to. Then the discrete version of (18) becomes

$$W_\varepsilon^2(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2) = \min_{\mathbf{M} \in \Pi(\boldsymbol{\mu}_1, \boldsymbol{\mu}_2)} \langle \mathbf{C} + \varepsilon \log \mathbf{M}, \mathbf{M} \rangle \quad (20)$$

where $\varepsilon > 0$ is a fixed regularization parameter. The convex problem (20) can be solved using the Sinkhorn recursions [51], [52] a.k.a. iterative proportional fitting procedure (IPFP) [53]. These recursions are motivated by the observation that the minimizer of (20) must be a diagonal scaling of the known matrix $\boldsymbol{\Gamma} := \exp((-C/2\varepsilon)) \in \mathbb{R}_{>0}^{d \times d}$ where the exponential is element-wise, i.e.,

$$\mathbf{M} = \text{diag}(\mathbf{v}_1) \boldsymbol{\Gamma} \text{diag}(\mathbf{v}_2) \quad (21)$$

for to-be-determined $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^d$.

Starting with some initial guess, the Sinkhorn recursions alternate between updating $\mathbf{v}_1$ and $\mathbf{v}_2$ until convergence

$$\mathbf{v}_1^{k+1} \leftarrow \boldsymbol{\mu}_1 \oslash (\boldsymbol{\Gamma} \mathbf{v}_2^k) \quad (22a)$$

$$\mathbf{v}_2^{k+1} \leftarrow \boldsymbol{\mu}_2 \oslash (\boldsymbol{\Gamma}^\top \mathbf{v}_1^{k+1}) \quad (22b)$$

for recursion index $k = 0, 1, \dots$; the symbol $\oslash$ denotes the element-wise (Hadamard) division. The updates (22a)–(22b) can be seen as alternating Kullback-Leibler projections [54], [55] with guaranteed linear convergence.

When $\varepsilon = 0$ in (20), we get an LP corresponding to the discrete version of (17). This LP has d^2 unknowns with

$d^2 + 2d$ constraints, and solving the same as standard network flow problem has $\tilde{O}(d^2(2d)^{1/2})$ complexity [56] which is impractical for large d . Furthermore, using (17) as the endpoint loss for training a PINN to learn the solution of (8) and (9), requires us to compute

$$\text{AutoDiff}_\theta W^2(\mu_i, \mu_i^{\text{epoch index}}(\boldsymbol{\theta})) \quad \forall i \in \{0, T\} \quad (23)$$

for each epoch of the training, where AutoDiff_θ refers to the standard reverse mode automatic differentiation with respect to PINN training parameter $\boldsymbol{\theta}$. Evaluating (23) then amounts to differentiating through a very large scale LP which is computationally challenging even for moderately large d .

In contrast, using (18) as the endpoint loss for training a PINN to learn the solution of (8)–(9), requires us to compute

$$\text{AutoDiff}_\theta W_\varepsilon^2(\mu_i, \mu_i^{\text{epoch index}}(\boldsymbol{\theta})) \quad \forall i \in \{0, T\} \quad (24)$$

for a fixed $\varepsilon > 0$. Because the Sinkhorn recursions (22) involve a series of differentiable linear operations, it is amenable to Pytorch auto-differentiation to support backpropagation. Thus using W_ε^2 instead of W^2 as the endpoint distributional losses incur lesser computational overhead allowing us to train PINNs for nontrivial GSBPs. This advantage of Sinkhorn losses over Wasserstein losses has also been pointed out in a different context in [57]. Rigorous consistency results have appeared in [58] showing that the derivatives of the iterates from Sinkhorn recursion computed through automatic differentiation, indeed converge to the derivatives of the corresponding Sinkhorn loss.

B. Proposed PINN Architecture

Our proposed architecture for the PINN is shown in Fig. 2. For the GSBP considered here, the state-time $\boldsymbol{\xi} := (\mathbf{x}, t)$ comprises the features that are inputs to the network, and the network output $\boldsymbol{\eta} := (\psi, \rho_{\text{opt}}^{\mathbf{u}}, \mathbf{u}_{\text{opt}})$ comprises of the value function, optimally controlled pdf, and optimal policy.

The proposed PINN is a fully connected feed-forward NN with multiple hidden layers, and we parameterize its output using the network parameter $\boldsymbol{\theta} \in \mathbb{R}^D$, i.e.,

$$\boldsymbol{\eta}(\boldsymbol{\xi}) \approx \mathcal{N}_{\text{Schrödinger Bridge}}(\boldsymbol{\xi}; \boldsymbol{\theta}) \quad (25)$$

where $\mathcal{N}_{\text{Schrödinger Bridge}}(\cdot; \boldsymbol{\theta})$ denotes the NN approximant parameterized by $\boldsymbol{\theta}$. Here D denotes the dimension of the parameter space, i.e., the total number of to-be-trained weight,




Fig. 2. Architecture of the PINN with the system state x , and the time t as the input features $\xi := (x, t)$. The network output η comprises of the value function, optimally controlled pdf, and optimal control policy, i.e., $\eta := (\psi, \rho_{\text{opt}}^u, u_{\text{opt}})$. The networks $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$ are fully trained from MD simulation.

bias, and scaling parameters for the NN. For all neurons, we use the tanh activation functions.

As mentioned in Section II, the explicit expressions for f and g , the drift and diffusion coefficients, are not available from first-principle physics. We learn these coefficients from MD simulation data (see Sections IV-A and IV-B). As shown in Fig. 2, the networks $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$, represent the learned drift f and the learned diffusion g , respectively, which are used to evaluate the loss function $\mathcal{L}_{\mathcal{N}_{\text{Schrödinger Bridge}}}$ for the PINN.

The PINN loss function $\mathcal{L}_{\mathcal{N}_{\text{Schrödinger Bridge}}}$ consists of the sum of the losses associated with the $m + 2$ equations in (8), and the losses associated with the boundary conditions (9). Specifically, let $\mathcal{L}_\psi$ be the MSE loss for the HJB PDE (8a). Likewise, let $\mathcal{L}_{\rho_{\text{opt}}^u}$ be the MSE loss for the FPK PDE (8b), and because the control policy has m components ($u_1, \dots, u_m$), let $\mathcal{L}_{u_{\text{opt}}}|_{j=1, \dots, m}$ be the corresponding MSE loss term for each control policy component in (8c).

However, the MSE losses are insufficient to capture the distributional mismatch for endpoint boundary conditions (9). Per Section III-A, we use the Sinkhorn losses as the boundary condition losses $\mathcal{L}_{\rho_0}$ and $\mathcal{L}_{\rho_T}$, and differentiate through the corresponding Sinkhorn recursions.

Thus,

$$\mathcal{L}_{\mathcal{N}_{\text{Schrödinger Bridge}}} := \mathcal{L}_{\rho_0} + \mathcal{L}_{\rho_T} + \mathcal{L}_\psi + \mathcal{L}_{\rho_{\text{opt}}^u} + \sum_{j=1}^m \mathcal{L}_{u_{\text{opt}}|_{j=1, \dots, m}} \quad (26)$$

where each summand loss term in (26) is evaluated on a set of N collocation points $\{\xi_i\}_{i=1}^N$ in the domain of the feature space $\Omega := \mathcal{X} \times [0, T]$, i.e., $\{\xi_i\}_{i=1}^N \subset \Omega$. For instance, the equation error losses are of the form

$$\mathcal{L}_\psi := \frac{1}{N} \sum_{i=1}^N \left(\left. \frac{\partial \psi}{\partial t} \right|_{\xi_i} - \frac{1}{2} \|u_{\text{opt}}\|_2^2 \right)_{\xi_i}^2 + \langle \nabla \psi, f \rangle_{\xi_i} + \langle G, \text{Hess}(\psi) \rangle_{\xi_i}^2$$

$$\mathcal{L}_{\rho_{\text{opt}}^u} := \frac{1}{N} \sum_{i=1}^N \left(\left. \frac{\partial \rho_{\text{opt}}^u}{\partial t} \right|_{\xi_i} + \nabla \cdot (\rho_{\text{opt}}^u f) \right)_{\xi_i}^2 - \langle \text{Hess}, G \rho_{\text{opt}}^u \rangle_{\xi_i}^2$$

$$\mathcal{L}_{u_{\text{opt}}|_{j=1, \dots, m}} := \frac{1}{N} \sum_{i=1}^N \left(\left. u_{j_{\text{opt}}} \right|_{\xi_i} - \frac{\partial}{\partial u_{j_{\text{opt}}}} (\langle \nabla \psi, f \rangle) + \langle G, \text{Hess}(\psi) \rangle \right)_{\xi_i}^2$$

where u_{opt}^j denotes the j th component of the optimal control u_{opt} .

We implemented the Sinkhorn recursions with the log-sum-exp (LSE) technique [59, Sec. 3] to maintain numerical stability at the expense of minor memory overhead. We employed mini-batching for sampling our PINN output, and used the same sample indices to sample from our prescribed ρ_0, ρ_T . The squared Euclidean distance matrix $\mathbf{C}$ mentioned in Section III-A was constructed from the output batch points.

We used the PINN software library [35] with a Pytorch backend to perform numerical experiments using the above loss functions. The PINN library [35] was not written for Schrödinger bridge-type problems, so we needed to modify it to suit our needs. One modification was to program PINN to compute loss between outputs and distributions directly and integrate the Sinkhorn iteration algorithm into the library. We also modified it to perform the mini-batching we needed. In summary, for training the PINN, the overall loss (26) was minimized over $\theta \in \mathbb{R}^D$ by solving

$$\theta^* = \underset{\theta \in \mathbb{R}^D}{\text{argmin}} \mathcal{L}_{\mathcal{N}_{\text{Schrödinger Bridge}}}(\{\xi_i\}_{i=1}^N; \theta). \quad (27)$$

Section IV details the simulation setup and reports the numerical results.




Fig. 3. (a) Initial disordered crystalline structure. (b) Final BCC structure with minor defects. These images were generated using OVITO [61].

IV. NUMERICAL CASE STUDY OF CONTROLLED ISOTROPIC COLLOIDAL SA IN AN NPT ENSEMBLE

We now present a numerical case study of a colloidal SA system where the drift coefficient $\mathbf{f}$ and the diffusion coefficient $\mathbf{g}$ are not analytically available, instead they are learned as NN representations $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$, from MD simulation data. Such learned representations are nonlinear in the state $\mathbf{x}$ and nonaffine in control $\mathbf{u}$. We then solve the GSBP (1) using the PINN architecture proposed in Section III to design a minimum effort controller steering the distribution in the order-parameter space to synthesize the body-centered cubic (BCC) crystal structure over the given time horizon. Fig. 3 shows an initial disordered structure and a final BCC structure.

A. System Description

We consider the *in-silico* representation of isotropic colloidal particles with identical Lennard-Jones interaction potentials within an NPT (isothermal-isobaric) ensemble. The Lennard-Jones potential is used to model particle interactions in the system and is defined as

$$U(r) := 4\epsilon \left(\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right) \quad (28)$$

where r denotes the particle radius, and ϵ denotes the depth of the potential energy well and thus quantifies the strength of attractive forces between particles. The symbol σ denotes the distance at which the potential energy is nullified, thereby demarcating the intermolecular potential's shift from attraction to repulsion depending on particle size [60, p. 234].

An ensemble of 2048 particles is initialized at a given temperature and pressure. While the positions of these particles may be considered the most natural states of a colloidal SA system, they result in an unmanageably high-dimensional state space. To circumvent this difficulty, we seek a lower-dimensional representation. To this end, the Steinhardt bond order parameters $\langle C_{10} \rangle$ and $\langle C_{12} \rangle$ are used in this work, which are directly defined in terms of the particle positions. To calculate these parameters [62] from the MD simulation data, we proceed through a series of steps, as discussed below.

We first extract the positional information for each particle from the MD simulation data (see Section IV-B). Next, we identify the neighbors for each particle based on the Voronoi method [63]. Using this information, we calculate the spherical harmonics, Y_{lm} , indexed by two quantum numbers, viz. the azimuthal or orbital quantum number, denoted by l , and the magnetic quantum number, denoted by m .

The azimuthal quantum number defines the shape of the orbital, and for our context $l \in [1, 12]$. The magnetic quantum number represents the orientation of the orbital in space, and for our context $m \in [-l, l]$, see e.g., [64, p. 545]. Accordingly, the spherical harmonics are defined as

$$Y_{lm}(\theta, \phi) := \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{im\phi} \quad (29)$$

where θ and ϕ represent the polar and azimuthal angles, respectively. In (29), the P_l^m denote the associated Legendre polynomials [65, pp. 331–339], which is a class of functions that arise in the solution to Laplace's equation in spherical coordinates.

Let $v(i)$ denote the number of neighbors of particle i , and let r_{ij} signify the positional vector between particle i and its neighbor j . Subsequently, the l th bond order parameter $C_l(i)$, for each particle i , is computed as [62]

$$C_l(i) = \left(\frac{4\pi}{2l+1} \sum_{m=-l}^l \left| \frac{1}{v(i)} \sum_{j=1}^{v(i)} Y_{lm}(r_{ij}) \right|^2 \right)^{\frac{1}{2}} \quad (30)$$

where the index $i \in [0, v]$ and v represents the total number of particles in the ensemble ($v = 2048$ in our case study). Furthermore, index $j \in [0, v(i)]$. In (30), normalizing by the number of neighbors ensures that the final system order parameter is size-independent and thus, scalable across different systems. That is, the normalization ensures that $C_l(i) \in [0, 1]$.

Next, the individual bond order parameters $C_l(i)$ are averaged over all particles in the ensemble to calculate the averaged l th Steinhardt bond order parameter

$$\langle C_l \rangle = \frac{1}{v} \sum_{i=1}^v C_l(i) \quad (31)$$

which can be used to describe the state of a colloidal SA system.* Therefore, the physics-based order parameters serve as a reduced-dimensionality conduit that enhances the efficiency and effectiveness of our subsequent analyses by circumventing the need to work with high-dimensional particle position data.†

In this work, we specifically choose the order parameters $\langle C_{10} \rangle$ and $\langle C_{12} \rangle$ for their efficacy in distinguishing between the BCC and the face-centered cubic (FCC) structures. The values of $\langle C_{10} \rangle$ and $\langle C_{12} \rangle$ for defect-free assembled BCC and FCC structures do not overlap, which enables differentiation between the two structure types.

In summary, the controlled dynamics of our colloidal SA system is described by the SDE (1b), where the state and inputs are defined as

$$\begin{aligned} \mathbf{x} &:= (\langle C_{10} \rangle, \langle C_{12} \rangle) \in \mathcal{X} \equiv [0, 1]^2 \\ \mathbf{u} &:= (\text{temperature}, \text{pressure}) \in \mathcal{U}. \end{aligned}$$

We denote the components of the optimal control policy $\mathbf{u}_{\text{opt}}$ as $u_1^{\text{opt}}, u_2^{\text{opt}}$, respectively.

*In actual SA systems, image analysis techniques can be used to locate and track particle centers, as well as compute local and global order parameters in real-time (e.g., see [23]).

†In this work, the Steinhardt bond order parameters were calculated using the Python package Freud [66].

TABLE II

COMPARISON OF DIFFERENT MODEL ARCHITECTURES AND HYPERPARAMETERS FOR LEARNING THE NN REPRESENTATIONS $\mathcal{N}_{\text{DRIFT}}$ AND $\mathcal{N}_{\text{DIFFUSION}}$ FOR THE DRIFT $\mathbf{f}$ AND DIFFUSION $\mathbf{g}$, RESPECTIVELY. THE DIFFERENT ARCHITECTURES VARY IN THE NUMBER OF HIDDEN LAYERS AND THEIR NODES, ALL USING $\tanh$ ACTIVATION FUNCTION. ARCHITECTURE 1 EMPLOYS ONE HIDDEN LAYER WITH 200 NODES, ARCHITECTURE 2 UTILIZES A HIDDEN LAYER OF 1000 NODES, AND ARCHITECTURE 3 DEPLOYS SIX HIDDEN LAYERS WITH 200 NODES EACH

Model number	Learning rate	Batch size (of epoch)	Model architecture	Validation loss	Training loss
1	10^{-3}	1/4	1	0.390	0.057
2	10^{-2}	1/4	1	0.460	0.024
3	10^{-4}	1/4	1	1.260	0.120
4	10^{-3}	1	1	3.890	0.200
5	10^{-3}	1/4	1	0.200	0.016
6	10^{-3}	1/4	2	9.290	0.072
7	10^{-3}	1/4	3	0.030	0.003
8	10^{-4}	1/4	3	0.031	0.007
9	10^{-3}	1	3	0.110	0.015

B. Learning $\mathbf{f}$ and $\mathbf{g}$

To learn the NN representations $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$, which model the drift and diffusion coefficients $\mathbf{f}$ and $\mathbf{g}$ in the SDE (1b), we performed MD simulations for the above-described system using the Python package HOOMD-blue [67] with final time $T = 200$ s. The simulation data consisted of 200 state trajectories, i.e., the trajectories of the order parameters $\langle C_{10} \rangle$, $\langle C_{12} \rangle$ for $t \in [0, T]$, that represent the time evolution of position of the $\nu = 2048$ particles of the colloidal SA system, mentioned earlier in Section IV-A. Each state trajectory was generated under different linear temperature and pressure ramp rates (i.e., $\mathbf{u}$), which were sampled using a Latin Hypercube design and scaled to $[-0.005, 0.005]$, the input range for the simulation. To generate the training and test data for learning the NN models, the state trajectories were sampled 500 times.

Building on our earlier work [29], $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$ were trained on the MD data with a controlled neural SDE. The NNs are designed to be functions of the current time $t \in [0, T]$, the system state $\mathbf{x}$, and the control input $\mathbf{u}$. The networks are passed to (1b) to predict the state evolution. The MSE loss is computed for each time step, and the learning process aims to minimize the total MSE loss between the networks' predicted states and the actual observed states from the MD simulation trajectories. For model optimization, we used the Adam optimizer [68] which adjusts the learning rate on a per-parameter basis. The learning rate was initially set to a predefined constant, as per Table II, and was subsequently adjusted using an exponential learning rate scheduler with a decay rate of 0.999. This scheduler reduces the learning rate multiplicatively after each epoch. The data was partitioned into a 70/20/10 distribution for the training, testing, and validation subsets, respectively. The implementation of these networks was done with the torchsde [69] Python package.

To determine the best architecture for the NNs $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$, we used hyperparameter (depth and width, batch size, learning rate) turning as detailed in Table II; a total of nine models were trained and evaluated. All of the NN




Fig. 4. Validation losses for nine different NN models $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$, for the drift and diffusion terms in the SDE (1b), with legend numbers corresponding to model numbers in Table II.

architectures follow a sequential design of fully connected layers, with five input units and an output layer of two units. The architectures vary in the number of hidden layers and their nodes, all using $\tanh$ activation functions. Architecture 1 employs one hidden layer with 200 nodes; architecture 2 utilizes a hidden layer of 1000 nodes; and architecture 3 deploys six hidden layers with 200 nodes each. The batch size, defining the number of samples to be processed before updating the model, is tuned for learning. Lastly, we adjust the learning rate, a factor determining how much the model's parameters should be adjusted with respect to the calculated error, for balanced and steady learning without risking instability or slow convergence. The MSE was used as the loss function for all models.

Fig. 4 shows the validation MSE loss for all models, which are evaluated by using the $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$ in (1b) to predict the state $\mathbf{x}$, and then comparing the predicted states with those obtained from MD simulations. These validation results demonstrate that all models converge, indicating that the training time was sufficient. As seen in Table II and Fig. 4, model 7 exhibited the best performance evidenced by its minimal validation loss. Consequently, we used model 7 for representing the colloidal SA dynamics in the form of (1b). Its corresponding $\mathcal{N}_{\text{Drift}}$ and $\mathcal{N}_{\text{Diffusion}}$ are used for the optimal control synthesis for the GSBP. On an NVIDIA GTX 1080, each model undergoes training that, on average, takes 10 s per training step. To complete 100 epochs, this process requires approximately 1.2 h per model. The approximate inference time for the model is 0.0123 s.

C. Controller Synthesis

With the $\mathbf{f}$, $\mathbf{g}$ learned as per Section IV-B for the colloidal SA system described in Section IV-A, we considered the GSBP (5a) over fixed time horizon $[0, T]$, where the final time $T = 200$ s, the initial state $\mathbf{x}(t = 0) \sim \rho_0 = \mathcal{N}(\mathbf{m}_0, \mathbf{\Sigma}_0)$, and the terminal state $\mathbf{x}(t = T) \sim \rho_T = \mathcal{N}(\mathbf{m}_T, \mathbf{\Sigma}_T)$. Here, the notation $\mathcal{N}(\mathbf{m}, \mathbf{\Sigma})$ stands for a joint Gaussian distribution with mean vector $\mathbf{m}$ and covariance matrix $\mathbf{\Sigma}$. We used

$$\mathbf{m}_0 = (0.2, 0.2)^\top, \quad \mathbf{m}_T = (0.4, 0.375)^\top, \quad \mathbf{\Sigma}_0 = \mathbf{\Sigma}_T = 0.1\mathbf{I}_2. \quad (32)$$




Fig. 5. PINN residuals in solving the conditions of optimality (8)–(9) for the simulation in Section IV-C.




Fig. 6. Results for the GSBP simulation detailed in Section IV-C over time $t \in [0, 200]$. The color denotes the value of the plotted variable; see colorbar (dark red = high, light yellow = low). (a) Contour plots of the optimally controlled state PDFs $\rho_{\text{opt}}^u(t, \mathbf{x})$ over the state space $[0, 1]^2$. (b) Contour plots of the optimal control component $u_{1,\text{opt}}(t, \mathbf{x})$ over the state space $[0, 1]^2$. (c) Contour plots of the optimal control component $u_{2,\text{opt}}(t, \mathbf{x})$ over the state space $[0, 1]^2$. (d) Contour plots of the value function $\psi(t, \mathbf{x})$ over the state space $[0, 1]^2$.

In particular, the statistics of the initial state $\mathbf{x}(t = 0) \sim \mathcal{N}(\mathbf{m}_0, \Sigma_0)$ is chosen to coincide with that used in the MD simulation in Section IV-B. The mean $\mathbf{m}_T$ for the target

terminal state $\mathbf{x}(t = T) \sim \mathcal{N}(\mathbf{m}_T, \Sigma_T)$ was chosen to represent the BCC crystal structure. Hence, the control objective for the GSBP represents the problem of designing a minimum




Fig. 7. 150 random sample paths resulting from closed-loop simulations using the learned optimal policy $\mathbf{u}_{\text{opt}}(t, \langle C_{10} \rangle, \langle C_{12} \rangle)$. (a) Optimally controlled $\langle C_{10} \rangle$ state trajectories. (b) Optimally controlled $\langle C_{12} \rangle$ state trajectories.

effort Markovian controller that provably steers the stochastic order parameters in a way to synthesize BCC crystal structure over the prescribed time horizon.

We used the PINN $\mathcal{N}_{\text{Schrödinger Bridge}}$ proposed in Section III-B for numerically solving the GSBP conditions of optimality (8)–(9). For our PINN implementation, the domain for state-time collocation is $\Omega = [0, 1]^2 \times [0, 200]$. Our network consisted of four hidden layers, each containing 70 neurons, all with tanh activation functions. We trained the PINN for 100 000 epochs using the Adam optimizer [68] with a learning rate of 10^{-3} . All our training were performed on a computing platform with NVIDIA Quadro p1000, 640 Cuda cores, and 64 GB RAM. For the collocation, we used $N = 3000$ pseudorandom samples, drawn using Sobol distribution, between the endpoint boundary conditions at $t = 0$ and $t = 200$. We also uniformly randomly sampled 3000 samples every 20 000 epochs to satisfy compute constraints. For computing the Sinkhorn losses at the endpoint boundary conditions, we used an entropic regularization parameter of $\varepsilon = 0.1$ as in (20). For the computing platform mentioned above, training the proposed PINN on average takes 2 s per epoch, so to complete 100 000 epochs, it takes approximately 55.5 h.

Fig. 5 shows the PINN residuals in (26), and Fig. 6 shows the corresponding GSBP solutions obtained from the trained PINN. In particular, Fig. 6(a) shows the evolution of the optimally controlled transient joint PDFs $\rho_{\text{opt}}^u(t, \mathbf{x})$ interpolating the fixed ρ_0, ρ_T mentioned above. Notice that, even though the initial and terminal stochastic states are both chosen to have Gaussian statistics, the transient joints in Fig. 6(a) are non-Gaussian. This is expected since the learned $\mathbf{f}, \mathbf{g}$, as well as the optimal controller $\mathbf{u}_{\text{opt}}$ [see Fig. 6(b) and (c)], are nonlinear in state. A comparison of Fig. 6(b) and (c) with Fig. 6(c) also shows that the optimal controls are high (resp. low) in regions where the value function ψ changes rapidly, i.e., when the (sub)gradient of ψ is large (resp. small).

To further illustrate the GSBP results, we performed a closed loop sample path simulation for 150 initial state samples $\mathbf{x}(t = 0) \sim \rho_0$ (with the same ρ_0 mentioned before) using the learned optimal control policy $\mathbf{u}_{\text{opt}}(t, \langle C_{10} \rangle, \langle C_{12} \rangle)$ that provably steers the given ρ_0 from $t = 0$ to the given ρ_T (BCC crystal) at $t = T = 200$ s. The corresponding closed-loop state sample paths shown in Fig. 7 demonstrate that the optimal policy indeed steers the controlled stochastic state from around (0.2, 0.2) to around (0.4, 0.375) with high probability, as specified per problem data (32).

For the closed-loop simulations, we constructed a k -d tree [70] (with leaf size = 2) for fast querying of the PINN-trained optimal control policy $\mathbf{u}_{\text{opt}}(t, \langle C_{10} \rangle, \langle C_{12} \rangle)$. This construction takes 1.785 s. During the numerical integration of the SDE, querying the optimal control policy takes 0.227 ms. Without the k -d tree construction, this querying is $1000\times$ slower (approximately 0.22 s). With k -d tree-based querying, to simulate a closed-loop sample path as in Fig. 7 using the Euler-Maruyama scheme with 500 equispaced time steps in $[0, T]$, taking approximately 177 s. These experiments suggest that the proposed control approach is practically viable for colloidal SA.

V. CONCLUSION

The work presented here proposes “neural Schrödinger bridge”—a novel NN-based learning and control framework for solving the GSBP, which is a fixed time horizon stochastic optimal control problem with constraints on endpoint distributions and controlled SDEs. Our work is motivated by the problem of minimum effort controlled colloidal SA, where the controlled dynamics is usually not available from first principle physics and instead learned from the MD simulation data.

Our contributions go beyond connecting the colloidal SA problem with the GSBP. Since the drift and diffusions obtained from data-driven learning of the controlled neural SDEs are nonlinear in state, nonaffine in control, and have explicit

time-dependence, the standard computational techniques for solving GSBPs available in the literature no longer apply in this setting. In fact, we show that the conditions of optimality for such GSBPs are very different from those studied in the literature in that here we are led to solve a system of $m + 2$ coupled PDEs with two boundary conditions, where m is the number of control inputs. This system of PDEs we derive is new, and is of independent interest in the theory of Schrödinger bridge and related stochastic optimal control problems. To numerically solve this nonstandard system, we propose a custom variant of PINN, and demonstrate its effectiveness on a data-driven colloidal SA case study. Our results should be of broad interest to control and machine learning researchers using diffusion models for learning and control.

REFERENCES

- [1] J. A. Paulson, A. Mesbah, X. Zhu, M. C. Molaro, and R. D. Braatz, "Control of self-assembly in micro- and nano-scale systems," *J. Process Control*, vol. 27, pp. 38–49, Mar. 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0959152414002595>
- [2] X. Tang and M. A. Grover, "Control of microparticle assembly," *Annu. Rev. Control, Robot., Auto. Syst.*, vol. 5, no. 1, pp. 491–514, May 2022.
- [3] X. Tang, Y. Xue, and M. A. Grover, "Colloidal self-assembly with model predictive control," in *Proc. Amer. Control Conf.*, Jun. 2013, pp. 4228–4233.
- [4] X. Tang, B. Rupp, Y. Yang, T. D. Edwards, M. A. Grover, and M. A. Bevan, "Optimal feedback controlled assembly of perfect crystals," *ACS Nano*, vol. 10, no. 7, pp. 6791–6798, Jul. 2016.
- [5] E. Schrödinger, "Über die umkehrung der naturgesetze," *Sitzungsberichte der Preuss. Phys. Math. Klasse*, vol. 10, pp. 144–153, Jan. 1931.
- [6] E. Schrödinger, "Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique," *Annales de L'institut Henri Poincaré*, vol. 2, no. 4, 1932, pp. 269–310.
- [7] J.-D. Benamou and Y. Brenier, "A computational fluid mechanics solution to the monge-kantorovich mass transfer problem," *Numerische Math.*, vol. 84, no. 3, pp. 375–393, Jan. 2000.
- [8] Y. Chen, T. T. Georgiou, and M. Pavon, "On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint," *J. Optim. Theory Appl.*, vol. 169, no. 2, pp. 671–691, 2016.
- [9] K. F. Caluya and A. Halder, "Wasserstein proximal algorithms for the Schrödinger bridge problem: Density control with nonlinear drift," *IEEE Trans. Autom. Control*, vol. 67, no. 3, pp. 1163–1178, Mar. 2022.
- [10] K. F. Caluya and A. Halder, "Reflected Schrödinger bridge: Density control with path constraints," in *Proc. Amer. Control Conf. (ACC)*, 2021, pp. 1137–1142.
- [11] Y. Chen, T. T. Georgiou, and M. Pavon, "Stochastic control liaisons: Richard Sinkhorn meets gaspard Monge on a Schrödinger bridge," *SIAM Rev.*, vol. 63, no. 2, pp. 249–313, Jan. 2021.
- [12] V. De Bortoli, J. Thornton, J. Heng, and A. Doucet, "Diffusion Schrödinger bridge with applications to score-based generative modeling," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 17695–17709.
- [13] G. Wang, Y. Jiao, Q. Xu, Y. Wang, and C. Yang, "Deep generative learning via Schrödinger bridge," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 10794–10804.
- [14] T. Chen, G.-H. Liu, and E. A. Theodorou, "Likelihood training of Schrödinger bridge using forward-backward SDEs theory," 2021, *arXiv:2110.11291*.
- [15] F.-A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah, "Diffusion models in vision: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 9, pp. 10850–10869, Sep. 2023.
- [16] M. Pavon, G. Trigila, and E. G. Tabak, "The data-driven Schrödinger bridge," *Commun. Pure Appl. Math.*, vol. 74, no. 7, pp. 1545–1573, 2021.
- [17] F. Vargas, P. Thodoroff, A. Lamacraft, and N. Lawrence, "Solving Schrödinger bridges via maximum likelihood," *Entropy*, vol. 23, no. 9, p. 1134, 2021.
- [18] A. Stromme, "Sampling from a Schrödinger bridge," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2023, pp. 4058–4067.
- [19] K. F. Caluya and A. Halder, "Finite horizon density steering for multi-input state feedback linearizable systems," in *Proc. Amer. Control Conf. (ACC)*, Jul. 2020, pp. 3577–3582.
- [20] S. Haddad, K. F. Caluya, A. Halder, and B. Singh, "Prediction and optimal feedback steering of probability density functions for safe automated driving," *IEEE Control Syst. Lett.*, vol. 5, no. 6, pp. 2168–2173, Dec. 2021.
- [21] I. Nodoozi and A. Halder, "Schrödinger meets Kuramoto via Feynman-Kac: Minimum effort distribution steering for noisy nonuniform Kuramoto oscillators," in *Proc. IEEE 61st Conf. Decis. Control (CDC)*, Dec. 2022, pp. 2953–2960.
- [22] I. Nodoozi, J. O'Leary, A. Mesbah, and A. Halder, "A physics-informed deep learning approach for minimum effort stochastic control of colloidal self-assembly," in *Proc. Amer. Control Conf. (ACC)*, May 2023, pp. 609–615.
- [23] J. J. Juárez and M. A. Bevan, "Feedback controlled colloidal self-assembly," *Adv. Funct. Mater.*, vol. 22, no. 18, pp. 3833–3839, Sep. 2012.
- [24] Y. Gao and R. Lakerveld, "Feedback control for defect-free alignment of colloidal particles," *Lab Chip*, vol. 18, no. 14, pp. 2099–2110, 2018.
- [25] S. Whitelam and I. Tamblyn, "Learning to grow: Control of material self-assembly using evolutionary reinforcement learning," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 101, no. 5, May 2020, Art. no. 052604.
- [26] J. O'Leary, M. M. Khare, and A. Mesbah, "Novelty search for neuroevolutionary reinforcement learning of deceptive systems: An application to control of colloidal self-assembly," in *Proc. Amer. Control Conf. (ACC)*, San Diego, CA, USA, May 2023, pp. 2776–2781.
- [27] M. Dijkstra and E. Luijten, "From predictive modelling to machine learning and reverse engineering of colloidal self-assembly," *Nature Mater.*, vol. 20, no. 6, pp. 762–773, Jun. 2021.
- [28] J. O'Leary, J. A. Paulson, and A. Mesbah, "Stochastic physics-informed neural ordinary differential equations," *J. Comput. Phys.*, vol. 468, Nov. 2022, Art. no. 111466, doi: [10.1016/j.jcp.2022.111466](https://doi.org/10.1016/j.jcp.2022.111466).
- [29] R. Mao, J. O'Leary, A. Mesbah, and J. Mittal, "A deep learning framework discovers compositional order and self-assembly pathways in binary colloidal mixtures," *JACS Au*, vol. 2, no. 8, pp. 1818–1828, Aug. 2022.
- [30] A. Lizano and X. Tang, "Convolutional neural network-based colloidal self-assembly state classification," *Soft Matter*, vol. 19, no. 19, pp. 3450–3457, 2023.
- [31] E. Hopf, "The partial differential equation $u_t + uu_x = \mu_{xx}$," *Commun. Pure Appl. Math.*, vol. 3, pp. 201–230, Sep. 1950.
- [32] J. D. Cole, "On a quasi-linear parabolic equation occurring in aerodynamics," *Quart. Appl. Math.*, vol. 9, no. 3, pp. 225–236, 1951.
- [33] Y. Chen, T. Georgiou, and M. Pavon, "Entropic and displacement interpolation: A computational approach using the Hilbert metric," *SIAM J. Appl. Math.*, vol. 76, no. 6, pp. 2375–2396, Jan. 2016.
- [34] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, Feb. 2019.
- [35] L. Lu, X. Meng, Z. Mao, and G. E. Karniadakis, "DeepXDE: A deep learning library for solving differential equations," *SIAM Rev.*, vol. 63, no. 1, pp. 208–228, Jan. 2021.
- [36] T. Koshizuka and I. Sato, "Neural Lagrangian Schrödinger bridge: Diffusion modeling for population dynamics," in *Proc. 11th Int. Conf. Learn. Represent.*, 2022.
- [37] B. Kim, G. Kwon, K. Kim, and J. Chul Ye, "Unpaired image-to-image translation via neural Schrödinger bridge," 2023, *arXiv:2305.15086*.
- [38] M. Fazlyab, M. Morari, and G. J. Pappas, "Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming," *IEEE Trans. Autom. Control*, vol. 67, no. 1, pp. 1–15, Jan. 2022.
- [39] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of Lipschitz constants for deep neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 1–12.
- [40] B. Oksendal, *Stochastic Differential Equations: An Introduction With Applications*. Berlin, Germany: Springer, 2013.
- [41] A. Friedman, *Partial Differential Equations of Parabolic Type*. Mineola, NY, USA: Courier Dover, 2008.
- [42] G. A. Pavliotis, *Stochastic Processes and Applications*. Berlin, Germany: Springer, 2016.

- [43] A. Halder and R. Bhattacharya, "Dispersion analysis in hypersonic flight during planetary entry using stochastic Liouville equation," *J. Guid., Control, Dyn.*, vol. 34, no. 2, pp. 459–474, Mar. 2011.
- [44] B. Jamison, "Reciprocal processes," *Zeitschrift Für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, vol. 30, no. 1, pp. 65–86, 1974.
- [45] A. Blaquiere, "Controllability of a Fokker-Planck equation, the Schrödinger system, and a related stochastic optimal control (revised version)," *Dyn. Control*, vol. 2, no. 3, pp. 235–253, 1992.
- [46] A. Wakolbinger, "Schrödinger bridges from 1931 to 1991," in *Proc. 4th Latin Amer. Congr. Probab. Math. Statist.*, 1992, pp. 61–79.
- [47] C. Villani, *Topics in Optimal Transportation*, vol. 58, 1st ed. Providence, RI, USA: American Mathematical Society, 2003.
- [48] L. Kantorovich, "On the transfer of masses (in Russian)," in *Doklady Akademii Nauk*, vol. 37, no. 2, 1942, pp. 227–229.
- [49] G. Carlier, V. Duval, G. Peyré, and B. Schmitzer, "Convergence of entropic schemes for optimal transport and gradient flows," *SIAM J. Math. Anal.*, vol. 49, no. 2, pp. 1385–1418, Jan. 2017.
- [50] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 26, 2013, pp. 1–9.
- [51] R. Sinkhorn, "A relationship between arbitrary positive matrices and doubly stochastic matrices," *Ann. Math. Statist.*, vol. 35, no. 2, pp. 876–879, Jun. 1964.
- [52] R. Sinkhorn, "Diagonal equivalence to matrices with prescribed row and column sums," *Amer. Math. Monthly*, vol. 74, no. 4, pp. 402–405, Apr. 1967.
- [53] W. E. Deming and F. F. Stephan, "On a least squares adjustment of a sampled frequency table when the expected marginal totals are known," *Ann. Math. Statist.*, vol. 11, no. 4, pp. 427–444, Dec. 1940.
- [54] L. M. Bregman, "The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming," *USSR Comput. Math. Math. Phys.*, vol. 7, no. 3, pp. 200–217, Jan. 1967.
- [55] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré, "Iterative Bregman projections for regularized transportation problems," *SIAM J. Sci. Comput.*, vol. 37, no. 2, pp. 1111–1138, Jan. 2015.
- [56] Y. T. Lee and A. Sidford, "Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\text{vrnk})$ iterations and faster algorithms for maximum flow," in *Proc. IEEE 55th Annu. Symp. Found. Comput. Sci.*, Oct. 2014, pp. 424–433.
- [57] A. Genevay, G. Peyré, and M. Cuturi, "Learning generative models with Sinkhorn divergences," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2018, pp. 1608–1617.
- [58] E. Pauwels and S. Vaiter, "The derivatives of Sinkhorn-Knopp converge," 2022, *arXiv:2207.12717*.
- [59] T. Viehmann, "Implementation of batched Sinkhorn iterations for entropy-regularized Wasserstein loss," 2019, *arXiv:1907.01729*.
- [60] D. McQuarrie, *Statistical Mechanics* (G—Reference, Information and Interdisciplinary Subjects Series). Melville, NY, USA: Univ. Science Books, 2000.
- [61] A. Stukowski, "Visualization and analysis of atomistic simulation data with OVITO—The open visualization tool," *Model. Simul. Mater. Sci. Eng.*, vol. 18, no. 1, Jan. 2010, Art. no. 015012.
- [62] P. J. Steinhardt, D. R. Nelson, and M. Ronchetti, "Bond-orientational order in liquids and glasses," *Phys. Rev. B, Condens. Matter*, vol. 28, no. 2, pp. 784–805, Jul. 1983.
- [63] C. Rycroft, "VORO++: A three-dimensional Voronoi cell library in C++," Lawrence Berkeley National Lab. (LBNL), Berkeley, CA, USA, Tech. Rep. LBNL-1432E, 2009.
- [64] J. Kovac, "Physical chemistry: A molecular approach (McQuarrie, Donald A.; Simon, John D.)," *J. Chem. Educ.*, vol. 75, no. 5, p. 545, 1998.
- [65] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions With Formulas, Graphs, and Mathematical Tables*, vol. 55. Washington, DC, USA: US Government Printing Office, 1968.
- [66] V. Ramasubramani, B. D. Dice, E. S. Harper, M. P. Spellings, J. A. Anderson, and S. C. Glotzer, "Freud: A software suite for high throughput analysis of particle simulation data," *Comput. Phys. Commun.*, vol. 254, Sep. 2020, Art. no. 107275.
- [67] J. A. Anderson, J. Glaser, and S. C. Glotzer, "HOOMD-blue: A Python package for high-performance molecular dynamics and hard particle Monte Carlo simulations," *Comput. Mater. Sci.*, vol. 173, Feb. 2020, Art. no. 109363.
- [68] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.

- [69] X. Li, T.-K. L. Wong, R. T. Q. Chen, and D. Duvenaud, "Scalable gradients for stochastic differential equations," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2020, pp. 3870–3882.
- [70] J. L. Bentley, "Multidimensional binary search trees used for associative searching," *Commun. ACM*, vol. 18, no. 9, pp. 509–517, Sep. 1975.



Iman Nodozi (Graduate Student Member, IEEE) received the B.S. degree in electrical engineering from the Hamedan University of Technology, Hamedan, Iran, in 2013, the M.S. degree in electrical engineering from Imam Khomeini International University, Qazvin, Iran, in 2016. He is currently pursuing the Ph.D. degree in electrical and computer engineering with the University of California at Santa Cruz, Santa Cruz, CA, USA.

His primary research interests include stochastic systems, control, optimization, and machine learning.



Charlie Yan received the M.S. degree in electrical and computer engineering from UC Santa Cruz, Santa Cruz, CA, USA, in 2023.

He currently works at Ames Research Center, Mountain View, CA, USA, on the 2024 VIPER Lunar Mission. His primary research interests include stochastic control and optimization in robotics and fabrication.



Mira Khare received the B.S. degree in chemical engineering from Oregon State University, Corvallis, OR, USA, in 2021. She is currently pursuing the Ph.D. degree in chemical engineering with the University of California at Berkeley, Berkeley, CA, USA.

Her primary research interests include optimization and control of nanoscale systems using machine learning techniques.



Abhishek Halder (Senior Member, IEEE) is an Associate Professor at the Department of Aerospace Engineering, and a member of the Translational AI Center, Iowa State University, Ames, IA, USA. He is a Visiting Associate Professor at the Department of Applied Mathematics, University of California at Santa Cruz, Santa Cruz, CA, USA.

His research interests include stochastic systems, control, optimization and machine learning with application focus on large scale cyber-physical systems.



Ali Mesbah (Senior Member, IEEE) is an Associate Professor of chemical and biomolecular engineering at the University of California at Berkeley, Berkeley, CA, USA. His research interests include learning-based analysis and control of uncertain systems, with applications to materials processing and manufacturing systems.

Dr. Mesbah is a Senior Member of the AIChE.