GENERATIVE MODELING VIA *Nonlinear* SCHRÖDINGER BRIDGES AND BRANCHING DIFFUSIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose a new generative modeling framework grounded in forward-backward partial differential equations (PDEs). These equations arise naturally in mean-field game (MFG) theory, and are closely related to nonlinear Schrödinger systems. Our main contribution is to reinterpret this connection within the setting of generative modeling and to develop an efficient numerical methodology for its practical implementation. Specifically, we derive and characterize a nonlinear extension of the classical Schrödinger bridge problem through a coupled system of PDEs, yielding a model we call *SBNL (Schrödinger Bridge NonLinear)*. To compute its solution, we design a simulation algorithm based on branching diffusions and the probabilistic representation of semilinear PDEs. The resulting method is embarrassingly parallel, scalable to high dimensions, and applicable to a broad class of nonlinear interaction functionals. We validate the approach through numerical experiments, including recovery of the classical Schrödinger bridge, nonlinear extensions with convex penalties (Gross-Pitaevskii) as well as potential interaction terms, and high-dimensional generative tasks such as image synthesis on the MNIST and FFHQ human face dataset. Furthermore, the same framework embeds score-based generative modeling, diffusion models, and tractable gaussian Schrödinger bridges.

1 INTRODUCTION

The Schrödinger bridge (SB) problem (Schrödinger, 1931; Föllmer, 1988; Léonard, 2014) has recently attracted significant attention in generative modeling due to its ability to learn an optimally controlled stochastic dynamics that transports between two arbitrary endpoint distributions. By minimizing relative entropy with respect to a reference diffusion (typically Brownian motion), the SB framework reveals deep connections between optimal transport, stochastic control, and score-based generative modeling (Benamou et al., 2019; Song et al., 2021; Chen et al., 2023; De Bortoli et al., 2021).

When both terminal costs and the evolving marginal law are controlled, the problem extends beyond classical (linear) Schrödinger bridges to mean-field control and mean-field games (MFG) (Lasry & Lions, 2007; Cardaliaguet; Guéant et al., 2011; Guéant, 2012). In this regime, interactions through the marginal distribution induce a semilinear forward–backward structure.

Existing iterative schemes for SB, such as iterative proportional fitting (IPF) / iterative Markovian fitting (IMF) and related diffusion-based approaches (Chen et al., 2023; De Bortoli et al., 2021; Shi et al., 2023), are primarily designed for *linear* entropic transport. These methods rely on alternating optimization procedures, which can be computationally expensive and may suffer from error accumulation. More recent approaches based on Gaussian mixture energy models (Korotin et al., 2023; Gushchin et al., 2024) improve scalability in some settings but remain limited in expressivity in high dimensions. Even recent extensions based on controlled stochastic dynamics (Alouadi et al., 2026) are not well suited to law-dependent running costs, which induce intrinsically coupled semilinear dynamics.

Although recent works in MFG and generalized Schrödinger bridge formulations clarify the structure of the coupling in terms of forward–backward potentials (Liu et al., 2022; Tang et al., 2026), they often do not yield objectives that are both theoretically grounded and directly amenable to unbiased stochastic optimization in high dimensions.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Our contributions are as follows:

1. We introduce a variational objective for semilinear mean-field Schrödinger bridges with law-dependent running and terminal costs, and characterize optimality through coupled $(\varphi, \bar{\varphi})$ PDEs (equivalently, a Lasry - Lions (m, v) system). We refer to the resulting model as *SBNL (Schrödinger Bridge NonLinear)*.
2. We develop a marked-branching stochastic algorithm for the two-type $(\varphi, \bar{\varphi})$ coupling, prove unbiasedness and L^2-integrability of the estimator, and train the potentials jointly (rather than via alternating refitting), with branching used both during training and sampling.
3. We report experiments ranging from bridge recovery and soft terminal penalties to Gross-Pitaevskii-type couplings and high-dimensional latent image transport.

1.1 Related work

Linear Schrödinger bridges. The Schrödinger bridge problem is classically linked to entropic optimal transport (Léonard, 2014; Villani, 2003; Benamou et al., 2019), with efficient Sinkhorn-type and iterative fitting procedures in the *linear* setting. These methods scale well when dynamics are fixed between prescribed marginals. Recent neural approaches approximate bridges through forward-backward SDEs (Chen et al., 2023) or diffusion-style training (De Bortoli et al., 2021; Shi et al., 2023); they mainly target linear bridges and often rely on expensive alternating refits (IPF/IMF). A complementary, simulation-free line parameterizes the Schrödinger potentials with Gaussian mixture forms, yielding lightweight solvers (Korotin et al., 2023; Gushchin et al., 2024) at the cost of limited expressivity. Even recent drift-and-volatility extensions (Schrödinger-Bass bridges) (Alouadi et al., 2026) still assume (approximately) linear entropic transport and are not tailored to running costs that depend on the evolving marginal law.

Control formulations and mean-field games (MFG). Once running or terminal terms depend on the marginal law, the problem leaves the linear SB regime and sits in mean-field control / MFG (Lasry & Lions, 2007; Cardaliaguet), with forward-backward PDE couplings studied in (Guéant et al., 2011; Swiecicki et al., 2016; Bonnemain et al., 2020; Benamou et al., 2019). Extending Sinkhorn-type or alternating schemes to such coupled HJB-Fokker-Planck systems is challenging in high dimension. Recent interaction/unbalanced frameworks (Liu et al., 2024; Zhang et al., 2025; Tang et al., 2026) encode growth, interaction, or state costs through composite weighted losses and staged training. Trajectory-level optimal-control views, e.g. adjoint matching (Domingo-Enrich et al., 2025), couple learning to long-horizon simulation and can be costly. We retain entropic regularization while centering *law-dependent* running costs, starting from the *non linear* Sinkhorn system with jointly learned coupled Schrödinger potentials, using branching Monte Carlo.

Branching representations for semilinear PDEs. Branching diffusion methods provide probabilistic representations of semilinear PDEs via weighted particle systems (McKean, 1975; Henry-Labordère et al., 2019; Henry-Labordère & Touzi, 2021). We extend these ideas to the coupled forward-backward system arising in mean-field Schrödinger bridges, where the interaction between the two potentials introduces a time-reflected dependency. This enables a scalable estimator for both the potentials and their gradients, which we integrate into a neural generative framework.

2 Semilinear Schrödinger bridges

We denote by μ_0 and μ_T laws on $\mathbb{R}^d$. Let $(X_t)_{t\in[0,T]}$ be controlled by a drift $\alpha_\cdot$ under $\mathbb{P}$, with $X_0 \sim \mu_0$ and $\mathrm{d}X_t = \alpha_t\,\mathrm{d}t + \mathrm{d}W_t$.

2.1 Formulation

As a generalization of the classical Schrödinger bridge formulation, we consider the following variational problem:

$$P := \inf_{\alpha_\cdot}\ \mathbb{E}_{X_0\sim\mu_0}\left[\int_0^T \left(\frac{1}{2c}|\alpha_t|^2 + f(m(t,\cdot))\right)\mathrm{d}t\ +\ g(m(T,\cdot))\right], \tag{1}$$

where $c > 0$, $m(t)$ denotes the law of X_t, and f and g are functionals. The infimum is taken over the control α, where m is constrained to satisfy the Fokker-Planck PDE.

The three terms appearing in P play different roles:

- The quadratic control cost $\frac{1}{2}|\alpha_t|^2$ penalizes deviations of the dynamics from the reference Brownian motion. By Girsanov's theorem, this term corresponds exactly to the relative entropy $H(\mathbb{P} \mid \mathbb{P}^0)$ between the controlled law $\mathbb{P}$ of the diffusion and the Wiener measure $\mathbb{P}^0$ (Föllmer, 1988). This quadratic structure is also crucial analytically: it is the *only* choice of control penalty for which the associated Hamilton-Jacobi-Bellman (HJB) equation admits a tractable solution through a Cole-Hopf transformation. For this reason, essentially all generative modeling frameworks based on stochastic control can be viewed as particular cases of the Schrödinger bridge.
- The functional $f(m(t,\cdot))$ introduces a *running cost* depending directly on the marginal law of the process at each time t. Such terms allow for mean-field type interactions, density regularization, or penalization of undesirable behaviors (e.g. concentration, mass spread) during the evolution. Classical examples include (see the review (Cardaliaguet)):
  - Quadratic interaction: $f(m) = \frac{\kappa}{2}\int_{\mathbb{R}^d} m(x)^2\,\mathrm{d}x$, which penalizes high concentrations and favors dispersion.
  - Entropy-type term: $f(m) = \int_{\mathbb{R}^d} m(x)\,\log m(x)\,\mathrm{d}x$, encouraging spreading of the distribution.
  - Potential interaction: $f(m) = \int_{\mathbb{R}^d} V(x)\,m(x)\,\mathrm{d}x$, where V is an external field acting on the density (as in (Liu et al., 2024)).
  - Pairwise interaction kernel: $f(m) = \frac{1}{2}\iint_{\mathbb{R}^d\times\mathbb{R}^d} K(x-y)\,m(x)\,m(y)\,\mathrm{d}x\mathrm{d}y$, modeling attractive or repulsive mean-field effects.
- The terminal cost $g(m(T,\cdot))$ enforces constraints or preferences on the final law. In the hard-constraint setting of the classical Schrödinger bridge, g is chosen so that $m(T,\cdot) = \mu_T$, i.e. the terminal distribution matches the prescribed target. In the relaxed setting, g can encode convex penalties measuring the discrepancy between the final distribution and the target. For use below, we denote by g^* the Legendre transform (convex conjugate) of g:

$$g^*(v(T,\cdot)) := \sup_{m(T,\cdot)} \Big\{\mathbb{E}_{X\sim m(T,\cdot)}[v(T,X)] - g(m(T,\cdot))\Big\}.$$

  Note that in practice it does not make sense to perfectly calibrate to the empirical measure, which is only a finite sample approximation of the true law. Instead, one seeks both proximity to the empirical distribution and sufficient regularity of the solution, precisely what the soft constraint achieves.

Remark 2.1 (State dependence and MFG link). Running and terminal costs that depend explicitly on the state x can be rewritten in the law-only form (1) by integrating against m; the resulting forward-backward MFG optimality system and its relation to Theorem 2.1 are recorded in Appendix E.

2.2 Main result

Theorem 2.1 (Semilinear Schrödinger bridge)**.** *We assume* (A1)-(A3) *from Appendix D. The optimal controlled process of (1) is given under a measure* $\mathbb{P}$ *by*

$$\mathrm{d}X_t = \nabla \log \varphi(t, X_t)\,\mathrm{d}t + \mathrm{d}W_t, \tag{2}$$

with $\mathrm{Law}(X_t) = \varphi(t,\cdot)\bar{\varphi}(T-t,\cdot)$ *and* $F := \frac{\delta f}{\delta m(t,x)}$ *where the pair* $(\varphi, \bar{\varphi})$ *satisfies the PDE system*

$$\partial_t\varphi_t + \tfrac{1}{2}\Delta\varphi_t = c\,\varphi_t\,F(t,x,\varphi_t\bar{\varphi}_{T-t}), \tag{3}$$

$$\partial_t\bar{\varphi}_t + \tfrac{1}{2}\Delta\bar{\varphi}_t = c\,\bar{\varphi}_t\,F(T-t,x,\bar{\varphi}_t\varphi_{T-t}), \tag{4}$$

The terminal conditions $(\varphi(T,\cdot), \bar{\varphi}(T,\cdot))$ *are characterized by the coupled* nonlinear Sinkhorn system

$$\begin{cases} \varphi(0,x)\,\bar{\varphi}(T,x) = \mu_0(x), \\ \bar{\varphi}(0,x)\,\varphi(T,x) = G^{-1}\big(-\frac{1}{c}\log\varphi(T,\cdot),\,x\big), \end{cases} \tag{5}$$

where $G := \frac{\delta g}{\delta m(T,x)}$ *and* $\varphi(0,\cdot), \bar{\varphi}(0,\cdot)$ *are obtained by propagating the terminal potentials backward through the PDE system* (3)-(4).

The proof is deferred to Appendix A.

Remark 2.2 (Nonlinear Sinkhorn interpretation). When $f \equiv 0$ with a hard terminal constraint, (5) reduces to the usual Sinkhorn system, with $(\varphi, \bar{\varphi})$ propagated by *linear* heat kernels and matched to μ_0, μ_T. Our formulation adds two nonlinear ingredients: the semilinear propagation (3)-(4) through the running cost F, and the terminal coupling $\varphi(T,\cdot) \mapsto G^{-1}(-\frac{1}{c}\log\varphi(T,\cdot),\cdot)$ induced by g. The values $\varphi(0,x), \bar{\varphi}(0,x)$ are obtained from the PDE system via the branching representation of Section 3.

A rigorous analysis of the PDE system is already available in (Lasry & Lions, 2007); here, for completeness, we include a formal sketch for readers less familiar with the mean-field game literature. In order to solve the coupled *nonlinear Sinkhorn system*, a key step is the evaluation of the quantity $\varphi(0, X_0)$. This computation will be performed via a branching Monte Carlo representation, which we develop and analyze in the following section.

Remark 2.3 (Forward-backward system). Although Theorem 2.1 is stated in terms of the system (3)-(4) with terminal conditions, it can equivalently be expressed in the usual forward-backward form familiar from mean field games. Indeed, defining the density $m(t,x) := \varphi(t,x)\bar{\varphi}(T-t,x)$, one checks that m satisfies the forward Fokker-Planck equation (see proof for details)

$$\partial_t m + \nabla\cdot\big(m\,\nabla\log\varphi\big) - \tfrac{1}{2}\Delta m = 0,$$

while $v := -\frac{1}{c}\log\varphi$ satisfies the backward Hamilton-Jacobi equation

$$\partial_t v + \tfrac{1}{2}\Delta v - \tfrac{c}{2}\,|\nabla v|^2 + F(t,x,m) = 0$$

Thus the pair (m, v) solves the standard forward-backward MFG system of Lasry - Lions type. Our formulation in terms of $(\varphi, \bar{\varphi})$ is algebraically equivalent but particularly well-suited for the branching Monte-Carlo representation introduced below.

Remark 2.4 (Soft constraint and matching marginal μ_T). Consider the terminal cost

$$g(m(T,\cdot)) = \lambda\, D_f(\mu_T, m(T,\cdot)), \qquad \lambda > 0,$$

with $f(u) := -\log u - 1$; this is the (reverse) Kullback-Leibler divergence $\mathrm{KL}(\mu_T \| m(T,\cdot))$ up to an additive constant. In this case, the convex conjugate g^* is

$$g^*(v(T,\cdot)) = \lambda\,\mathbb{E}_{X_T\sim\mu_T}\Big[e^{\,v(T,X_T)/\lambda}\Big], \tag{6}$$

so that the terminal coupling in the nonlinear Sinkhorn system (5) becomes explicit:

$$G^{-1}\big(-\tfrac{1}{c}\log\varphi(T,\cdot),x\big) = \mu_T(x)\,\varphi(T,x)^{-a}, \qquad a := \frac{1}{\lambda c}. \tag{7}$$

The system (5) then reads

$$\begin{cases} \varphi(0,x)\,\bar{\varphi}(T,x) = \mu_0(x), \\ \bar{\varphi}(0,x)\,\varphi(T,x) = \mu_T(x)\,\varphi(T,x)^{-a}. \end{cases} \tag{8}$$

In our numerical experiments, we consider this formulation for various values of λ. In the limit $\lambda \to \infty$ (i.e. $a \to 0$), the terminal penalty enforces the hard constraint $m(T,\cdot) = \mu_T$, and the terminal coupling in (8) reduces to the standard marginal constraint

$$\bar{\varphi}(0,x)\,\varphi(T,x) = \mu_T(x), \tag{9}$$

so that the nonlinear Sinkhorn system recovers the classical Schrödinger bridge.

Appendix F records an analogous soft terminal condition based on a p-Wasserstein discrepancy.

Remark 2.5 (Special cases of semilinear Schrödinger bridges). Appendix H records several tractable specializations: linear bridges with confining potentials and explicit heat kernels on Coxeter chambers (including Dyson and Bru-Wishart limits), the Gaussian semilinear bridge for log-linear F (Bunne et al., 2023) and a relative-entropy terminal cost (Proposition H.1), lifted bridges with Gyöngy's Markov projection (Proposition H.2), and score-based / DDPM models as the exact $F \equiv 0$ solution with $\bar{\varphi} \equiv 1$ and noised-data prior (Appendix H.2), with proofs in Appendices I and J.

3 BRANCHING DIFFUSIONS

The semilinear PDE system (3)-(4) is generally analytically intractable, and thus requires numerical approximation. A deterministic Sinkhorn-type iterative method based on PDE techniques was introduced by Guéant (2012); however, this approach suffers from the curse of dimensionality and is therefore not suitable in the context of generative modeling. An alternative and more powerful approach is provided by *marked branching diffusion representations*, originating in McKean's probabilistic solution of the KPP equation (McKean, 1975) and subsequently extended to broad classes of semilinear and nonlinear PDEs (Henry-Labordère, 2012; Henry-Labordère et al., 2019; Henry-Labordère & Touzi, 2021). These representations have been successfully applied to parabolic, elliptic, and even hyperbolic PDEs, including, notably, the complex Gross-Pitaevskii equation developed in Appendix G.

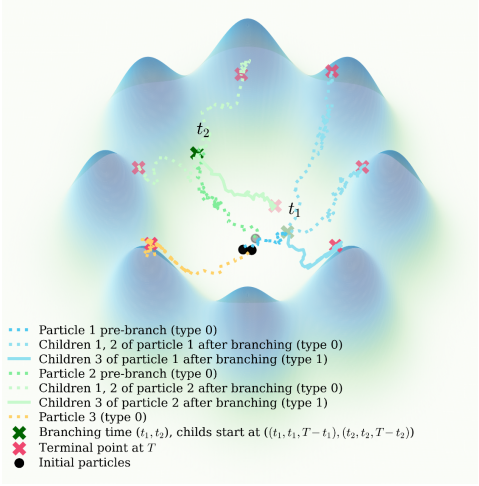

Figure 1: Branching paths on an 8-Gaussian terminal potential, illustrating Gross-Pitaevskii-type branching (Appendix M).

The key idea is to encode the nonlinear terms of the PDE into a branching mechanism: particles follow Brownian dynamics until a random exponential branching time, at which point they reproduce according to rules determined by the nonlinearity F. The resulting branching trees provide unbiased MC estimators of the PDE solution. This stochastic representation is *embarrassingly parallelizable*, which makes it particularly effective for large-scale simulations and high-dimensional problems.

More generally, when F is analytic in m, we truncate its power series at order K, yielding a polynomial nonlinearity amenable to branching. Without loss of generality, we write it in the *canonical form*

$$F(t,x,m) = 1 - \sum_{k=0}^{K} a_k(t,x)\, m^k := 1 - F_p(t,x,m), \tag{10}$$

where the coefficients $(a_k)_{k=0,\cdots,K}$ are functions of t and $x \in \mathbb{R}^d$. Thus the system (3)-(4) becomes

$$\begin{cases} \partial_t \varphi_t + \frac{1}{2}\Delta\varphi_t = c\,\varphi_t\Big(1 - F_p\big(t,x,\varphi_t\bar{\varphi}_{T-t}\big)\Big), \\ \partial_t \bar{\varphi}_t + \frac{1}{2}\Delta\bar{\varphi}_t = c\,\bar{\varphi}_t\Big(1 - F_p\big(T-t,x,\bar{\varphi}_t\varphi_{T-t}\big)\Big). \end{cases} \tag{11}$$

3.1 BRANCHING REPRESENTATION

We use a marked branching Brownian motion with two types of particles: white vertex $\circ$ represents φ, black vertex $\bullet$ represents $\bar{\varphi}$. The branching rate is fixed to be $\beta = c$, and we choose arbitrarily probabilities $p_k > 0$ for $k = 0, 1, \ldots, K$ with $\sum_{k=0}^{K} p_k = 1$. Each particle moves as a Brownian motion and lives an exponential time $\tau \sim \mathrm{Exp}(c)$. At a branching/death time s, draw $K^\star$ with probability $p_{K^\star}$ and apply

$$\begin{cases} \text{White parent } (\circ): & \circ_s \longrightarrow (k{+}1)\ \circ_s\ + k\,\bullet_{T-s}, \quad k = 0,1,\ldots,K, \\ \text{Black parent } (\bullet): & \bullet_s \longrightarrow k\ \circ_{T-s}\ + (k{+}1)\,\bullet_s, \quad k = 0,1,\ldots,K. \end{cases}$$

Here, the subscript denotes the (re-)starting time of each particle: offspring with the same type as the parent start at time s, and offspring with the opposite type start at time $T - s$.

Each simulation carries a multiplicative *weight*. At the beginning of the simulation the root particle at (t, x) is initialized with weight $\mathcal{W} = 1$. At every branching event occurring at time s with outcome $K^\star$, the weight is updated as

$$\mathcal{W} \longrightarrow \mathcal{W} \times \frac{a_{K^\star}(s, X_s)}{p_{K^\star}}. \tag{12}$$

where X_s is the position of the particle at time s. Thus the global weight of a branching tree is the product of the factors associated with all branchings. The weight is reset to 1 at the beginning of each independent Monte-Carlo simulation. At time T, each surviving white (resp. black) particle at position X_T^i (resp. $\bar{X}_T^j$) contributes the factor $\varphi(T, X_T^i)$ (resp. $\bar{\varphi}(T, \bar{X}_T^j)$). We denote by

$$\begin{cases} N_T^{\text{white}} & \text{the number of white particles at time } T, \text{ with positions } X_T^1, \dots, X_T^{N_T^{\text{white}}}, \\ \bar{N}_T^{\text{black}} & \text{the number of black particles at time } T, \text{ with positions } \bar{X}_T^1, \dots, \bar{X}_T^{\bar{N}_T^{\text{black}}}. \end{cases}$$

Each position X_T^i or $\bar{X}_T^j$ is generated by simulating the Brownian trajectories according to the branching mechanism until time T.

Proposition 3.1 (Unbiased stochastic representation for φ)**.** *Assume that $\varphi(T,\cdot), \bar{\varphi}(T,\cdot) \in L^\infty(\mathbb{R}^d)$ and that the coefficients $(a_k)_{k=0}^K$ belong to $L^\infty([0,T]\times\mathbb{R}^d)$. Set*

$$\Phi := \mathcal{W} \prod_{i=1}^{N_T^{\text{white}}} \varphi(T, X_T^i) \prod_{j=1}^{\bar{N}_T^{\text{black}}} \bar{\varphi}(T, \bar{X}_T^j).$$

The mild solution φ to system (3)–(4) *admits the unbiased square-integrable stochastic representation*

$$\varphi(t,x) = \mathbb{E}_{t,x}\big[\Phi\big].$$

The proof is deferred to Appendix L; illustrative Feynman-Galton-Watson diagrams in Appendix M, and a numerical sanity check of the time-reflected coupling $u(t,\cdot)u(T-t,\cdot)$ in Appendix K.

4 Stochastic-optimization form of the *nonlinear* Sinkhorn system

To enforce the two nonlinear Sinkhorn conditions (5), we match the associated measures using the Maximum Mean Discrepancy (MMD) (Gretton et al., 2012), which we now recall.

Definition 4.1 (Maximum Mean Discrepancy)**.** Let k be a positive definite kernel. For two probability laws P, Q on $\mathbb{R}^d$, the MMD admits the closed form

$$\text{MMD}_k^2(P,Q) = \mathbb{E}_{X,X'\sim P}[k(X,X')] + \mathbb{E}_{Y,Y'\sim Q}[k(Y,Y')] - 2\,\mathbb{E}_{X\sim P;Y\sim Q}[k(X,Y)],$$

where X, X' (resp. Y, Y') are independent copies from P (resp. Q).

Remark 4.1 (Characteristic kernel). If k is characteristic, then $\text{MMD}_k(P,Q) = 0 \iff P = Q$.

Background on MMD and RKHS is provided in Appendix C.

Proposition 4.1 (MMD objective)**.** *Let k be a characteristic kernel and define*

$$\mathcal{L}_{\text{MMD}} := \text{MMD}_k^2\Big(\varphi(0,\cdot)\,\bar{\varphi}(T,\cdot),\, \mu_0\Big) + \text{MMD}_k^2\Big(\bar{\varphi}(0,\cdot)\,\varphi(T,\cdot),\, G^{-1}\big(-\tfrac{1}{c}\log\varphi(T,\cdot),\cdot\big)\Big). \quad (13)$$

Then $\mathcal{L}_{\text{MMD}} \geq 0$, and $\mathcal{L}_{\text{MMD}} = 0 \iff (\varphi(T,\cdot), \bar{\varphi}(T,\cdot))$ satisfy the nonlinear Sinkhorn system (5)*.*

4.1 Meta-algorithm for training semilinear Schrödinger bridges

We parametrize the terminal potentials by neural networks $\varphi_\theta(x) := \varphi_\theta(T,x)$ and $\bar{\varphi}_\eta(x) := \bar{\varphi}_\eta(T,x)$. Let μ^* be a reference density that is easy to sample and evaluate.

For $Z \sim \mu^*$, we define $\widehat{A}_{\theta,\eta}(Z, \circ \text{ (resp. } \bullet))$ using a white (resp. black) rooted branching tree starting from $(0, Z)$:

$$\widehat{A}_{\theta,\eta}(Z,\circ) := \mathcal{W} \prod_{i=1}^{N_T^{\text{white}}} \varphi_\theta(Z_T^i) \prod_{j=1}^{\bar{N}_T^{\text{black}}} \bar{\varphi}_\eta(\bar{Z}_T^j),$$

with analogous definition for a black-rooted tree.

Algorithm 1: Training semilinear SB

Input: Laws μ_0, μ_T, μ^*, rate c, rule F, penalty $\lambda > 0$, kernel k, step γ, batch size B.

Output: Trained potentials $(\varphi_\theta, \bar{\varphi}_\eta)$.

Initialize θ, η;

while *not converged* **do**

  Draw $\{X_0^{(b)}\}_{b=1}^B \sim \mu_0$, $\{X_T^{(b)}\}_{b=1}^B \sim \mu_T$, and $\{Z^{(b)}\}_{b=1}^B \sim \mu^*$;

  Simulate white and black trees from each $Z^{(b)}$ and form $\widehat{A}_{\theta,\eta}(Z^{(b)},\circ)$ and $\widehat{A}_{\theta,\eta}(Z^{(b)},\bullet)$;

  $\widehat{\mathcal{L}}_{\text{MMD}} \leftarrow \frac{1}{B^2}\sum_{b,b'=1}^B J(b,b')$; *// all pairs*

  $(\theta,\eta) \leftarrow (\theta,\eta) - \gamma\nabla_{(\theta,\eta)}\widehat{\mathcal{L}}_{\text{MMD}}$;

In the KL terminal case of Remark 2.4, using the reweighted kernel $\tilde{k}(x,y) := \mu^*(x)\,k(x,y)\,\mu^*(y)$, which remains characteristic provided $\mu^* > 0$ almost everywhere, the MMD objective reads

$$\min_{\theta,\eta} \mathcal{L}_{\mathrm{MMD}}(\theta,\eta) := \min_{\theta,\eta} \mathbb{E}[J]. \tag{14}$$

where

$$\begin{aligned} J = {} & \widehat{A}_{\theta,\eta}(Z,\circ)\bar{\varphi}_\eta(Z)\widehat{A}_{\theta,\eta}(Z',\circ)\bar{\varphi}_\eta(Z')k(Z,Z') - 2\widehat{A}_{\theta,\eta}(Z,\circ)\bar{\varphi}_\eta(Z)k(Z,X_0')\mu^*(X_0') \\ & + \mu^*(X_0)k(X_0,X_0')\mu^*(X_0') + \widehat{A}_{\theta,\eta}(Z,\bullet)\varphi_\theta(Z)\widehat{A}_{\theta,\eta}(Z',\bullet)\varphi_\theta(Z')k(Z,Z') \\ & - 2\widehat{A}_{\theta,\eta}(Z,\bullet)\bar{\varphi}_\eta(Z)\varphi_\theta(X_T')^{\frac{-1}{\lambda c}}k(Z,X_T')\mu^*(X_T') + (\varphi_\theta(X_T)\varphi_\theta(X_T'))^{\frac{-1}{\lambda c}}\,\mu^*(X_T)k(X_T,X_T')\mu^*(X_T'). \end{aligned} \tag{15}$$

Here, $(Z,Z') \sim \mu^* \otimes \mu^*$, $(X_0,X_0') \sim \mu_0 \otimes \mu_0$, and $(X_T,X_T') \sim \mu_T \otimes \mu_T$ are drawn independently. Each term is estimated on mini-batches and minimized by stochastic gradient descent (Algorithm 1). A Gaussian kernel and a suitable reference density μ^* provide stable default choices (Appendix B).

4.2 Meta-algorithm for sampling with semilinear Schrödinger bridges

Once the potentials $(\varphi_\theta, \bar{\varphi}_\eta)$ have been learned, we can generate new samples by simulating the optimal controlled process (Algorithm 2): under $\mathbb{P}$, $\mathrm{d}X_t = \nabla \log \varphi_\theta(t,X_t)\,\mathrm{d}t + \mathrm{d}W_t$. We employ the branching diffusion representation of Proposition 3.1 to obtain unbiased Monte-Carlo estimates of $\nabla \log \varphi_\theta(t,X_t)$ at (t,X_t). Differentiation with respect to the spatial input x is then performed automatically using *autograd*, yielding a stochastic estimator of the score field. With this estimator at hand, we discretize the dynamics using the Euler-Maruyama scheme (or, optionally, a higher-order discretization).

Algorithm 2: Sampling from semilinear SB

Input: Trained $(\varphi_\theta, \bar{\varphi}_\eta)$, rate c, rule F, Euler steps N, horizon T, samples M.
Output: Approximate bridge path $(X_{t_k})_{k=0}^N$.

Initial $X_0 \sim \mu_0$, $\Delta t \leftarrow T/N$;
for $k = 0$ **to** $N-1$ **do**
  Simulate M white-rooted trees from (t_k, X_{t_k});
  Form $\widehat{\varphi}_k(x)$ of $\varphi(t_k,x)$:

$$\widehat{\varphi}_k(x) \leftarrow \tfrac{1}{M}\sum_{m=1}^{M} \mathcal{W}^m \prod_{i=1}^{N_T^{m,\mathrm{w}}} \varphi_\theta(X_{T,x}^{m,i}) \prod_{j=1}^{\bar{N}_T^{m,\mathrm{b}}} \bar{\varphi}_\eta(\bar{X}_{T,x}^{m,j})$$

  $v_k \leftarrow \nabla_x \log \widehat{\varphi}_k(x)|_{x=X_{t_k}}$; *// autograd*
  Draw $\xi_k \sim \mathcal{N}(0, I_d)$;
  $X_{t_{k+1}} \leftarrow X_{t_k} + \Delta t\, v_k + \sqrt{\Delta t}\,\xi_k$;

5 Numerical experiments

5.1 Implementation details

Each of $\varphi_\theta, \bar{\varphi}_\eta : \mathbb{R}^d \to (0,\infty)$ is a positive-scalar feedforward network. Unless stated otherwise, all use 3 to 8 fully connected layers with nonlinear activations. Widths are 16 for $d = 2$, 32 for $d = 50$, and 64 for $d = 512$ for latent FFHQ experiments. Potentials are trained jointly by stochastic gradient ascent on the MMD objective 14 (Adam, $\gamma = 10^{-4}$, $B = 64$). We use automatic differentiation of $\log \varphi_\theta$ for sampling (Section 4.2); see Appendix N. We train on 10^4 samples in 2D and 10^5 samples in 50D, for 300 epochs in both cases. Sampling uses between $N = 50$ and $N = 200$ Euler steps and $M \in \{10^3, 10^4\}$ branching trees per score evaluation (higher M in higher dimension). All runs use NVIDIA Tesla T4 GPUs; bridge training takes $\sim$30 s for 300 epochs on 10^4 points in 2D, and generation $\sim$3 s for 50 steps.

5.2 Results

Classical Schrödinger bridges ($F = 0$, $\lambda = \infty$, $d = 2$). We begin with controlled experiments in low dimensions to validate that our algorithm recovers the *classical Schrödinger bridge*. Specifically, we set $F \equiv 0$ and enforce a hard terminal constraint (see Remark 2.4 with $\lambda = \infty$), so that the problem reduces to entropic optimal transport.

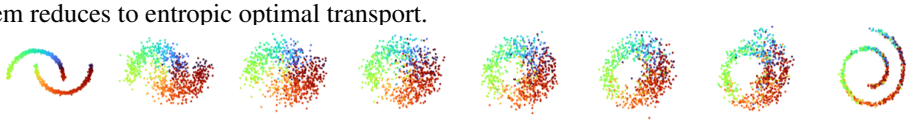

Figure 2: Temporal evolution of $\mathrm{Law}(X_t)$ from Moons $\to$ Swiss (source μ_0 on the left, target μ_T on the right).

We compare SBNL against DeepGSB (Liu et al., 2022), LightSB (Korotin et al., 2023), LightSBB-M (Alouadi et al., 2026), DSBM (Shi et al., 2023), and BranchSBM (Tang et al., 2026) on a Gaussian $\rightarrow$ mixture of 8 Gaussians benchmark, selecting methods that are closest to our setting in terms of modeling assumptions and problem formulation. Figure 3a displays terminal samples, while Table 3b reports quantitative distances to μ_T. SBNL captures all modes and achieves the best mixture W_2, whereas BranchSBM attains the lowest empirical W_2. Overall, SBNL remains competitive across evaluation metrics. Additional per-mode coverage results and full comparisons are provided in Appendix O.2.

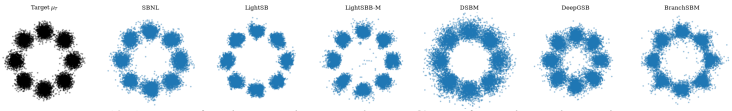

(3a) Terminal samples on the 8-Gaussians benchmark.

Method	MW_2	W_2^{emp}	E	$\|\Delta\bar{x}\|$	$\|\Delta\Sigma\|_F$
SBNL	$\mathbf{1.52}_{\pm 0.07}$	$1.14_{\pm 0.23}$	$0.023_{\pm 0.012}$	$0.29_{\pm 0.04}$	$2.97_{\pm 0.22}$
BranchSBM	$1.57_{\pm 0.14}$	$\mathbf{1.05}_{\pm 0.12}$	$\mathbf{0.018}_{\pm 0.009}$	$\mathbf{0.08}_{\pm 0.05}$	$\mathbf{1.74}_{\pm 0.19}$
LightSBB-M	$1.55_{\pm 0.27}$	$1.26_{\pm 0.15}$	$0.018_{\pm 0.009}$	$0.16_{\pm 0.07}$	$1.79_{\pm 0.39}$
LightSB	$1.54_{\pm 0.24}$	$1.22_{\pm 0.12}$	$0.024_{\pm 0.009}$	$0.11_{\pm 0.06}$	$4.69_{\pm 0.49}$
DSBM	$2.54_{\pm 0.15}$	$1.61_{\pm 0.12}$	$0.044_{\pm 0.017}$	$0.23_{\pm 0.09}$	$12.05_{\pm 0.84}$
DeepGSB	$2.48_{\pm 0.19}$	$2.02_{\pm 0.11}$	$0.199_{\pm 0.020}$	$0.24_{\pm 0.08}$	$25.93_{\pm 0.24}$

(3b) Comparison (MW_2 : closed-form Wasserstein, W_2^{emp} : empirical W_2 on 10^4 samples, E : energy distance, and $\|\Delta\bar{x}\|$ (resp. $\|\Delta\Sigma\|_F$) : Euclidean (resp. Frobenius) error on the sample mean (resp. covariance)).

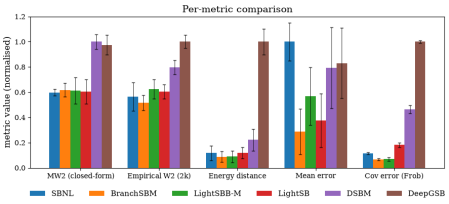

(3c) Visual comparison across all metrics, each column-wise max-normalized to 1 (lower is better).

High-dimensional mixtures ($F(m) = 0$, $\lambda = \infty$, $d = 50$). We use the same d=50 Gaussian-mixture benchmark as in (Shi et al., 2023) (Figures 21-22). The empirical mean, standard deviation, and skewness of $\text{Law}(X_T)$ converge to those of μ_T within 50 training epochs (Figure 23). We observe no variance collapse, attributable to the unbiased branching estimator of $\nabla \log \varphi_\theta$ avoiding the bias that score-matching pipelines accumulate at high d (De Bortoli et al., 2021).

Soft terminal conditions ($F = 0$, $\lambda < \infty$, $d = 2$). We next relax the hard constraint by introducing a finite penalty parameter λ. Figure 4a shows $\text{Law}(X_t)$ for decreasing values of λ, and Figure 4b reports the induced Kullback-Leibler divergence at the terminal time. As expected, $D_{\text{KL}}(\mu_T, \text{Law}(X_T))$ decreases monotonically with λ, confirming that stronger penalties enforce closer alignment with the target law.



(4a) Generation for $\lambda \in \{10^5, 10^4, 5{\times}10^3, 10^3, 10^2\}$.

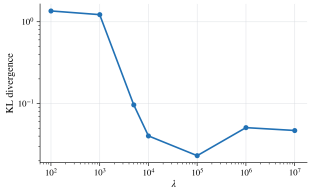

(4b) Terminal $D_{\text{KL}}(\mu_T, \text{Law}(X_T))$ versus λ.

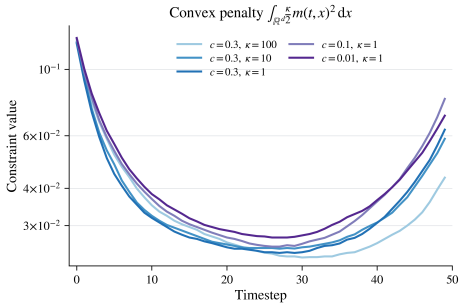

(4c) Evolution of the convex penalty during generation for $c \in \{0.01, 0.1, 0.3\}$ and $\kappa \in \{1, 10, 100\}$.

Gross-Pitaevskii penalty ($F(m) = \kappa m$, $\lambda = \infty$, $d = 2$). We also investigate the convex penalty functional $f(m(t,\cdot)) := \frac{\kappa}{2} \int_{\mathbb{R}^d} m(t,x)^2 \, \mathrm{d}x$, so that $F(m) = \kappa m$. Larger κ and c increase branching, making the dynamics more sensitive to the nonlinearity (Figure 4c).

Potential running cost ($F(t,x,m) = V(x)$, $\lambda = \infty$, $d = 2$). We finally illustrate a *state-dependent* running cost of potential type, $f(m(t,\cdot)) := \int_{\mathbb{R}^d} V(x)\, m(t,x) \, \mathrm{d}x$, so that $F(t,x,m) = V(x)$. This lets the environment *shape* the transport path: the optimal bridge routes mass around the obstacle while still matching both marginals. Figure 5a (resp. Figure 5b) shows the evolving samples (resp. density estimation) of $\text{Law}(X_t)$ deflecting around the high-potential region, and Figure 5c display a 3D evolution. Further potential landscapes are reported in Appendix O.3.

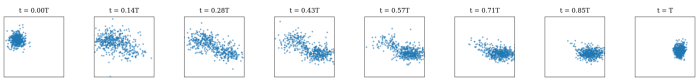

(5a) Samples Law(X_t) deflecting around high-V.

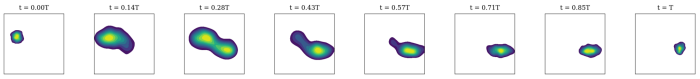

(5b) Estimated density $m(t,\cdot)$ snapshots, $t \in [0,T]$.

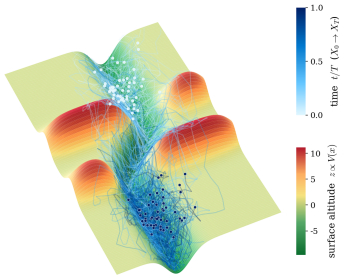

(5c) 3D evolution.

MNIST in latent space ($F(m)=0$, $\lambda=\infty$, $d=32$)**.** On MNIST we run the bridge in the d=32 latent space of a pretrained VAE on $(1,24,24)$ grayscale images, with μ_0,μ_T the codes of digits 1 and 3 and 6×10^3 samples per class for training. Figure 6 shows Law(X_t) and decoded trajectories; intermediate reconstructions stay visually smooth and digit-like.

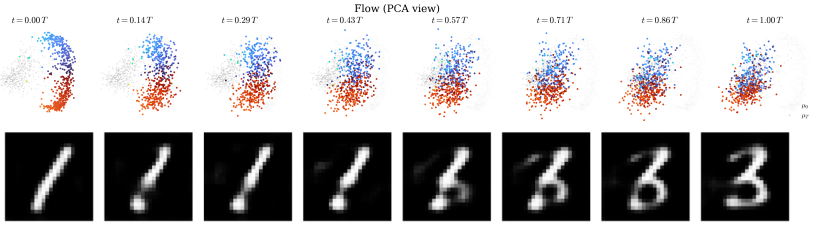

Figure 6: Temporal evolution of the latent variables and their corresponding decoded representations.

FFHQ young-to-old ($F(m)=0$, $\lambda=\infty$, $d=512$)**.** We evaluate FFHQ young-to-old transfer in the d=512 ALAE latent space (Pidhorskyi et al., 2020; Alouadi et al.), using 3×10^4 samples per age group. Figure 7 shows decoded trajectories and, for each baseline, the source child and its decoded terminal sample. SBNL preserves identity and attributes, while several baselines collapse to a near-canonical face; the extended panel is in Appendix O.7.



(7a) Comparison of source (left) and terminal samples.

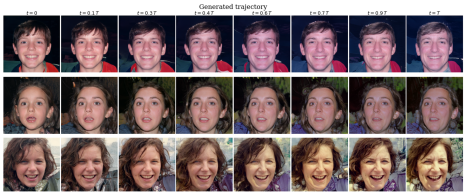

(7b) Decoded trajectories during sampling.

FID of the different methods

SBNL	LightSB	LightSBB-M	DSBM	DeepGSB	BranchSBM-F	BranchSBM-M
36.03	59.90	40.76	38.92	66.39	61.96	39.54

Figure 7: FFHQ young-to-old in ALAE latent space. Lower FID is better.

6 Discussion

We formulate semilinear mean-field Schrödinger bridges with law-dependent $f(m(t,\cdot))$, $g(m(T,\cdot))$ and characterize optimality through coupled $(\varphi,\bar{\varphi})$ PDEs, equivalently a Lasry - Lions (m,v) system – beyond the linear entropic setting of Sinkhorn and iterative Markovian fitting. We recast the terminal conditions as a *nonlinear* Sinkhorn system, which we solve by Maximum Mean Discrepancy, using branching diffusions to obtain unbiased stochastic representations of φ. Our experiments cover bridge recovery, Gross-Pitaevskii-type couplings, potential interactions, and latent transport.

The numerics are encouraging, but we do not claim systematic superiority over mature diffusion Schrödinger-bridge pipelines in the linear setting. Training in the nonlinear regime still requires careful tuning of the hyperparameters (branching rate c, kernel bandwidths, penalty λ) to remain stable, and both the parallelization of the branching simulation and the variance of the underlying Monte Carlo estimators leave room for improvement. We see this work as a new variational and optimality system whose first mesh-free realization invites variance reduction and hybridization with standard training pipelines.

REFERENCES

A. Alouadi, P. Henry-Labordère, G. Loeper, H. Pham, and N. Touzi. Bridging Schrödinger and Bass for generative modelling. Preprint.

Alexandre Alouadi, Pierre Henry-Labordère, Grégoire Loeper, Othmane Mazhar, Huyên Pham, and Nizar Touzi. Lightsbb-m: Bridging schrödinger and bass for generative diffusion modeling, 2026. URL https://arxiv.org/abs/2601.19312.

Jean-David Benamou, Guillaume Carlier, Simone Di Marino, and Luca Nenna. An entropy minimization approach to second-order variational mean-field games. *Mathematical Models and Methods in Applied Sciences*, 29(8):1553–1583, 2019. doi: 10.1142/S0218202519500283.

Thomas Bonnemain, Thierry Gobron, and Denis Ullmo. Schrödinger approach to mean field games with negative coordination. *SciPost Physics*, 9(4):059, 2020. doi: 10.21468/SciPostPhys.9.4.059.

Charlotte Bunne, Ya-Ping Hsieh, Marco Cuturi, and Andreas Krause. The schrödinger bridge between gaussian measures has a closed form. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 5802–5833. PMLR, 2023. URL https://proceedings.mlr.press/v206/bunne23a.html.

Pierre Cardaliaguet. Notes on mean field games (from p.-l. lions' lectures at collège de france). URL https://www.ceremade.dauphine.fr/~cardaliaguet/MFG20130420.pdf.

Thierry Cazenave. *Semilinear Schrödinger equations*, volume 10 of *Courant Lecture Notes in Mathematics*. American Mathematical Society, 2003.

Tianrong Chen, Guan-Horng Liu, and Evangelos A. Theodorou. Likelihood training of schrödinger bridge using forward-backward sdes theory, 2023. URL https://arxiv.org/abs/2110.11291.

Claude Cohen-Tannoudji, Bernard Diu, and Franck Laloë. *Mécanique quantique, Tome 3*. CNRS Editions, Paris, 2017. ISBN 9782705652404.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pp. 17695–17709, 2021.

Julie Delon and Agnès Desolneux. A wasserstein-type distance in the space of gaussian mixture models. *SIAM Journal on Imaging Sciences*, 13(2):936–970, 2020.

Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky T. Q. Chen. Adjoint matching: Fine-tuning flow and diffusion generative models with memoryless stochastic optimal control, 2025. URL https://arxiv.org/abs/2409.08861.

Bruno Dupire. Pricing with a smile. *Risk*, 7(1):18–20, 1994.

Hans Föllmer. Random fields and diffusion processes. In P.-L. Hennequin (ed.), *École d'Été de Probabilités de Saint-Flour XV–XVII, 1985–87*, volume 1362 of *Lecture Notes in Mathematics*, pp. 101–203. Springer, 1988. doi: 10.1007/BFb0086180.

David J. Grabiner. Brownian motion in a weyl chamber, non-colliding particles, and random matrices. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 35(2):177–204, 1999. doi: 10.1016/S0246-0203(99)80003-0.

Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012. URL http://jmlr.org/papers/v13/gretton12a.html.

Olivier Guéant. Mean field games equations with quadratic hamiltonian: A specific approach. *Mathematical Models and Methods in Applied Sciences*, 22(9):1250022, 2012. doi: 10.1142/S0218202512500222.

Olivier Guéant, Jean-Michel Lasry, and Pierre-Louis Lions. Mean field games and applications. In *Paris-Princeton Lectures on Mathematical Finance 2010*, volume 2003 of *Lecture Notes in Mathematics*, pp. 205–266. Springer, 2011. doi: 10.1007/978-3-642-14660-2_3.

Nikita Gushchin, Sergei Kholkin, Evgeny Burnaev, and Alexander Korotin. Light and optimal schrödinger bridge matching, 2024. URL https://arxiv.org/abs/2402.03207.

Julien Guyon and Pierre Henry-Labordère. *Nonlinear Option Pricing*. Chapman & Hall/CRC, 2011.

István Gyöngy. Mimicking the one-dimensional marginal distributions of processes having an itô differential. *Probability Theory and Related Fields*, 71(4):501–516, 1986.

Pierre Henry-Labordère. Cutting cva's complexity, 2012. URL https://ssrn.com/abstract=1995502. Risk magazine.

Pierre Henry-Labordère and Nizar Touzi. Branching diffusion representation for nonlinear cauchy problems and monte carlo approximation. *The Annals of Applied Probability*, 31(5):2350–2375, October 2021. doi: 10.1214/20-AAP1649.

Pierre Henry-Labordère, Nadia Oudjane, Xiaolu Tan, Nizar Touzi, and Xavier Warin. Branching diffusion representation of semilinear pdes and monte carlo approximation. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 55(1):184–210, 2019. doi: 10.1214/17-AIHP880.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, pp. 6840–6851, 2020.

Alexander Korotin, Nikita Gushchin, and Evgeny Burnaev. Light schrödinger bridge, 2023. URL https://arxiv.org/abs/2310.01174.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.

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.

Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos A. Theodorou. Deep generalized schrödinger bridge, 2022. URL https://arxiv.org/abs/2209.09893.

Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos A. Theodorou, and Ricky T. Q. Chen. Generalized schrödinger bridge matching, 2024. URL https://arxiv.org/abs/2310.02233.

Henry P. McKean. Application of brownian motion to the equation of kpp. *Communications on Pure and Applied Mathematics*, 28(3):323–331, 1975.

Stanislav Pidhorskyi, Donald Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders, 2020. URL https://arxiv.org/abs/2004.04467.

Erwin Schrödinger. Über die umkehrung der naturgesetze. *Sitzungsberichte der Preussischen Akademie der Wissenschaften. Physikalisch-mathematische Klasse*, pp. 144–153, 1931.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion Schrödinger bridge matching. In *Advances in Neural Information Processing Systems*, volume 36, pp. 77995–78009, 2023.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.

Catherine Sulem and Pierre-Louis Sulem. *The nonlinear Schrödinger equation: self-focusing and wave collapse*, volume 139 of *Applied Mathematical Sciences*. Springer, 1999.

Igor Swiecicki, Thierry Gobron, and Denis Ullmo. Schrödinger approach to mean field games. *Physical Review Letters*, 116(12):128701, 2016. doi: 10.1103/PhysRevLett.116.128701.

Gábor J. Székely and Maria L. Rizzo. Testing for equal distributions in high dimension. *InterStat*, 5(16.10): 1249–1272, 2004.

Sophia Tang, Yinuo Zhang, Alexander Tong, and Pranam Chatterjee. Branched schrödinger bridge matching, 2026. URL https://arxiv.org/abs/2506.09007.

Cédric Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, 2003.

Zhenyi Zhang, Zihan Wang, Yuhao Sun, Tiejun Li, and Peijie Zhou. Modeling cell dynamics and interactions with unbalanced mean field schrödinger bridge, 2025. URL https://arxiv.org/abs/2505.11197.

A PROOF OF THEOREM 2.1

Proof. We start from the variational problem

$$P = \inf_{\alpha.} \mathbb{E}_{X_0 \sim \mu_0} \left[\int_0^T \left(\frac{1}{2c} |\alpha_t|^2 + f(m(t,\cdot)) \right) \mathrm{d}t \; + \; g(m(T,\cdot)) \right],$$

where the state process X_t follows the controlled diffusion

$$\mathrm{d}X_t = \alpha_t \, \mathrm{d}t + \mathrm{d}W_t,$$

and the marginal density $m(t,x)$ satisfies the Fokker-Planck equation

$$\partial_t m(t,x) = -\nabla_x \cdot \big(\alpha(t,x)\, m(t,x)\big) + \tfrac{1}{2}\Delta m(t,x), \quad m(0,x) = \mu_0(x).$$

Step 1: Lagrangian formulation. Introducing a Lagrange multiplier $v(t,x)$ for the Fokker-Planck constraint, we obtain by weak duality :

$$P \geq \sup_{v} \inf_{\alpha., m} \left\{ \int_0^T \int_{\mathbb{R}^d} \frac{1}{2c} |\alpha(t,x)|^2 m(t,x) \, \mathrm{d}x \, \mathrm{d}t + \int_0^T f(m(t,\cdot)) \, \mathrm{d}t + g(m(T,\cdot)) \right.$$
$$\left. - \int_0^T \int_{\mathbb{R}^d} v(t,x) \Big(\partial_t m + \nabla_x \cdot (\alpha m) - \tfrac{1}{2}\Delta m \Big) \mathrm{d}x \, \mathrm{d}t \right\}.$$

Step 2: Integration by parts. After integrating by parts in time and space (boundary terms at spatial infinity vanish), this becomes

$$P \geq \sup_{v} \inf_{\alpha., m} \left\{ \int_0^T \int_{\mathbb{R}^d} \frac{1}{2c} \, |\alpha(t,x)|^2 \, m(t,x) \, \mathrm{d}x \, \mathrm{d}t + \int_0^T f\big(m(t,\cdot)\big) \, \mathrm{d}t \right.$$
$$+ \int_0^T \int_{\mathbb{R}^d} \Big(\partial_t v(t,x) + \alpha(t,x) \cdot \nabla_x v(t,x) + \tfrac{1}{2}\Delta v(t,x) \Big) \, m(t,x) \, \mathrm{d}x \, \mathrm{d}t$$
$$\left. + \, g\big(m(T,\cdot)\big) - \int_{\mathbb{R}^d} v(T,x) \, m(T,x) \, \mathrm{d}x + \int_{\mathbb{R}^d} v(0,x) \, \mu_0(x) \, \mathrm{d}x \right\}. \tag{16}$$

Step 3: First-order conditions. Optimizing with respect to α yields

$$\alpha^*(t,x) = -c \, \nabla_x v(t,x).$$

The dual problem therefore becomes

$$P \geq \sup_{v} \left\{ \int_0^T f^*\big(\mathcal{L}v(t,\cdot)\big) \, \mathrm{d}t - g^*(v(T,\cdot)) + \mathbb{E}_{X_0 \sim \mu_0}\big[v(0,X_0)\big] \right\}, \tag{17}$$

where

$$\begin{cases} \mathcal{L}v := \partial_t v - \dfrac{c}{2}|\nabla_x v|^2 + \dfrac{1}{2}\Delta v, \\ f^*(\psi) := \inf\limits_{m} \left\{ f(m) + \displaystyle\int_{\mathbb{R}^d} \psi(x) \, m(x) \, \mathrm{d}x \right\}, \\ g^*(\phi) := \sup\limits_{m} \left\{ \displaystyle\int_{\mathbb{R}^d} \phi(x) \, m(x) \, \mathrm{d}x - g(m) \right\}. \end{cases}$$

Guessing the sup over $v(t,\cdot)$; we define the coupled forward-backward system

$$\begin{cases} \partial_t v^* + \dfrac{1}{2}\Delta v^* - \dfrac{c}{2}|\nabla_x v^*|^2 + F\big(t,x,m^*(t,x)\big) = 0, \qquad v^*(T,\cdot) \text{ given}, \\ \partial_t m^* = c \, \nabla_x \cdot \big(m^* \nabla_x v^*\big) + \dfrac{1}{2}\Delta m^*, \qquad m^*(0,\cdot) = \mu_0. \end{cases} \tag{18}$$

which is equivalent to the equation on $v^*(t,\cdot)$ only :

$$\partial_t m^* = c\,\nabla_x\cdot\big(m^*\nabla_x v^*\big) + \frac{1}{2}\Delta m^*, \qquad \text{with} \quad m^* = f^{*\prime}(\mathcal{L}v^*).$$

The existence and uniqueness of such a PDE system is guaranteed by Lasry - Lions result under assumptions (A1) to (A3) (Appendix D).

Using such a $v^*(t,\cdot)$:

$$P \geq \sup_{v^*(T,\cdot)} \left\{ \int_0^T f^*\big(\mathcal{L}v^*(t,\cdot)\big)\,\mathrm{d}t - g^*(v^*(T,\cdot)) + \mathbb{E}_{X_0\sim\mu_0}\big[v^*(0,X_0)\big] \right\}. \tag{19}$$

Reversing the integration by parts gives

$$P \geq \sup_{v^*(T,\cdot)} \left\{ \int_0^T\!\!\int_{\mathbb{R}^d} \frac{c}{2}\,|\nabla_x v^*|^2\, m^*\,\mathrm{d}x\,\mathrm{d}t + \int_0^T f(m^*(t,\cdot))\,\mathrm{d}t + g(m^*(T,\cdot)) \right\}. \tag{20}$$

On the other hand, the primal problem reads

$$P = \inf_{\alpha,m} \left\{ \int_0^T\!\!\int_{\mathbb{R}^d} \frac{1}{2c}|\alpha(t,x)|^2 m(t,x)\,\mathrm{d}x\,\mathrm{d}t + \int_0^T f(m(t,\cdot))\,\mathrm{d}t + g(m(T,\cdot)) \right\}.$$

Choosing

$$\alpha^*(t,x) = -c\,\nabla_x v^*(t,x), \qquad m = m^*,$$

gives with arbitrary sup over $v^*(T,\cdot)$:

$$P \leq \sup_{v^*(T,\cdot)} \left\{ \int_0^T\!\!\int_{\mathbb{R}^d} \frac{c}{2}\,|\nabla_x v^*|^2\, m^*\,\mathrm{d}x\,\mathrm{d}t + \int_0^T f(m^*(t,\cdot))\,\mathrm{d}t + g(m^*(T,\cdot)) \right\}. \tag{21}$$

The primal problem can now be written as

$$P \;=\; \sup_{v^*(T,\cdot)} \left\{ \int_0^T f^*\big(\mathcal{L}v^*(t,\cdot)\big)\,\mathrm{d}t - g^*\big(v^*(T,\cdot)\big) + \mathbb{E}_{X_0\sim\mu_0}\big[v^*(0,X_0)\big] \right\},$$

Combining (20) and (21), we conclude that equality holds. Hence the pair (v^*, m^*) solving the coupled forward-backward system attains the value of the primal optimization problem.

Step 4: Optimality condition. Taking the functional derivative with respect to $v^*(T,x)$, we obtain

$$\frac{\delta P}{\delta v^*(T,x)} = \int_0^T\!\!\int_y m^*(t,y)\frac{\delta \mathcal{L}v^*(t,y)}{\delta v^*(T,x)}\,\mathrm{d}t\,\mathrm{d}y - G^{-1}\big(v^*(T,\cdot),x\big) + \frac{\delta \mathbb{E}_{X_0\sim\mu_0}\big[v^*(0,X_0)\big]}{\delta v^*(T,x)}.$$

Applying the chain rule to $\mathcal{L}v^*$ yields

$$\frac{\delta \mathcal{L}v^*(t,y)}{\delta v^*(T,x)} = \partial_t \frac{\delta v^*(t,y)}{\delta v^*(T,x)} - c\,\nabla_y \frac{\delta v^*(t,y)}{\delta v^*(T,x)}\cdot\nabla_y v^*(t,y) + \frac{1}{2}\Delta_y \frac{\delta v^*(t,y)}{\delta v^*(T,x)}.$$

Integrating by parts in the time integral and collecting the boundary contributions, we get

$$\begin{aligned}
\frac{\delta P}{\delta v^*(T,x)} &= \int_0^T\!\!\int_y \underbrace{\Big(-\partial_t \cdot + c\nabla_y\big(\cdot\,\nabla_y v^*(t,y)\big) + \frac{1}{2}\Delta_y \cdot \Big)\Big(m^*(t,y)\Big)}_{=0\ (18)} \frac{\delta v^*(t,y)}{\delta v^*(T,x)}\,\mathrm{d}t\,\mathrm{d}y \\
&\quad + \int_y \left[m^*(t,y)\frac{\delta v^*(t,y)}{\delta v^*(T,x)} \right]_{t=0}^{t=T} \mathrm{d}y - G^{-1}\big(v^*(T,\cdot),x\big) + \frac{\delta \mathbb{E}_{X_0\sim\mu_0}\big[v^*(0,X_0)\big]}{\delta v^*(T,x)} \\
&= m^*(T,x) - \underbrace{\int_y m^*(0,y)\frac{\delta v^*(0,y)}{\delta v^*(T,x)}\mathrm{d}y}_{\frac{\delta \mathbb{E}_{X_0\sim\mu_0}\left[v^*(0,X_0)\right]}{\delta v^*(T,x)}} - G^{-1}\big(v^*(T,\cdot),x\big) + \frac{\delta \mathbb{E}_{X_0\sim\mu_0}\big[v^*(0,X_0)\big]}{\delta v^*(T,x)} \\
&= m^*(T,x) - G^{-1}\big(v^*(T,\cdot),x\big).
\end{aligned}$$

We get the optimal terminal condition

$$v^*(T,x) = G(m^*(T,\cdot),x)\,. \tag{22}$$

Step 5: Logarithmic transformation. Setting

$$v^*(t,x) = -\tfrac{1}{c}\log\varphi(t,x),$$

transforms (18) into

$$\partial_t\varphi + \tfrac{1}{2}\Delta\varphi = c\,\varphi\,F, \qquad F := \tfrac{\delta f}{\delta m(t,x)}.$$

The corresponding Fokker-Planck equation with optimal control (A) becomes

$$\partial_t m = \nabla\big(m(c\nabla v)\big) + \tfrac{1}{2}\Delta m. \tag{23}$$

Looking for a factorization of the form $m(t,x) = \varphi(t,x)\hat{\varphi}(t,x)$ yields the forward-backward system

$$\begin{cases} \partial_t\varphi + \tfrac{1}{2}\Delta\varphi = c\,\varphi F, & \varphi(T,\cdot) \text{ given}, \\ \partial_t\hat{\varphi} - \tfrac{1}{2}\Delta\hat{\varphi} = -c\,\hat{\varphi}F, & \hat{\varphi}(0,\cdot) \text{ given}, \end{cases}$$

with coupling condition $\varphi(0,x)\hat{\varphi}(0,x) = \mu_0(x)$. The forward PDE is transformed into a backward PDE by setting $\bar{\varphi}(t,\cdot) := \hat{\varphi}(T-t,\cdot)$.

$$\begin{cases} \partial_t\varphi_t + \tfrac{1}{2}\Delta\varphi_t = c\,\varphi_t\,F(t,x,\varphi_t\bar{\varphi}_{T-t}), & \varphi(T,\cdot) \text{ given}, \\ \partial_t\bar{\varphi}_t + \tfrac{1}{2}\Delta\bar{\varphi}_t = c\,\bar{\varphi}_t\,F(T-t,x,\bar{\varphi}_t\varphi_{T-t}), & \bar{\varphi}(T,\cdot) \text{ given}, \end{cases}$$

Together with the coupling conditions inherited from the factorization and the terminal optimality condition of Step 4, this yields the first-order *nonlinear Sinkhorn system*

$$\begin{cases} \varphi(0,x)\,\bar{\varphi}(T,x) = \mu_0(x), \\ \bar{\varphi}(0,x)\,\varphi(T,x) = G^{-1}(v^*(T,\cdot),x), \end{cases} \tag{24}$$

which characterizes the terminal potentials $(\varphi(T,\cdot),\bar{\varphi}(T,\cdot))$.

This completes the proof. □

B STOCHASTIC-OPTIMIZATION FORMS OF THE *nonlinear* SINKHORN SYSTEM

Using the branching representation of Proposition 3.1, the nonlinear Sinkhorn conditions read

$$\begin{cases} \mathbb{E}_{x,\text{white}}\left[\mathcal{W}\displaystyle\prod_{i=1}^{N_T^{\text{white}}}\varphi(T,X_T^i)\prod_{j=1}^{\bar{N}_T^{\text{black}}}\bar{\varphi}(T,\bar{X}_T^j)\right]\bar{\varphi}(T,x) = \mu_0(x), \\ \mathbb{E}_{x,\text{black}}\left[\mathcal{W}\displaystyle\prod_{i=1}^{N_T^{\text{white}}}\varphi(T,X_T^i)\prod_{j=1}^{\bar{N}_T^{\text{black}}}\bar{\varphi}(T,\bar{X}_T^j)\right]\varphi(T,x) = G^{-1}(-\tfrac{1}{c}\log\varphi(T,\cdot),x). \end{cases} \tag{25}$$

We now recast the two nonlinear Sinkhorn conditions in sample-based forms that can be used to train neural terminal potentials. We parametrize the terminal potentials $\varphi(T,\cdot),\bar{\varphi}(T,\cdot)$ by neural networks and, with a slight abuse of notation, write

$$\varphi_\theta(x) := \varphi(T,x), \qquad \bar{\varphi}_\eta(x) := \bar{\varphi}(T,x),$$

and let μ^* be a reference probability density on $\mathbb{R}^d$ that is easy to sample and evaluate.

Following the notation of Section 4.1, for $Z \sim \mu^*$ a white- (resp. black-) rooted branching tree started at $(0,Z)$ yields the unbiased one-tree estimators $\widehat{A}_{\theta,\eta}(Z,\circ)$ (resp. $\widehat{A}_{\theta,\eta}(Z,\bullet)$), with

$$\begin{cases} \mathbb{E}\big[\widehat{A}_{\theta,\eta}(Z,\circ)\big] = A_{\theta,\eta}(Z,\circ) \approx \varphi(0,Z), \\ \mathbb{E}\big[\widehat{A}_{\theta,\eta}(Z,\bullet)\big] = A_{\theta,\eta}(Z,\bullet) \approx \bar{\varphi}(0,Z), \end{cases} \tag{26}$$

where the approximation replaces the true terminal potentials by their neural-network parametrization $(\varphi_\theta, \bar{\varphi}_\eta)$. The two boundary equations (5) thus read

$$\begin{cases} A_{\theta,\eta}(x, \circ)\, \bar{\varphi}_\eta(x) = \mu_0(x), \\ A_{\theta,\eta}(x, \bullet)\, \varphi_\theta(x) = G^{-1}\big(-\frac{1}{c}\log \varphi_\theta, x\big). \end{cases} \tag{27}$$

In the KL terminal case of (6), the second condition becomes $A_{\theta,\eta}(x, \bullet)\, \varphi_\theta(x) = \mu_T(x)\, \varphi_\theta(x)^{-a}$ with $a := 1/(\lambda c)$.

Possibility 1: RKHS/MMD matching. Each boundary equation matches a (reweighted) measure to a marginal, so we enforce them by minimizing the MMD between the two sides. Using a characteristic kernel k with RKHS $\mathcal{H}$ (Appendix C), we set

$$\mathcal{L}_{\mathrm{MMD}}(\theta, \eta) := \mathrm{MMD}_k^2\big(A_{\theta,\eta}(\cdot, \circ)\, \bar{\varphi}_\eta,\ \mu_0\big) + \mathrm{MMD}_k^2\big(A_{\theta,\eta}(\cdot, \bullet)\, \varphi_\theta,\ \varphi_\theta^{-a}\mu_T\big).$$

Estimating $A_{\theta,\eta}(\cdot, \circ), A_{\theta,\eta}(\cdot, \bullet)$ by the branching trees and importance-sampling against μ^*, the closed form of Definition 4.1 gives the sample-based objective (14)-(15), minimized directly by stochastic gradient descent (Algorithm 1); a mixture of Gaussian kernels is a stable default.

C RKHS AND MMD

For completeness, we recall the definitions and elementary properties of reproducing kernel Hilbert spaces (RKHS) and the maximum mean discrepancy (MMD) used in Section 4.1 and in the proof of Theorem 2.1. We refer to (Gretton et al., 2012) for a detailed treatment.

Reproducing kernel Hilbert space. Let $k : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ be a symmetric positive definite kernel, i.e. for any $n \in \mathbb{N}$, points $x_1, \ldots, x_n \in \mathbb{R}^d$ and scalars $c_1, \ldots, c_n \in \mathbb{R}$,

$$\sum_{i,j=1}^{n} c_i c_j\, k(x_i, x_j)\ \geq\ 0.$$

By the Moore-Aronszajn theorem, there exists a unique Hilbert space $\mathcal{H}_k$ of functions $f : \mathbb{R}^d \to \mathbb{R}$, called the *reproducing kernel Hilbert space* of k, such that:

- for every $x \in \mathbb{R}^d$, the function $k(x, \cdot) \in \mathcal{H}_k$;
- (*reproducing property*) for every $f \in \mathcal{H}_k$ and every $x \in \mathbb{R}^d$,

$$f(x)\ =\ \langle f,\, k(x, \cdot)\rangle_{\mathcal{H}_k}.$$

In particular, taking $f = k(y, \cdot)$ gives $\langle k(x, \cdot), k(y, \cdot)\rangle_{\mathcal{H}_k} = k(x, y)$, so that k acts as an inner product between feature maps $x \mapsto k(x, \cdot)$.

Maximum mean discrepancy. Given two probability measures P, Q on $\mathbb{R}^d$, the *maximum mean discrepancy* associated with k is defined as the largest difference in expectations over the unit ball of $\mathcal{H}_k$:

$$\mathrm{MMD}_k(P, Q)\ :=\ \sup_{\|f\|_{\mathcal{H}_k} \leq 1} \Big\{\ \mathbb{E}_{X\sim P}\big[f(X)\big] - \mathbb{E}_{Y\sim Q}\big[f(Y)\big]\ \Big\}. \tag{28}$$

Intuitively, MMD searches for a witness function f in the RKHS unit ball that best distinguishes samples from P and Q.

Kernel mean embedding. Assume the integrability condition $\mathbb{E}_{X\sim P}[\sqrt{k(X, X)}] < \infty$. The *kernel mean embedding* of P is the element of $\mathcal{H}_k$ defined by

$$\mu_P(\cdot)\ :=\ \mathbb{E}_{X\sim P}\big[k(X, \cdot)\big]\ \in\ \mathcal{H}_k. \tag{29}$$

By the reproducing property, it satisfies $\mathbb{E}_{X\sim P}[f(X)] = \langle f, \mu_P\rangle_{\mathcal{H}_k}$ for all $f \in \mathcal{H}_k$, i.e. integration against P is represented by an inner product with μ_P.

MMD as a distance between embeddings. Substituting $\mathbb{E}_P[f] = \langle f, \mu_P \rangle_{\mathcal{H}_k}$ and $\mathbb{E}_Q[f] = \langle f, \mu_Q \rangle_{\mathcal{H}_k}$ into (28), the supremum over the unit ball is attained at $f = (\mu_P - \mu_Q)/\|\mu_P - \mu_Q\|_{\mathcal{H}_k}$, and the MMD reduces to the RKHS distance between the mean embeddings:

$$\mathrm{MMD}_k(P, Q) \;=\; \big\|\mu_P - \mu_Q\big\|_{\mathcal{H}_k}. \tag{30}$$

Expansion via the reproducing property. Expanding the squared norm in (30) and using $\langle k(x,\cdot), k(y,\cdot)\rangle_{\mathcal{H}_k} = k(x,y)$ gives the closed-form expression

$$\mathrm{MMD}_k^2(P, Q) \;=\; \mathbb{E}_{X,X'\sim P}\big[k(X,X')\big] + \mathbb{E}_{Y,Y'\sim Q}\big[k(Y,Y')\big] - 2\,\mathbb{E}_{\substack{X\sim P\\ Y\sim Q}}\big[k(X,Y)\big], \tag{31}$$

where X, X' (resp. Y, Y') are independent copies drawn from P (resp. Q). This form only involves expectations of kernel evaluations and is therefore directly estimable from samples, which is what makes the objective (14) tractable.

Characteristic kernels. The kernel k is said to be *characteristic* if the embedding map $P \mapsto \mu_P$ is injective on the space of probability measures. Equivalently,

$$\mathrm{MMD}_k(P, Q) = 0 \quad\Longleftrightarrow\quad P = Q.$$

Classical examples of characteristic kernels on $\mathbb{R}^d$ include the Gaussian kernel $k(x,y) = \exp(-\|x - y\|^2/2\sigma^2)$ and, more generally, mixtures of Gaussian kernels across several bandwidths, which we use as a stable default in our experiments.

D Standing assumptions for well-posedness

We use the following conditions, under which the associated forward-backward mean-field system admits a unique classical solution on $[0, T]$ (and thus the $(\varphi, \bar{\varphi})$-system in Theorem 2.1 is well posed).

(A1) *Regularity and growth.* The functional derivatives

$$F(t, x, m) := \frac{\delta f}{\delta m(t, x)}, \qquad G(x, m) := \frac{\delta g}{\delta m(T, x)}$$

exist, are continuous in (t, x) and Lipschitz in m with respect to the 2-Wasserstein metric, uniformly on compact time intervals. They have at most linear (or sub-quadratic) growth in x.

(A2) *Lasry - Lions monotonicity.* For any probability densities m_1, m_2,

$$\int_{\mathbb{R}^d} \Big(F(t, x, m_1) - F(t, x, m_2)\Big)\,\big(m_1(x) - m_2(x)\big)\,\mathrm{d}x \;\geq 0, \quad \forall t \in [0, T],$$

and

$$\int_{\mathbb{R}^d} \Big(G(x, m_1) - G(x, m_2)\Big)\,\big(m_1(x) - m_2(x)\big)\,\mathrm{d}x \;\geq 0.$$

These monotonicity conditions ensure uniqueness of the mean-field equilibrium (Lasry & Lions, 2007).

(A3) *Initial data.* The initial law $m(0, \cdot) = \mu_0$ has sufficient regularity (e.g., $\mu_0 \in L^1 \cap L^\infty$ with finite second moment).

E State dependence and relation to mean-field games

One could allow the running and terminal costs to depend explicitly on the state variable x, e.g.

$$\int_0^T \int_{\mathbb{R}^d} \ell\big(t, x, m(t,\cdot)\big)\, m(t, \mathrm{d}x)\,\mathrm{d}t \quad \text{and} \quad \int_{\mathbb{R}^d} h\big(x, m(T,\cdot)\big)\, m(T, \mathrm{d}x).$$

This causes no loss of generality: by setting

$$f\big(m(t,\cdot)\big) := \int_{\mathbb{R}^d} \ell\big(t, x, m(t,\cdot)\big)\, m(t, \mathrm{d}x), \qquad g\big(m(T,\cdot)\big) := \int_{\mathbb{R}^d} h\big(x, m(T,\cdot)\big)\, m(T, \mathrm{d}x),$$

we recover exactly the functional form (1), so that any explicit x-dependence can be "integrated out" into a dependence on $m(t,\cdot)$ and $m(T,\cdot)$. The resulting stochastic optimal control with mean-field (law) costs is standard in the mean-field game (MFG) literature; the corresponding optimality system is a forward-backward pair of PDEs (Fokker-Planck for m and HJB for the value), coupled through the functional derivatives $F = \delta f/\delta m$ and $G = \delta g/\delta m$ (see, e.g., (Lasry & Lions, 2007)) as made explicit in Theorem 2.1.

Altogether, the functional (1) extends the original entropic optimal transport problem of Schrödinger to a rich class of mean-field type control problems. It preserves the central feature that the quadratic control cost corresponds to entropy minimization, thereby maintaining analytic tractability.

F Soft terminal constraint via Wasserstein cost

Remark F.1 (Soft constraint via Wasserstein cost). In analogy with the f-divergence penalization of Remark 2.4, one may also impose a soft terminal condition through an optimal transport discrepancy. Specifically, taking

$$g(m) \;=\; \lambda\, W_p^p(\mu_T, m), \qquad \lambda > 0,$$

the convex conjugate g^* can be expressed in terms of the c-transform associated with the cost $|x-y|^p$:

$$g^*(v) = -\,\lambda\, \mathbb{E}_{X_T\sim\mu_T}\big[v^{\,c}(X_T)\big], \qquad v^{\,c}(x) := \inf_{y\in\mathbb{R}^d}\big\{|x-y|^p - v(y)\big\}.$$

In practice, directly computing $v^{\,c}$ in $\mathbb{R}^d$ may be challenging; a common approach is to consider an entropic relaxation of the transport cost, which yields a simpler expression more amenable to computation while retaining the same geometric intuition.

G Link to the (complex) semilinear Schrödinger equation and Gross-Pitaevskii

Consider the classical semilinear Schrödinger equation (NLS) on physical time $\tau \in \mathbb{R}$ (and its complex conjugate):

$$i\,\partial_\tau\psi \;\; = -\tfrac{1}{2}\Delta\psi \;+\; c\,\psi\,F(|\psi|^2), \tag{32}$$

$$-\,i\,\partial_\tau\bar\psi \;\; = -\tfrac{1}{2}\Delta\bar\psi \;+\; c\,\bar\psi\,F(|\psi|^2), \tag{33}$$

where $|\psi|^2 = \bar\psi\,\psi$ and F is a (real-valued) nonlinearity acting on the density. The Gross-Pitaevskii (cubic NLS) corresponds to the choice $F(u) = u$ (see e.g. (Cohen-Tannoudji et al., 2017)):

$$i\,\partial_\tau\psi \;=\; -\tfrac{1}{2}\Delta\psi \;+\; c\,\psi\,|\psi|^2.$$

Perform a Wick rotation to imaginary time by setting $\tau = \,i\,t$ (so that $\partial_\tau = -i\,\partial_t$). Substituting into (32)-(33) yields

$$\partial_t\psi + \tfrac{1}{2}\Delta\psi = c\,\psi\,F(\bar\psi\psi), \tag{34}$$

$$-\partial_t\bar\psi + \tfrac{1}{2}\Delta\bar\psi = c\,\bar\psi\,F(\bar\psi\psi). \tag{35}$$

Identifying $(\psi,\bar\psi)$ with $(\varphi,\bar\varphi)$, we recover exactly our system (3)-(4) written in forward-backward form (see Remark 2.3). In the cubic case (Gross-Pitaevskii), the sign of c distinguishes two qualitatively different regimes:

- Defocusing / repulsive ($c > 0$). The nonlinearity disperses mass, counteracting concentration. Classically, this regime admits global-in-time solutions for broad classes of initial data (e.g., in H^1), scattering in low dimensions, and exhibits no finite-time blow-up (Cazenave, 2003). After Wick rotation (imaginary time), the system becomes a well-behaved gradient(-like) flow for the energy; in our forward-backward formulation (3)-(4) this corresponds to a monotone mean-field interaction.
- Focusing / attractive ($c < 0$). The nonlinearity promotes concentration and can balance dispersion to form *bright solitons*; it may also produce finite-time blow-up (collapse) for sufficiently concentrated initial data (depending on the spatial dimension and mass) (Sulem & Sulem, 1999). Under Wick rotation, the flow may descend a *nonconvex* energy landscape; in (3)-(4) this yields a mean-field term with negative monotonicity.

H SPECIAL CASES OF SEMILINEAR SCHRÖDINGER BRIDGES

H.1 CONFINING POTENTIALS

When the interaction term reduces to a potential $F(t,x,m) \equiv V(t,x)$, the Schrödinger system (3)-(4) becomes *linear*. In this case, the optimal drift admits a Feynman-Kac representation:

$$\mathrm{d}X_t = \nabla \log \mathbb{E}\left[\varphi\big(T, B_T^{t,X_t}\big)\, \exp\Big(-c\int_t^T V(s, B_s^{t,X_t})\,\mathrm{d}s\Big)\right] \mathrm{d}t + \mathrm{d}W_t, \tag{36}$$

where $(B_s^{t,x})_{s\geq t}$ is a Brownian motion starting from x at time t.

A natural choice in generative modeling is a *confining potential*

$$V(x) = g\,\mathbf{1}_{x\notin\mathcal{D}},$$

which restricts the diffusion to a prescribed domain $\mathcal{D}$ as $g\to\infty$. In this limit, the optimal diffusion (36) takes the form

$$\begin{aligned}\mathrm{d}X_t &= \nabla \log \mathbb{E}\Big[\,\varphi(T, B_T^{t,X_t})\,\mathbf{1}_{\{B_\cdot^{t,X_t}\in\mathcal{D}\}}\,\Big]\,\mathrm{d}t + \mathrm{d}W_t \\ &= \nabla \log\left(\int_{\mathcal{D}} \varphi(T,y)\, p_{T-t}^{\mathcal{D}}(X_t,y)\,\mathrm{d}y\right)\mathrm{d}t + \mathrm{d}W_t,\end{aligned}$$

where $p_\tau^{\mathcal{D}}(x,y)$ denotes the heat kernel on $\mathcal{D}$ with absorbing (Dirichlet) boundary conditions.

H.1.1 COXETER DOMAINS AND DYSON-TYPE DIFFUSIONS

When $\mathcal{D}$ is a *Coxeter chamber* associated with a finite reflection group (Grabiner, 1999), the heat kernel $p_\tau^{\mathcal{D}}(x,y)$ can be computed explicitly from the *Coxeter reflection principle*, which generalizes the Karlin-McGregor determinant formula and expresses it as a signed sum of one-dimensional Gaussian kernels:

$$p_\tau(x,y) = \frac{1}{\sqrt{2\pi\tau}}\,\exp\Big(-\frac{(x-y)^2}{2\tau}\Big).$$

Let $R\subset\mathbb{R}^d$ be a root system with positive roots R_+, and W the corresponding finite Coxeter group* generated by reflections $\omega_\alpha(x) = x - \frac{2\langle\alpha,x\rangle}{\langle\alpha,\alpha\rangle}\alpha$. The corresponding Coxeter chamber is

$$\mathcal{D} = \{x\in\mathbb{R}^d :\ \langle\alpha,x\rangle > 0 \text{ for all } \alpha\in R_+\}.$$

For the classical Coxeter types A and C, the corresponding chambers are

$$\begin{aligned}\mathcal{D}^A &:= \{x\in\mathbb{R}^d : x_1 < x_2 < \cdots < x_d\},\\ \mathcal{D}^C &:= \{x\in\mathbb{R}^d : 0 < x_1 < x_2 < \cdots < x_d\}.\end{aligned}$$

These domains impose symmetry-breaking constraints relevant to generative modeling. For type A, the chamber $\mathcal{D}^A$ enforces a strict coordinate ordering, providing a ranked representation that removes permutation ambiguity in exchangeable systems (e.g., sorted eigenvalues or ordered latent features). For type C, the positivity constraint ($x_i > 0$) confines the process to the positive orthant, suitable for modeling nonnegative and ordered latent variables such as intensities or squared singular values. Thus, Coxeter chambers like $\mathcal{D}^A$ and $\mathcal{D}^C$ serve as geometric priors for generative diffusions on manifolds with permutation or reflection symmetries.

*A *Coxeter group* is a finite group W generated by reflections ω_α across hyperplanes orthogonal to a finite set of nonzero vectors $R\subset\mathbb{R}^d$, called the *root system*. The reflections satisfy the relations $(\omega_\alpha\omega_\beta)^{m_{\alpha\beta}} = 1$ for integers $m_{\alpha\beta}\geq 2$, which determine the group structure via a *Coxeter matrix*. The subset $R_+\subset R$ of *positive roots* selects one of the two connected components of $\mathbb{R}^d\setminus\bigcup_{\alpha\in R} H_\alpha$, where $H_\alpha = \{x : \langle\alpha,x\rangle = 0\}$ are the reflection hyperplanes. The corresponding *Coxeter (or fundamental) chamber* is $\mathcal{D} = \{x\in\mathbb{R}^d : \langle\alpha,x\rangle > 0,\ \forall\alpha\in R_+\}$. When the root system is crystallographic, W coincides with a *Weyl group* associated with a simple Lie algebra, but more general (noncrystallographic) Coxeter types such as H_3, H_4, or $I_2(m)$ also fit this framework.

From the Coxeter reflection principle,[†] we obtain the general Coxeter–reflection formula

$$p_\tau^{\mathcal{D}}(x,y) = \sum_{w\in W} \det(w)\, p_\tau\big(x, wy\big), \qquad x,y\in\mathcal{D},$$

where $\det(w)\in\{\pm 1\}$ is the sign of the orthogonal transformation w. For $\mu_0=\delta_{X_0}$, we get

$$\mathrm{d}X_t = \nabla \log \int_{\mathcal{D}} \mu_T(dy)\, \frac{p_{T-t}^{\mathcal{D}}(X_t,y)}{p_T^{\mathcal{D}}(X_0,X_T)}\,\mathrm{d}t + \mathrm{d}W_t.$$

Remark H.1 (Dyson Brownian motion (Type A), $\lambda=\infty$). For type A and the hard constraint $\lambda=\infty$, if $\mu_0=\mu_T=\delta_{X_0}$, the optimal process supported on $\mathcal{D}^A$ is the non-colliding Brownian bridge

$$\mathrm{d}X_t^i = \left(\frac{X_0^i - X_t^i}{T-t} + \sum_{j\neq i}\frac{1}{X_t^i - X_t^j}\right)\mathrm{d}t + \mathrm{d}W_t^i, \qquad i=1,\dots,d.$$

In the long-time limit $T\to\infty$, the bridge term vanishes, and we recover the celebrated Dyson Brownian motion

$$\mathrm{d}X_t^i = \sum_{j\neq i}\frac{1}{X_t^i - X_t^j}\,\mathrm{d}t + \mathrm{d}W_t^i, \qquad i=1,\dots,d,$$

which describes the eigenvalue dynamics of Hermitian Brownian ensembles (GUE). Thus, the linear Schrödinger bridge framework naturally connects to classical random-matrix diffusions.

Remark H.2 (Bru-Wishart diffusion (Type C), $\lambda=\infty$). Analogously, for type C with $\mu_0=\mu_T=\delta_{X_0}$, the optimal diffusion confined to $\mathcal{D}^C$ reads

$$\mathrm{d}X_t^i = \left(\frac{X_0^i - X_t^i}{T-t} + \frac{1}{X_t^i} + 2X_t^i\sum_{j\neq i}\frac{1}{(X_t^i)^2-(X_t^j)^2}\right)\mathrm{d}t + \mathrm{d}W_t^i, \qquad i=1,\dots,d.$$

As $T\to\infty$, the bridge term disappears, yielding the Bru-Wishart process, which governs the eigenvalue dynamics of Laguerre (Wishart) random matrices.

H.2 Score-based generative modeling as an exact linear-SB solution ($F\equiv 0$, noised-data prior)

Take $F\equiv 0$ (classical Schrödinger bridge, set $\mu_T=p_{\mathrm{data}}$, and choose the reference marginal as a noisy version of μ_T: $\mu_0 := \mu_T * \mathcal{N}(0, T\,I_d)$. The $(\varphi,\bar\varphi)$ system of Theorem 2.1 collapses to two linear (backward) heat equations:

$$\partial_t\varphi + \tfrac{1}{2}\Delta\varphi = 0, \qquad \partial_t\bar\varphi + \tfrac{1}{2}\Delta\bar\varphi = 0.$$

We claim the exact solution is $\bar\varphi\equiv 1$ and

$$\varphi(t,x) \;=\; \big(p_{\mathrm{data}} * \mathcal{N}(0,(T-t)I_d)\big)(x) \;=:\; p_{T-t}(x).$$

[†] **Proof sketch.** Let W be a finite Coxeter group acting orthogonally on $\mathbb{R}^d$, generated by reflections $\{\omega_\alpha\}$ across the hyperplanes $H_\alpha=\{x:\langle\alpha,x\rangle=0\}$, and let $\mathcal{D}$ be a fundamental chamber. We verify $p_\tau^{\mathcal{D}}$ is the fundamental solution of the heat kernel on $\mathcal{D}$ with Dirichlet boundary conditions:

(i) Heat equation. Each term $p_\tau(x,wy)$ satisfies $(\partial_\tau - \frac{1}{2}\Delta_x)p_\tau = 0$ in x, hence so does their finite signed sum by linearity.

(ii) Boundary condition (reflection principle). For any generating reflection ω_α (with $\det(\omega_\alpha)=-1$) and any x on the mirror H_α, orthogonality gives $p_\tau(\omega_\alpha x, wy) = p_\tau\big(x,\omega_\alpha wy\big)$. Using the change of variable $w\mapsto\omega_\alpha w$ in the sum,

$$p_\tau^{\mathcal{D}}(\omega_\alpha x,y) = \sum_{w\in W}\det(w)\,p_\tau(\omega_\alpha x, wy) = \sum_{w\in W}\det(\omega_\alpha w)\,p_\tau(x,wy) = -\,p_\tau^{\mathcal{D}}(x,y).$$

Thus $p_\tau^{\mathcal{D}}$ is *odd* across each reflection plane and vanishes on $\partial\mathcal{D}$.

(iii) Initial condition. As $\tau\downarrow 0$, $p_\tau(x,wy)\to\delta(x-wy)$. If $x,y\in\mathcal{D}$, the only element with $wy=y$ is $w=e$ (by the fundamental-chamber property), so the sum tends to $\delta(x-y)$.

Indeed: (i) both solve the heat equations trivially; (ii) the coupling $\varphi(0,x)\bar{\varphi}(T,x) = \mu_0(x)$ holds since $\varphi(0,\cdot) = p_T = \mu_T * \mathcal{N}(0, T\,I_d) = \mu_0$; (iii) the hard terminal constraint $m(T,x) \propto \varphi(T,x) = p_{\mathrm{data}}(x)$ recovers μ_T. Consequently the optimal drift of Theorem 2.1 reduces to

$$\nabla \log \varphi(t, X_t) \;=\; \nabla \log p_{T-t}(X_t),$$

the *score function* of the Gaussian-smoothed data at noise level $T - t$. Simulating $\mathrm{d}X_t = \nabla \log \varphi(t, X_t)\,\mathrm{d}t + \mathrm{d}W_t$, $X_0 \sim \mu_0$, coincides with the reverse-time score-based sampler of Song et al. (2021); Ho et al. (2020): DDPM / score-based diffusion is the classical-SB solution with noised-data prior and $\bar{\varphi} \equiv 1$.

H.3 Analytical semilinear bridge between Gaussians for $F(m) = \kappa \log m$ and $G(m(T,\cdot), x) = \lambda \log \frac{m(T,x)}{\mu(T,x)}$

Proposition H.1 (Optimal diffusion between Gaussians)**.** *Let the initial and terminal laws be Gaussian,*

$$\mu_0 = \mathcal{N}(m_0, \Sigma_0), \qquad \mu_T = \mathcal{N}(m_T, \Sigma_T),$$

with $\Sigma_0, \Sigma_T \in \mathbb{S}^d_{++}$ *and* $m_0, m_T \in \mathbb{R}^d$. *We take* $F(m) = \kappa \log m$ *and* $G(m(T,\cdot), x) = \lambda \log \frac{m(T,x)}{\mu(T,x)}$. *Then the* optimal diffusion process *is given by*

$$dX_t \;=\; -\,c\,(A_t X_t + b_t)\,dt \;+\; dW_t,$$

The parameters

$$A_t \in \mathbb{R}^{d\times d}, \quad b_t \in \mathbb{R}^d, \quad (\hat{\mu}_t \in \mathbb{R}^d, \quad \hat{\Sigma}_t \in \mathbb{S}^d_{++}),$$

evolve according to the coupled forward-backward system of ODEs:

$$\begin{aligned}
\dot{A}_t &= 2c\,A_t^2 + \kappa\,\hat{\Sigma}_t^{-1}, \quad A_T = \lambda(\Sigma_T^{-1} - \hat{\Sigma}_T^{-1}) \\
\dot{b}_t &= 2c\,A_t b_t - \kappa\,\hat{\Sigma}_t^{-1}\hat{\mu}_t, \quad b_T = \lambda(\hat{\Sigma}_T^{-1}\hat{\mu}_T - \Sigma_T^{-1} m_T) \\
\dot{\hat{\mu}}_t &= -\,c\,(A_t\hat{\mu}_t + b_t), \quad \hat{\mu}_0 = m_0 \\
\dot{\hat{\Sigma}}_t &= -\,c\,(A_t\hat{\Sigma}_t + \hat{\Sigma}_t A_t) + I_d, \quad \hat{\Sigma}_0 = \Sigma_0
\end{aligned}$$

The proof is provided below in Appendix I. It proceeds by inserting a Gaussian ansatz for $m(t,\cdot)$ and a quadratic ansatz for $v(t,\cdot)$ into the forward Fokker-Planck and backward HJB equations, and then matching coefficient terms in x.

Remark. When $\kappa = 0$, this system reduces to the *classical Schrödinger bridge* between Gaussian marginals, yielding the well-known result that the bridge is a linear Gaussian diffusion with quadratic potentials governed by coupled Riccati-type ODEs (Bunne et al., 2023).

H.4 Lifted (semilinear) Schrödinger bridge and Gyöngy's Markov projection

Gyöngy's Markov projection theorem (Gyöngy, 1986) has played a central role in mathematical finance. It provides a way to construct a Markov diffusion in $\mathbb{R}^d$ with prescribed marginal distributions at each time. This idea underlies the calibration of *local volatility* and *local stochastic volatility* models, as pioneered by Dupire's formula (Dupire, 1994) and systematically developed in the monograph (Guyon & Henry-Labordère, 2011).

In this subsection, we apply Gyöngy's projection to the setting of Schrödinger bridges. The key observation is that one may lift the marginals μ_0, μ_T to an augmented space $\mathbb{R}^{d+1}$ (eventually higher), solve the Schrödinger bridge there, and then project back onto $\mathbb{R}^d$. Gyöngy's theorem then guarantees the existence of a Markov diffusion in $\mathbb{R}^d$ with the correct marginals μ_0, μ_T. Although this projected diffusion has the right one-time marginals, it is in general *not* itself a Schrödinger bridge. It therefore defines a new class of generative models, obtained by combining lifted SB with Markov projection.

$$\begin{array}{ccc}
\tilde{\mu}_0 & \xrightarrow{\text{SB in } \mathbb{R}^{d+1}} & \tilde{\mu}_T \\
{\scriptstyle projection}\big\downarrow & & \big\downarrow{\scriptstyle projection} \\
\mu_0 & \xrightarrow[\text{Markov proj.}]{} & \mu_T
\end{array}$$

Proposition H.2 (Lifted Schrödinger bridge and Gyöngy projection)**.** *Let $\mu_0, \mu_T \in \mathcal{P}(\mathbb{R}^d)$ be given marginals, and let $\tilde{\mu}_0, \tilde{\mu}_T \in \mathcal{P}(\mathbb{R}^{d+1})$ be their lifted versions via conditional distributions ν_0, ν_T:*

$$\tilde{\mu}_0(dx, dy) = \mu_0(dx)\,\nu_0(dy|x), \qquad \tilde{\mu}_T(dx, dy) = \mu_T(dx)\,\nu_T(dy|x).$$

Consider the Schrödinger bridge in $\mathbb{R}^{d+1}$ between $\tilde{\mu}_0$ and $\tilde{\mu}_T$ under the Wiener reference $\mathbb{P}^0$, with dynamics

$$d\tilde{X}_t = \nabla_{(x,y)} \log \varphi(t, \tilde{X}_t)\, dt + d\tilde{W}_t, \qquad \tilde{X}_t = (X_t, Y_t).$$

We define the projection $X_t \in \mathbb{R}^d$ by

$$dX_t \;=\; \frac{\displaystyle\int \nabla_x \varphi(t, X_t, y)\, \bar{\varphi}(T-t, X_t, y)\, \mathrm{d}y}{\displaystyle\int \varphi(t, X_t, y)\, \bar{\varphi}(T-t, X_t, y)\, \mathrm{d}y}\, \mathrm{d}t + \mathrm{d}W_t, \tag{37}$$

where $(\varphi, \bar{\varphi})$ are the Schrödinger potentials in $\mathbb{R}^{d+1}$. We have $\mathrm{Law}(X_0) = \mu_0$ and $\mathrm{Law}(X_T) = \mu_T$.

The proof is deferred to Appendix J.

I Proof of Proposition H.1

Proof. We proceed by a Gaussian-quadratic ansatz. Namely, we assume

$$m(t, \cdot) = \mathcal{N}\big(\hat{\mu}_t, \hat{\Sigma}_t\big), \qquad v(t, x) = \tfrac{1}{2} x^\top A_t x + b_t^\top x + a_t,$$

with time-dependent parameters $(A_t, b_t, a_t, \hat{\mu}_t, \hat{\Sigma}_t)$ to be determined.

Substituting these expressions into the Fokker-Planck equation (23) and into the Hamilton-Jacobi-Bellman equation (18), and matching the coefficients of the quadratic, linear and constant terms in x, one obtains after straightforward calculations the closed system of ODEs for $(A_t, b_t, \mu_t, \Sigma_t)$ stated in the proposition (Equation for a_t is not needed). We then use the terminal condition $v(T, x) = G(m(T, \cdot), x)$. We conclude by uniqueness of the forward-backward system.

□

J Proof of Proposition H.2

Proof. The drift of the projected diffusion is obtained by applying Gyöngy's mimicking theorem (Gyöngy, 1986) to the x-component of the lifted bridge. By definition,

$$b(t, x) = \mathbb{E}^{\tilde{\mathbb{P}}}\left[\nabla_x \log \varphi(t, X_t, Y_t) \,\middle|\, X_t = x\right].$$

then we use that $m(t, X_t, Y_t) \propto \varphi(t, X_t, Y_t)\bar{\varphi}(T-t, X_t, Y_t)$. This yields the ratio representation for $b(t, x)$, and thus the stated SDE. □

K Numerical sanity check: branching with time-reflected coupling

Marked branching for systems of coupled semilinear equations is established (Henry-Labordère, 2012; Henry-Labordère et al., 2019; Henry-Labordère & Touzi, 2021); the bridge system (3)-(4) is special in that the nonlinearity pairs $\varphi(t, \cdot)$ with $\bar{\varphi}(T-t, \cdot)$ (see the main text diagram). The scalar model below isolates $u(t, x)\, u(T-t, x)$ as an analytic benchmark for that $T-t$ product.

K.1 Sanity check

We illustrate the branching diffusion method on a simple one-dimensional semilinear PDE for a single potential u:

$$\begin{cases} \partial_t u(t, x) + \frac{1}{2}\, \partial_{xx} u(t, x) + u(t, x)\, u(T-t, x) \;=\; 0, & (t, x) \in \mathbb{R} \times \mathbb{R}, \\ u(T, x) \;=\; 3\, \mathrm{sech}^2(x) - 1 - \tanh\Big(\dfrac{T}{2}\Big). \end{cases} \tag{38}$$

The explicit solution is given by

$$u(t,x) \;=\; 3\,\mathrm{sech}^2(x) - 1 - \tanh\Big(t - \frac{T}{2}\Big),$$

so that $u(0,\cdot)$ is known in closed form. This example is used only as a numerical sanity check of the branching implementation. In this toy case, we work with a single family of particles (say "black" particles, representing u). We rewrite the nonlinearity $u(t,x)u(T-t,x)$ in (38) using the *canonical* form introduced in (10):

$$\partial_t u(t,x) + \frac{1}{2}\,\partial_{xx} u(t,x) + 1 \times \big(\tfrac{1}{2} \times 2u(t,x)\,u(T-t,x) + \tfrac{1}{2} \times 2u(t,x) - u(t,x)\big) \;=\; 0.$$

Figure 8 compares the analytical solution $u(0,x)$ as a function of x with the Monte Carlo branching algorithm for $T = 1$. The branching estimator closely matches the analytical curve, confirming the correctness of the branching scheme. To assess the impact of the nonlinear term, we also include the solution of the heat kernel with terminal condition $u(T,x)$.

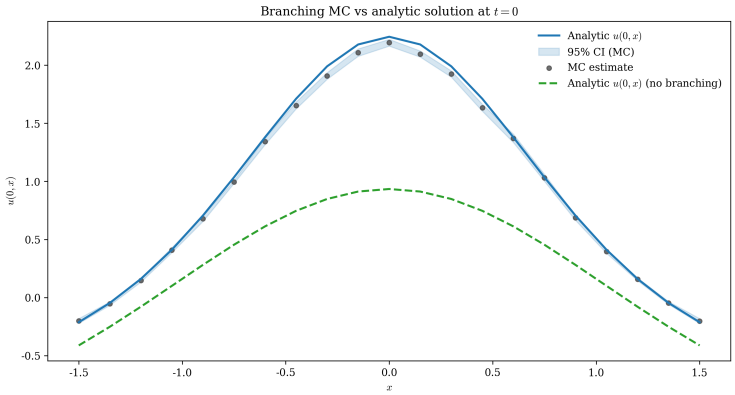

Figure 8: Comparison of the analytic solution $u(0,x)$ with the Monte Carlo estimate $\hat{u}_M(0,x)$, in the branching case and the heat kernel, highlighting the impact of the nonlinearity.

L Proof of Proposition 3.1

Proof. **Step 1. L^1-integrability.** Define $\mathcal{W}_\infty$ as the worst-case analogue of $\mathcal{W}$, obtained by replacing each branching factor $a_{K^\star}(s, X_s)/p_{K^\star}$ by $\|a_{K^\star}\|_{L^\infty([0,T]\times\mathbb{R}^d)}/p_{K^\star}$: it starts at 1 and at each branching event updates as

$$\mathcal{W}_\infty \longrightarrow \mathcal{W}_\infty \times \frac{\|a_{K^\star}\|_{L^\infty([0,T]\times\mathbb{R}^d)}}{p_{K^\star}}.$$

Then

$$\mathbb{E}\left[\left|\mathcal{W}\prod_{i=1}^{N_T^{\text{white}}}\varphi(T,X_T^i)\prod_{j=1}^{\bar{N}_T^{\text{black}}}\bar{\varphi}(T,\bar{X}_T^j)\right|\right] \le \mathbb{E}\Big[\mathcal{W}_\infty \|\varphi(T,\cdot)\|_\infty^{N_T^{\text{white}}}\,\|\bar{\varphi}(T,\cdot)\|_\infty^{\bar{N}_T^{\text{black}}}\Big] =: \varphi_\infty(0).$$

Define, for $r \ge 0$,

$$\|F\|_\infty^{(1)}(r) := 1 - \sum_{k=0}^{K} \|a_k\|_{L^\infty([0,T]\times\mathbb{R}^d)}\, r^k.$$

Taking a Feynman-Kac expectation over the branching tree, $\varphi_\infty(t)$ satisfies the coupled ODE system

$$\begin{cases} \dfrac{d\varphi_\infty(t)}{dt} = c\,\varphi_\infty(t)\,\|F\|_\infty^{(1)}(\varphi_\infty(t)\bar{\varphi}_\infty(T-t)), \quad \varphi_\infty(T) = \|\varphi(T,\cdot)\|_\infty, \\ \dfrac{d\bar{\varphi}_\infty(t)}{dt} = c\,\bar{\varphi}_\infty(t)\,\|F\|_\infty^{(1)}(\varphi_\infty(T-t)\bar{\varphi}_\infty(t)), \quad \bar{\varphi}_\infty(T) = \|\bar{\varphi}(T,\cdot)\|_\infty. \end{cases}$$

Setting $\hat{\varphi}_\infty(t) := \bar{\varphi}_\infty(T-t)$ decouples the time-reflection and yields

$$\begin{cases} \dfrac{d\varphi_\infty(t)}{dt} = c\,\varphi_\infty(t)\,\|F\|_\infty^{(1)}(\varphi_\infty(t)\hat{\varphi}_\infty(t)), & \varphi_\infty(T) = \|\varphi(T,\cdot)\|_\infty, \\ \dfrac{d\hat{\varphi}_\infty(t)}{dt} = -\,c\,\hat{\varphi}_\infty(t)\,\|F\|_\infty^{(1)}(\varphi_\infty(t)\hat{\varphi}_\infty(t)), & \hat{\varphi}_\infty(0) = \|\hat{\varphi}(0,\cdot)\|_\infty, \end{cases}$$

which in terms of $m_\infty := \varphi_\infty\hat{\varphi}_\infty$ and $v_\infty := -\frac{1}{c}\log\varphi_\infty$ becomes the trivial system $\mathrm{d}m_\infty/\mathrm{d}t = 0$, $\mathrm{d}v_\infty/\mathrm{d}t + \|F\|_\infty^{(1)}(m_\infty) = 0$. This admits a global-in-time bounded solution on $[0,T]$, proving existence of $\varphi_\infty, \bar{\varphi}_\infty$ and hence $\Phi \in L^1(\mathbb{P}^0)$.

Step 2. L^2-integrability. Similarly,

$$\mathbb{E}\left[\left|\mathcal{W}\prod_{i=1}^{N_T^{\mathrm{white}}}\varphi(T,X_T^i)\prod_{j=1}^{\bar{N}_T^{\mathrm{black}}}\bar{\varphi}(T,\bar{X}_T^j)\right|^2\right] \le \mathbb{E}\Big[\mathcal{W}_\infty^2\|\varphi(T,\cdot)\|_\infty^{2N_T^{\mathrm{white}}}\|\bar{\varphi}(T,\cdot)\|_\infty^{2\bar{N}_T^{\mathrm{black}}}\Big] =: \varphi_{\infty,2}(0).$$

At each branching event,

$$\mathcal{W}_\infty^2 \longrightarrow \mathcal{W}_\infty^2 \times \frac{\|a_{K^\star}\|^2_{L^\infty([0,T]\times\mathbb{R}^d)}}{p_{K^\star}^2}.$$

Define, for $r \ge 0$,

$$\|F\|_\infty^{(2)}(r) := 1 - \sum_{k=0}^{K}\frac{\|a_k\|^2_{L^\infty([0,T]\times\mathbb{R}^d)}}{p_k}\,r^k.$$

This leads to

$$\begin{cases} \dfrac{d\varphi_{\infty,2}(t)}{dt} = c\,\varphi_{\infty,2}\,\|F\|_\infty^{(2)}\big(\varphi_{\infty,2}(t)\bar{\varphi}_{\infty,2}(T-t)\big), & \varphi_{\infty,2}(T) = \|\varphi(T,\cdot)\|_\infty^2, \\ \dfrac{d\bar{\varphi}_{\infty,2}(t)}{dt} = c\,\bar{\varphi}_{\infty,2}\,\|F\|_\infty^{(2)}\big(\varphi_{\infty,2}(t)\bar{\varphi}_{\infty,2}(T-t)\big), & \bar{\varphi}_{\infty,2}(T) = \|\bar{\varphi}(T,\cdot)\|_\infty^2. \end{cases}$$

Integrating as before yields existence, hence $\Phi \in L^2(\mathbb{P}^0)$.

Step 3. Representation formula. Finally, by the strong Markov property and particle independence, together with $\Phi \in L^1$, we obtain

$$\varphi(t,x) = \mathbb{E}_{t,x}\big[\Phi\big],$$

which solves the semilinear system (3)–(4) (see (Henry-Labordère & Touzi, 2021) for details). Uniqueness of the mild solution concludes the proof. □

M FEYNMAN-GALTON-WATSON DIAGRAMS

Each realization of the branching system is a Galton-Watson tree: lines correspond to Brownian trajectories between branching events, and vertices represent branchings. This is directly analogous to tree Feynman diagrams: dashed lines depict white particles (φ), solid lines depict black particles ($\bar{\varphi}$), and trivalent vertices encode the cubic nonlinearity. For the real Gross-Pitaevskii case, $F(u) = u$, the only admissible branching is $k = 1$:

$$\circ \;\to\; 2\circ + 1\bullet, \qquad \bullet \;\to\; 1\circ + 2\bullet,$$

so each vertex corresponds to one incoming line splitting into three outgoing lines. Iterating these vertices yields a perturbative expansion in the coupling constant c: a diagram with n branchings contributes a weight of order c^n.

Importantly, if we truncate the Galton-Watson expansion to order 0 (no branching), the dynamics are given only by free Brownian propagation. In this case, the system reduces to the *classical*

Schrödinger bridge, and the PDEs collapse to the linear forward-backward heat equations. Higher-order branchings then represent nonlinear corrections beyond this baseline. Appendix K gives a scalar toy sanity check of the branching scheme for the $T-t$ coupling in (3)-(4).

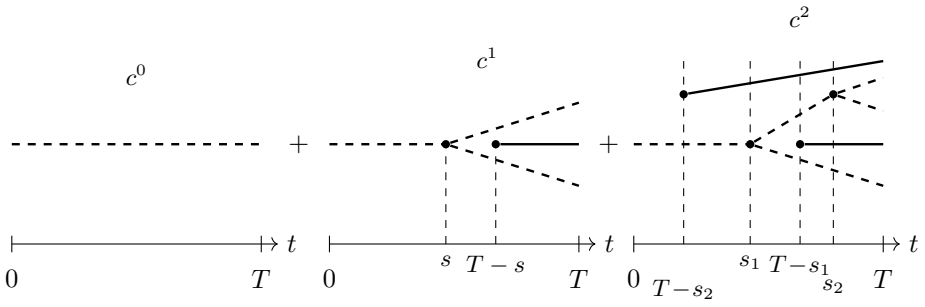

Figure 9: First three contributions in the perturbative expansion (Gross-Pitaevskii cubic branching). Dashed lines represent white particles (φ), solid lines represent black particles ($\bar{\varphi}$). Weights with $a_1 = p_1 = 1$: order c^0, no branching, i.e. classical Schrödinger bridge / linear heat kernel; order c^1 with one branching at s (white offspring starting at s, black offspring at $T - s$); order c^2 with two branchings at s_1, s_2 : offspring of the same type as the parent start at s_i, while offspring of the other type start at $T - s_i$. The horizontal time axes (all at the same vertical level) show time from 0 to T, with dashed vertical lines indicating the splitting times.

N Implementation details

Training and sampling both rely on stochastic quantities built from branching trees and on differentiating $\log \widehat{\varphi}(t, x)$ with respect to x (Algorithm 2). When coordinates have very different scales or heavy tails, (i) products of terminal factors in the branching estimator can span many orders of magnitude, inflating Monte Carlo variance, and (ii) autograd through $\log \varphi_\theta$ and the reweighted score becomes ill-conditioned relative to the isotropic unit-variance reference noise in the SDE. We therefore normalize inputs before bridge estimation, using one of the two mechanisms below depending on dimensionality and data modality.

N.1 Covariance whitening in data space

Let $X_0 \sim \mu_0$ and $X_T \sim \mu_T$. We define

$$\hat{X}_0 := C_0^{-\frac{1}{2}}\big(X_0 - \mathbb{E}^{\mu_0}[X_0]\big), \quad C_0 := \mathrm{Cov}(\mu_0),$$

$$\hat{X}_T := \sigma\, C_T^{-\frac{1}{2}}\big(X_T - \mathbb{E}^{\mu_T}[X_T]\big), \quad C_T := \mathrm{Cov}(\mu_T), \quad \sigma^2 = 2,$$

train the bridge on $\hat{\mu}_0 := \mathrm{Law}(\hat{X}_0)$ and $\hat{\mu}_T := \mathrm{Law}(\hat{X}_T)$, and map generated samples back to the original coordinates by the inverse affine map. If μ_0 and μ_T are Gaussian, then $\hat{\mu}_0 \sim \mathcal{N}(0, I)$, $\hat{\mu}_T \sim \mathcal{N}(0, \sigma^2 I)$ and $\hat{\mu}_T = \hat{\mu}_0 \star \mathcal{N}(0, I)$, which matches the classical Schrödinger bridge with $(\varphi_{\mathrm{NN}}, \bar{\varphi}_{\mathrm{NN}}) = (0, 0)$ at horizon $T = 1$. In practice we set $T = 2$ to retain nontrivial interpolation; when μ_T is close to Gaussian, initializing $(\varphi_{\mathrm{NN}}, \bar{\varphi}_{\mathrm{NN}}) = (0, 0)$ is already near optimal with little additional training.

N.2 Latent representations for imaging

For high-dimensional images, empirical covariances are singular or badly conditioned (e.g. near-constant background pixels), so $C_T^{-1/2}$ is unstable and pixel-space whitening is inappropriate. We instead encode images with a pretrained VAE (MNIST) or ALAE (FFHQ), yielding low-dimensional codes that are approximately standardized and better aligned with the isotropic diffusion reference; the bridge is trained in that latent space and decoded for visualization.

O RESULTS DETAILS

O.1 2D SYNTHETIC EXPERIMENTS

We adopt synthetic two-dimensional distributions commonly used in the literature (De Bortoli et al., 2021):

- $\mu_0 :=$ **Gaussian** $\to \mu_T :=$ **Mixture of Gaussians.** The source μ_0 is a centered Gaussian with variance $\sigma^2 I$ ($\sigma = 5$), and the target μ_T is a mixture of 8 Gaussians.
- $\mu_0 :=$ **Moon curves** $\to \mu_T :=$ **Swiss curves.** The source μ_0 is a composition of two Moon-shaped distributions, and the target μ_T is a Swiss curve.

The recovered Schrödinger bridges for these two cases are shown in Figure 10. The temporal evolution of the interpolating flow illustrates how the source progressively deforms into the target (Figures 2).

Training dynamics are summarized in Figure 27, see Appendix P.1. As the number of epochs increases, the generated terminal distribution $\mathrm{Law}(X_T)$ gradually converges toward the target law μ_T (Figure 27). At the same time, the potentials $(\varphi_\theta, \bar{\varphi}_\eta)$ sharpen as training progresses, reflecting improved discrimination between source and target (Figures 28 and 29, see Appendix P.1).

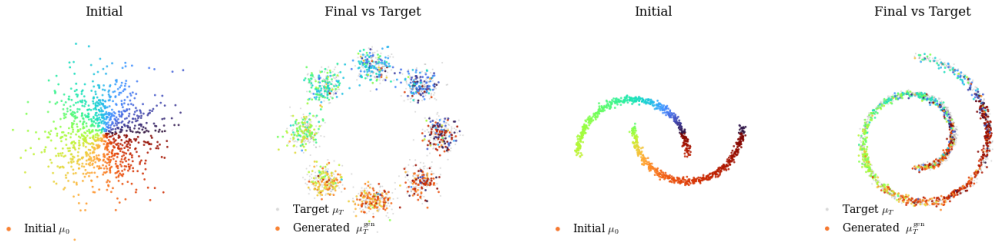

(10a) Gaussian $\to$ mixture of Gaussians.

(10b) Moon curves $\to$ Swiss curve.

Figure 10: Classical Schrödinger bridges in $\mathbb{R}^2$.

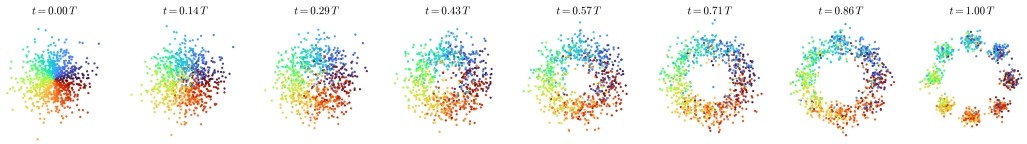

Figure 11: Extended view of the temporal evolution of $\mathrm{Law}(X_t)$ (Moons $\to$ Swiss).

O.2 COMPARISON ON THE 2D BENCHMARK

We compare SBNL with LightSB (Korotin et al., 2023), LightSBB-M (Alouadi et al., 2026), DSBM (Shi et al., 2023) and BranchSBM (Tang et al., 2026) on the Gaussian $\to$ mixture of 8 Gaussians benchmark from Appendix O. We report (i) terminal samples for each method (Figure 12), (ii) the closed-form mixture W_2 distance between the generated terminal law and μ_T (Figure 14a), and (iii) the per-mode coverage of the 8 components (Figure 15).

O.2.1 COMPUTATION OF THE COMPARISON METRICS

Let $\hat{\mu}_T$ denote the empirical distribution of $N = 10^4$ generated terminal samples and μ_T the target mixture of 8 isotropic Gaussians in $\mathbb{R}^2$.

Mixture Wasserstein MW_2. We first fit a Gaussian mixture with $K = 8$ components to the generated samples by expectation–maximization (sklearn's `GaussianMixture`, full covariances,

restarted 5 times, best log-likelihood kept), obtaining $\hat{p} = \sum_{k=1}^{K} \hat{w}_k \, \mathcal{N}(\hat{m}_k, \hat{\Sigma}_k)$. We compare it to the known target mixture $p^\star = \sum_{k=1}^{K} w_k^\star \, \mathcal{N}(m_k^\star, \Sigma_k^\star)$ using the *mixture Wasserstein* distance (Delon & Desolneux, 2020)

$$MW_2^2(\hat{p}, p^\star) \;=\; \min_{\pi \in \Pi(\hat{w}, w^\star)} \; \sum_{k,\ell} \pi_{k\ell} \, W_2^2\big(\mathcal{N}(\hat{m}_k, \hat{\Sigma}_k), \, \mathcal{N}(m_\ell^\star, \Sigma_\ell^\star)\big),$$

where the W_2 distance between Gaussians has the closed form $W_2^2(\mathcal{N}(m_1, \Sigma_1), \mathcal{N}(m_2, \Sigma_2)) = \|m_1 - m_2\|^2 + \mathrm{Tr}\big(\Sigma_1 + \Sigma_2 - 2(\Sigma_1^{1/2}\Sigma_2\Sigma_1^{1/2})^{1/2}\big)$, so MW_2 reduces to a small linear program on the $K \times K$ cost matrix and is computed in closed form.

Sliced Wasserstein SW_2. We use the standard Monte Carlo estimator of the sliced 2-Wasserstein distance,

$$SW_2^2(\hat{\mu}_T, \mu_T) \;=\; \frac{1}{L} \sum_{\ell=1}^{L} W_2^2\big(P_{\theta_\ell \sharp} \hat{\mu}_T, \; P_{\theta_\ell \sharp} \mu_T\big),$$

where θ_ℓ are $L = 200$ directions drawn uniformly on $\mathbb{S}^1$ and $P_\theta x = \theta^\top x$; each 1D W_2^2 is evaluated exactly on the sorted projected samples (with 10^4 samples for both $\hat{\mu}_T$ and μ_T).

Empirical W_2^{emp}. We solve the discrete optimal transport between $\hat{\mu}_T$ and an i.i.d. sample of size 10^4 from μ_T using the POT library (exact LP solver on the squared Euclidean cost), reporting $W_2^{\text{emp}} = (\sum_{ij} \pi_{ij}^\star \|x_i - y_j\|^2)^{1/2}$.

Energy distance E. We use the kernel-based energy distance (Székely & Rizzo, 2004), estimated by U-statistics on the same 10^4 samples from $\hat{\mu}_T$ and μ_T:

$$E^2(\hat{\mu}_T, \mu_T) \;=\; 2\,\mathbb{E}\|X - Y\| - \mathbb{E}\|X - X'\| - \mathbb{E}\|Y - Y'\|, \quad X, X' \sim \hat{\mu}_T, \; Y, Y' \sim \mu_T.$$

$E \geq 0$, with $E = 0$ iff the two distributions coincide.

Moment and coverage diagnostics. The mean error $|\Delta\mu| = \|\mathbb{E}_{\hat{\mu}_T}[X] - \mathbb{E}_{\mu_T}[X]\|$ and the covariance error $\|\Delta\Sigma\|_F = \|\mathrm{Cov}(\hat{\mu}_T) - \mathrm{Cov}(\mu_T)\|_F$ are computed from empirical moments. The per-mode coverage assigns each generated sample to its nearest component center in μ_T and reports the resulting empirical frequencies (Figure 15); the uniform reference is $1/K = 1/8$.

Confidence intervals. Each metric is a random variable of the generated sample, so every value reported in Table 3b and Figures 14a-14b is accompanied by a bootstrap confidence interval. For every method and every metric we draw $B = 5$ resamples with replacement of both the prediction tensor (size $N = 10^4$) and a freshly drawn reference sample from μ_T, then recompute the entire metric pipeline on each replica-including refitting the 8-component GMM with a fresh seed for MW_2, so that variability from the GMM fit is also captured. The reported point estimate is the mean across replicas and the $\pm$ value is the sample standard deviation (with Bessel correction). For the empirical Wasserstein distance, the $10^4 \times 10^4$ EMD is too heavy inside the bootstrap loop, so we down-sample to 2000 points per side in each replica (hence the W_2^{emp} label in Figure 14b). Per-mode coverage error bars in Figure 15 are obtained the same way. Each method is evaluated from a single trained checkpoint, so the reported intervals reflect *sample/estimator* uncertainty rather than training-seed variability.

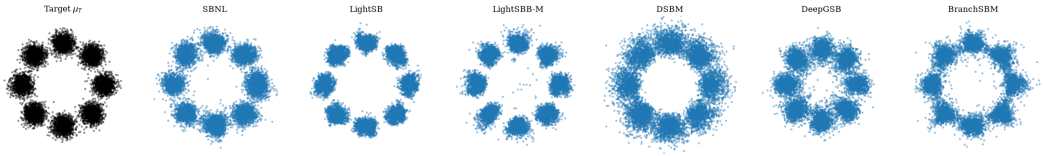

Figure 12: Terminal samples on Gaussian $\to$ mixture of 8 Gaussians for each method.

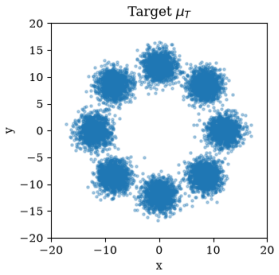

(13a) Target μ_T

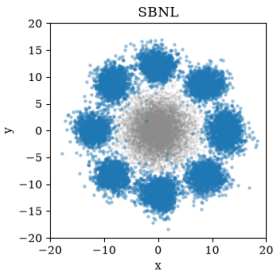

(13b) SBNL

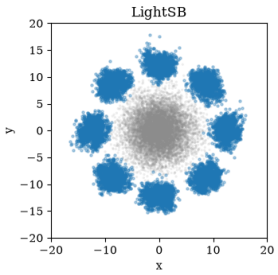

(13c) LightSB

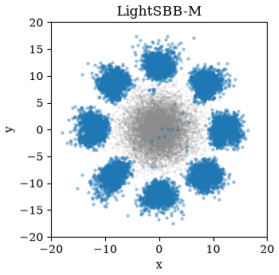

(13d) LightSBB-M

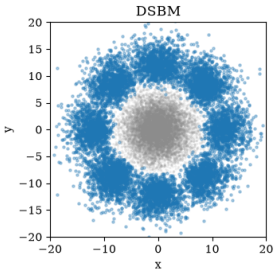

(13e) DSBM

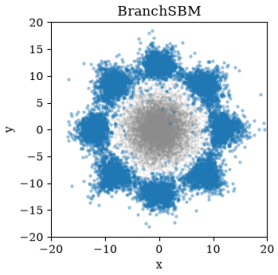

(13f) BranchSBM

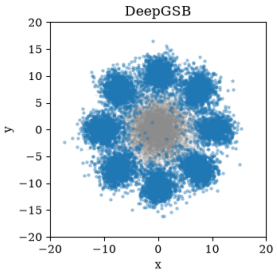

(13g) DeepGSB

Figure 13: Terminal samples per method on the Gaussian → mixture of 8 Gaussians benchmark.

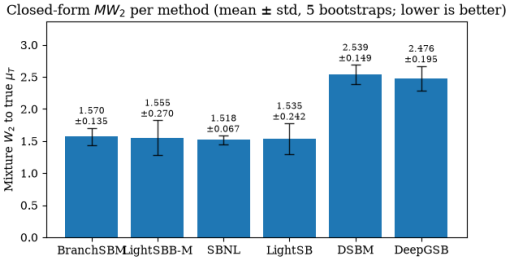

(14a) Closed-form mixture W_2 between generated terminal law and μ_T.

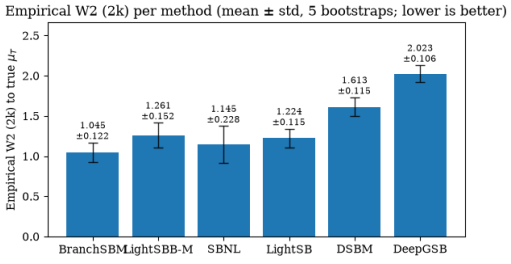

(14b) Empirical W_2 between generated terminal law and μ_T.

Figure 14: Per-method comparison on the Gaussian → mixture of 8 Gaussians benchmark (lower is better; mean ± std over 5 bootstraps).

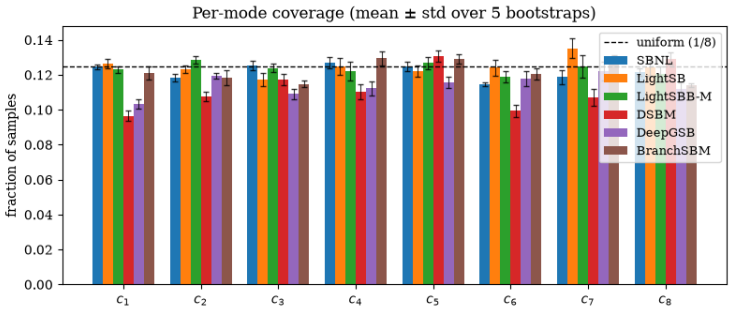

Figure 15: Per-mode coverage of the 8 target components. The dashed line marks the uniform reference $1/8$.

O.3 Potential running cost

We now consider a *state-dependent* running cost of potential type,

$$f\big(m(t,\cdot)\big) = \int_{\mathbb{R}^d} V(x)\, m(t,x)\,\mathrm{d}x \qquad \Longrightarrow \qquad F(t,x,\varphi,\bar{\varphi}) = V(x),$$

so that the canonical PDE (3) reads $\partial_t\varphi + \frac{1}{2}\Delta\varphi = c\,\varphi\,\big(1-(1-V(x))\big)$. The interaction is thus fully specified by a single **callable** $V:\mathbb{R}^d\to\mathbb{R}$, passed to the model and evaluated on the fly during both training and sampling to encode the environment.

Branching mechanism. Matching the canonical form $F = 1 - F_p$ with $F_p(t,x,m) = 1 - V(x)$ gives a single degree-zero coefficient $a_0(x) = 1 - V(x)$ (and $a_k = 0$ for $k \geq 1$). With $p_0 = 1$, the only branching outcome is that a particle continues as a single copy of the *same* type, and each branching event reweights the tree by

$$\mathcal{W} \;\longrightarrow\; \mathcal{W}\cdot\frac{a_0(X_s)}{p_0} = \mathcal{W}\cdot\big(1-V(X_s)\big).$$

There are no offspring and no type flips: the branching estimator of Proposition 3.1 reduces to a pure Feynman-Kac weighting of Brownian paths (Appendix H, confining potentials), with V bounded so that $1 - V > 0$ keeps the weights positive.

Training and sampling. Because F is state- but not density-dependent, the system (3)-(4) is linear and $\varphi(0,x)$ is estimated by the reweighted trees above; the potentials are trained with the MMD objective (14). Two practical ingredients make the importance-sampled estimator stable. First, the interior weights $W = \widehat{A}\,\bar{\varphi}_\eta/\mu^*$ (resp. $\widehat{B}\,\varphi_\theta/\mu^*$) are heavy-tailed when μ^* is too narrow, so we take a *wide* Gaussian reference μ^* covering the region actually explored by the branching diffusion. Second, we use *self-normalized* importance sampling (dividing each weight vector by its own sum, gradient preserved), which turns the interior expectations into low-variance ratio estimators comparing the *normalized* residual measures. At sampling time, the drift $\nabla\log\varphi$ is estimated by the same weighted trees (Algorithm 2), so V deflects trajectories away from high-V regions.

Endpoint offset trick. When V has an asymmetric geometry (e.g. gates at different heights, a curved valley), we place μ_0 and μ_T at *different* anchor points aligned with the entrance/exit of the low-energy corridor, rather than symmetrically. This gives the score a clear valley to follow through the obstacle. Transporting mass over a distance $\sim\ell$ requires the drift to integrate to $\sim\ell$, so we also use a longer horizon T (with a small rate c): too small a T leaves the drift no time to route the mass around the barrier.

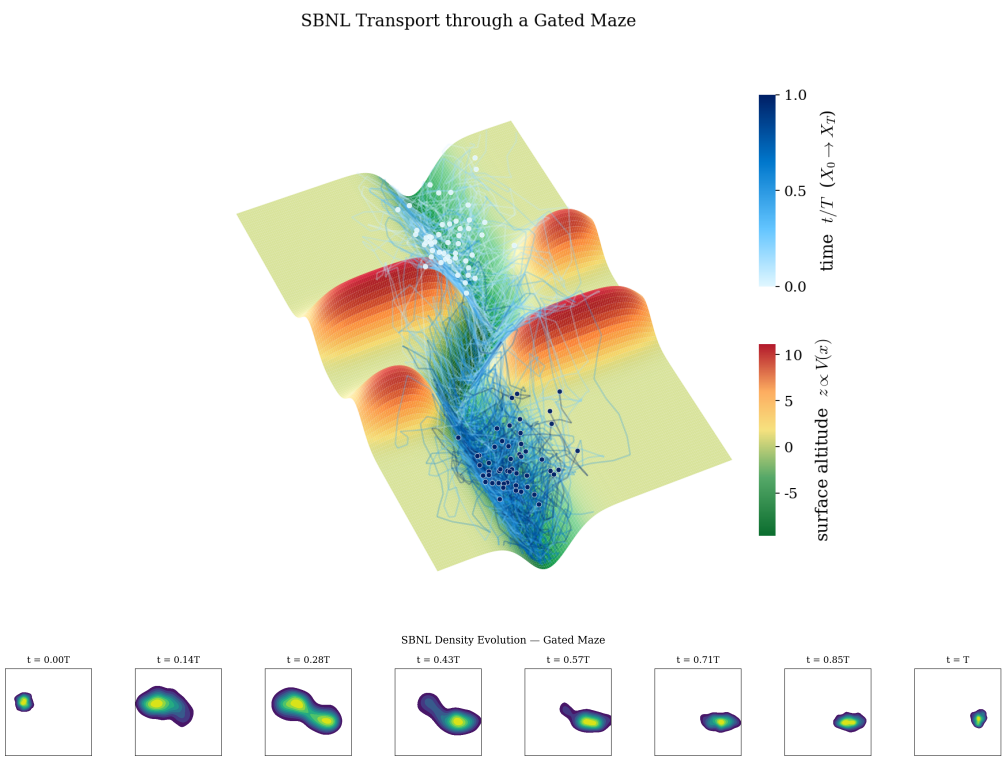

Figure 16: Gated maze: SBNL transport and estimated density evolution $m(t,\cdot) = \varphi(t,\cdot)\bar{\varphi}(T-t,\cdot)$.

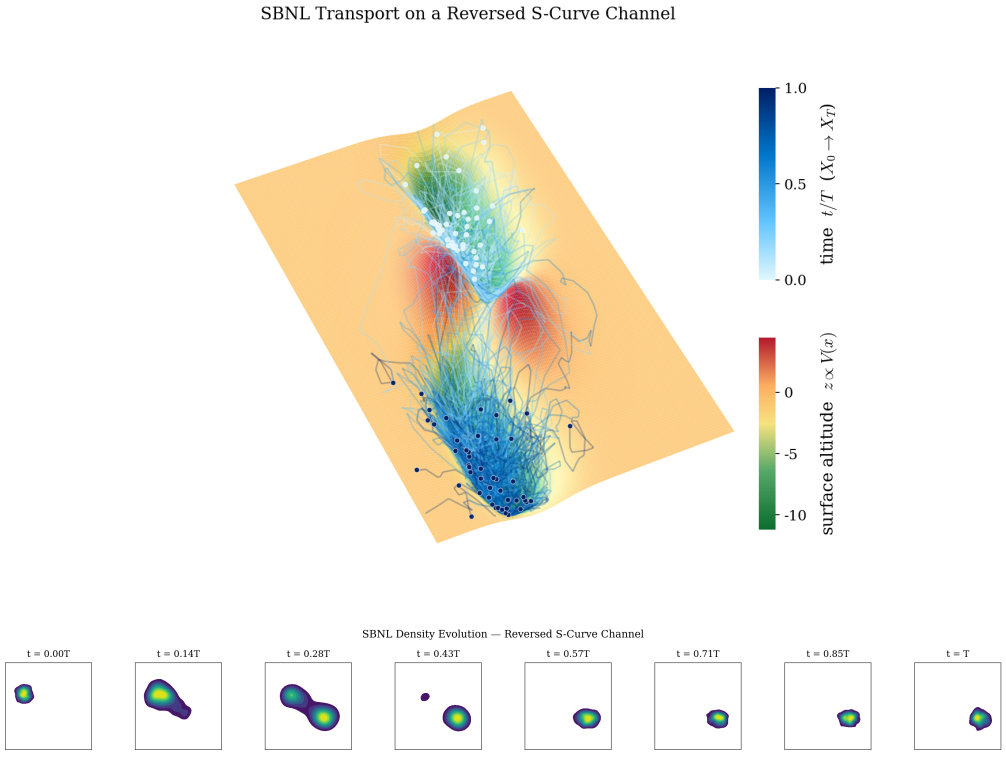

Figure 17: S-curve: SBNL transport and estimated density evolution $m(t,\cdot) = \varphi(t,\cdot)\bar{\varphi}(T-t,\cdot)$.

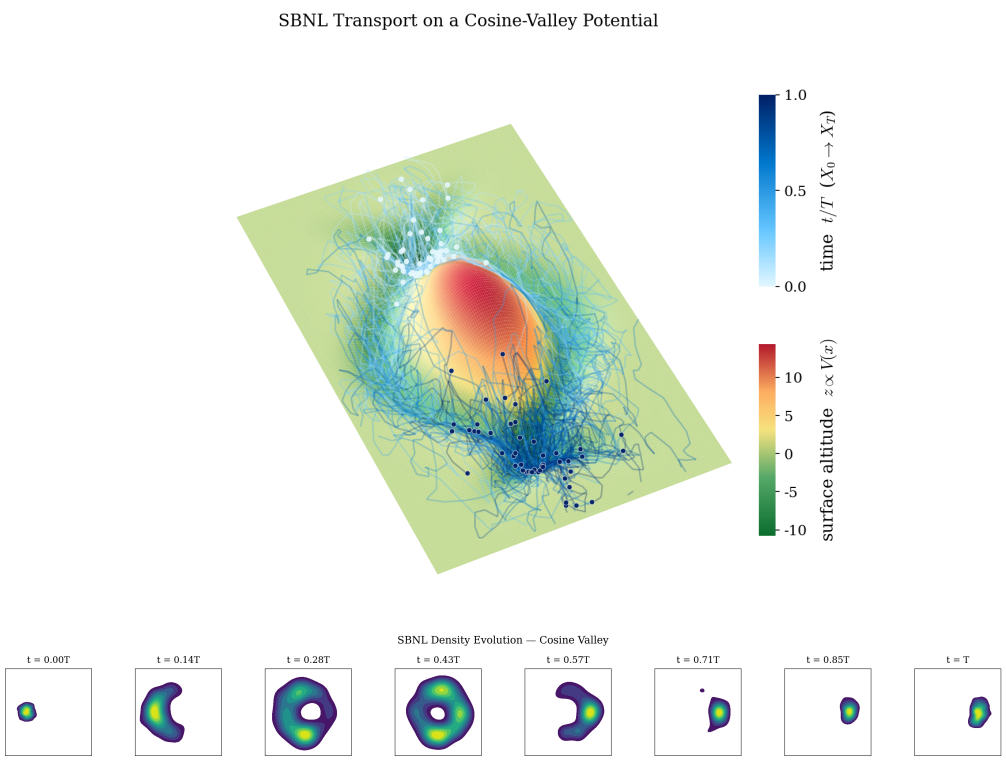

Figure 18: Valley: SBNL transport and estimated density evolution $m(t,\cdot) = \varphi(t,\cdot)\bar{\varphi}(T-t,\cdot)$.

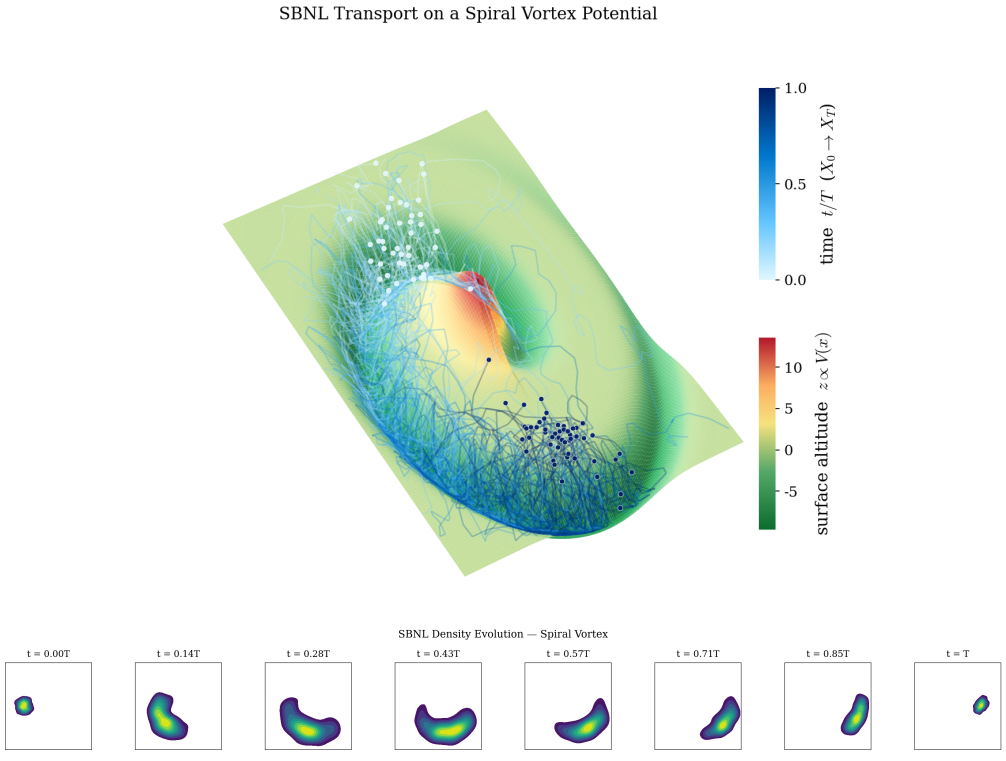

Figure 19: Spiral: SBNL transport and estimated density evolution $m(t,\cdot) = \varphi(t,\cdot)\bar{\varphi}(T-t,\cdot)$.

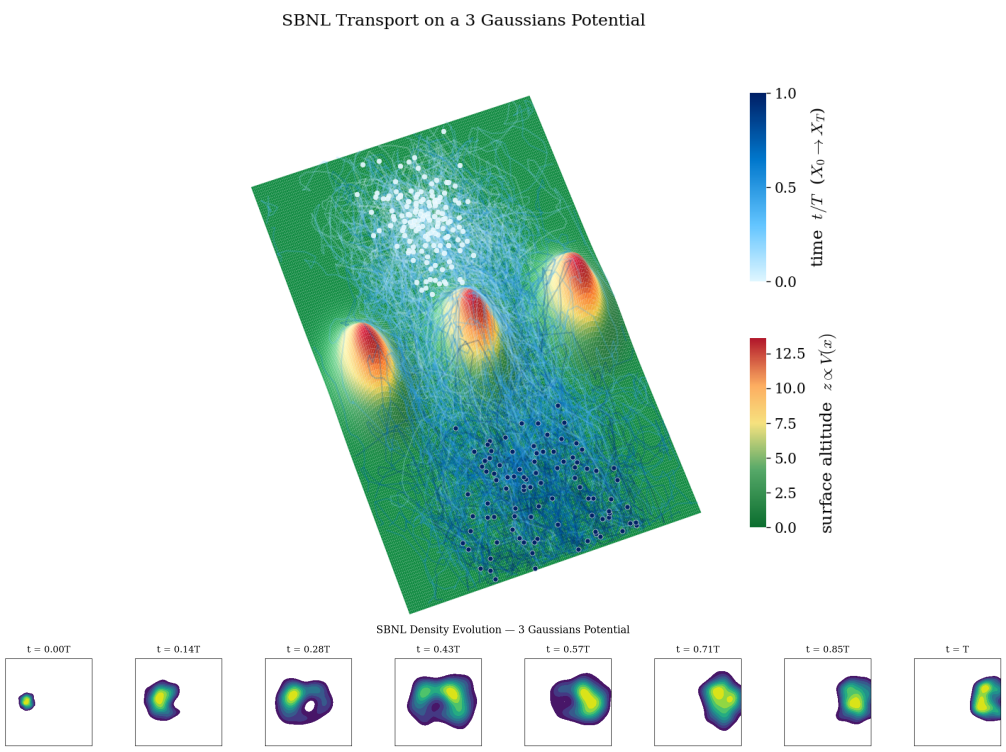

Figure 20: Three Gaussians: SBNL transport and estimated density evolution $m(t, \cdot) = \varphi(t, \cdot)\bar{\varphi}(T - t, \cdot)$.

O.4 50D SYNTHETIC EXPERIMENTS

We consider a nonlinear bridge between Gaussian mixtures in $\mathbb{R}^{50}$ with $c = 0.1$. Samples and trajectories are visualized by projection onto the first two PCA directions, with the PCA basis fit on pooled validation draws from μ_0 and μ_T. Figure 21 compares generated and target marginals after training; Figure 22 shows the evolution of projected trajectories.

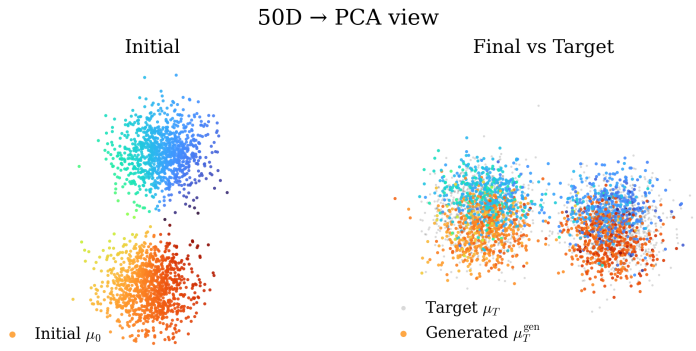

Figure 21: Generated vs. true samples in $\mathbb{R}^{50}$ (PCA projection).

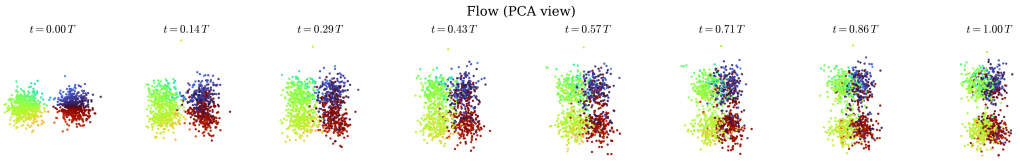

Figure 22: Temporal evolution in $\mathbb{R}^{50}$ (PCA projection).

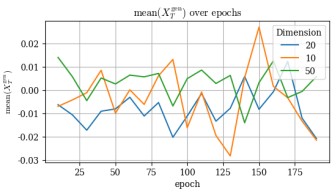

(23a) Mean of the generated samples across different dimensions (target mean: 0).

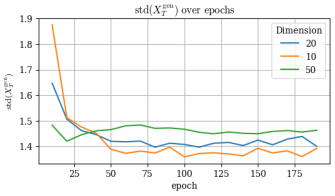

(23b) Std of the generated samples across different dimensions (target std: $\sqrt{2}$).

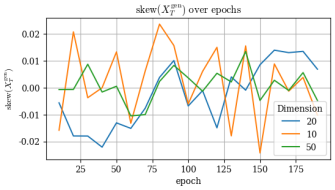

(23c) Skewness of the generated samples across different dimensions (target skewness: 0).

Figure 23: Evolution of the first three moments (mean, standard deviation, skewness) of the generated samples over training epochs, for different latent dimensions.

O.5 MNIST IN LATENT SPACE

We consider a nonlinear bridge between MNIST digits in latent space with $c = 0.1$. Samples and trajectories are visualized by projection onto the first two PCA directions, with the PCA basis fit on pooled validation draws from μ_0 and μ_T (Figure 6). Figure 24a projects the d=32 latent codes onto the first two PCA directions (basis fit on pooled validation samples from μ_0 and μ_T): the left panel shows the initial law μ_0, and the right panel overlays generated terminal latents on the target μ_T, illustrating that the bridge maps the source geometry into the target bulk while preserving local structure visible in the point colors. Figure 24b compares generated and target digits after training.

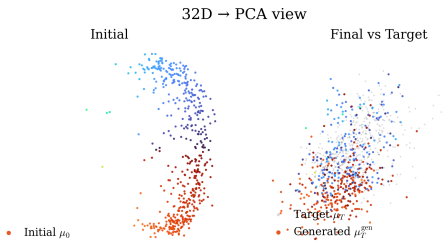

(24a) PCA view of μ_0 (left) and of generated terminal samples vs. target μ_T (right; target in gray, generated colored).

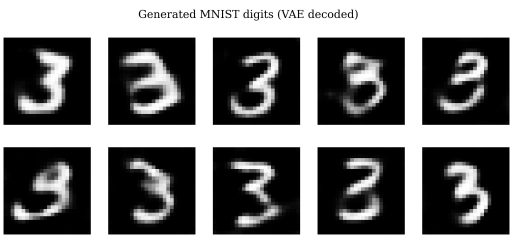

(24b) Generated MNIST digits obtained by decoding latent interpolations.

Figure 24: MNIST latent bridge (d=32): latent PCA visualization and decoded samples.

O.6 FFHQ IN LATENT SPACE



Figure 25: FFHQ young-to-old: decoded trajectories during sampling.

O.7 Comparison on FFHQ

We extend the FFHQ young-to-old comparison from Figure 7a to a larger panel of source children and all baselines, in ALAE latent space. Figure 26 shows, for each row, the source child (μ_0) on the left, the terminal samples produced by SBNL, LightSB, LightSBB-M, DSBM and BranchSBM (with separate female and male conditioning for BranchSBM), and a reference adult image drawn at random from μ_T on the right.

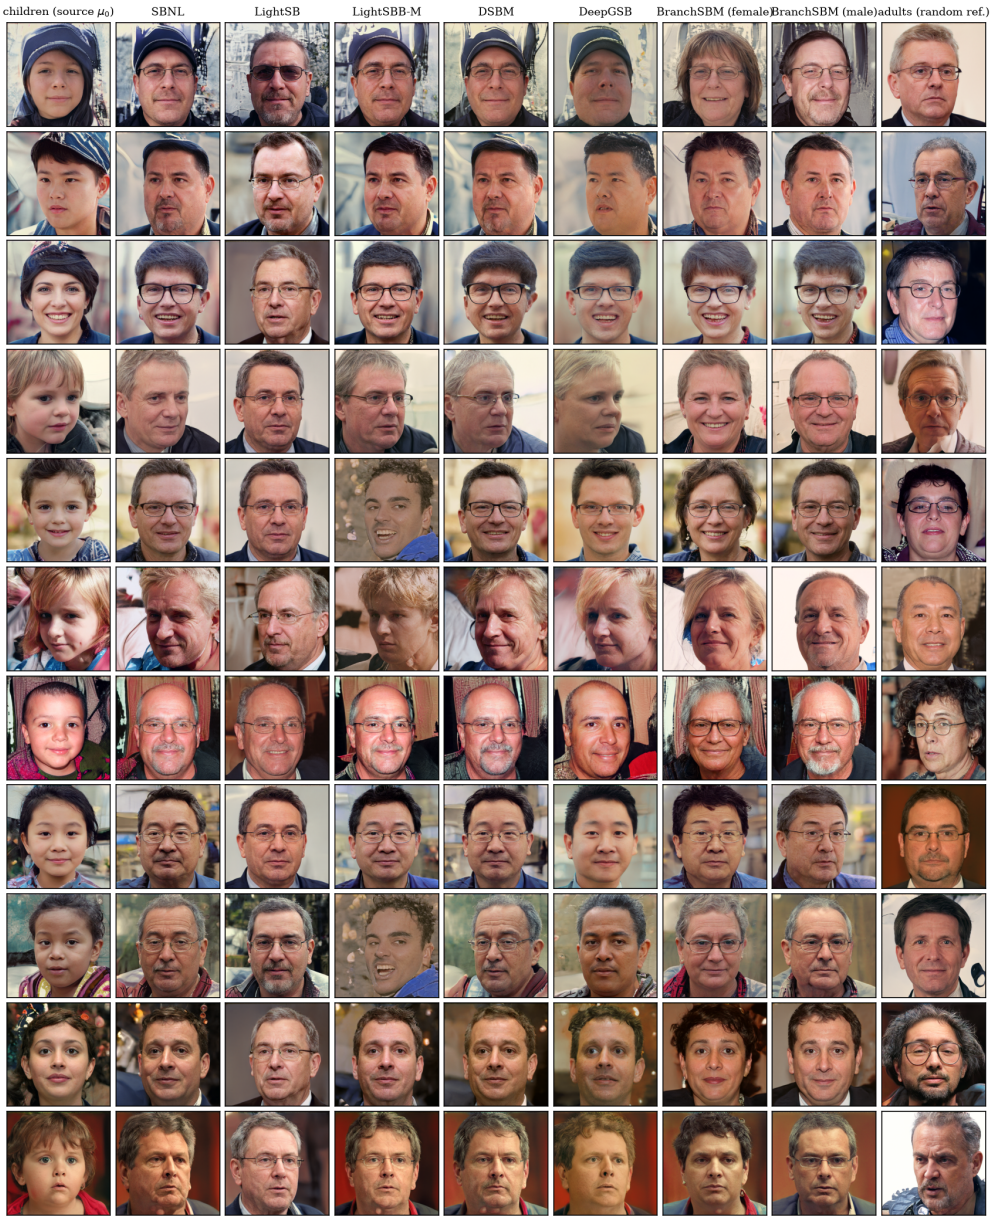

Figure 26: FFHQ young-to-old comparison in ALAE latent space, extended panel.

P ADDITIONAL FIGURES

P.1 EVOLUTION OF THE POTENTIALS DURING TRAINING

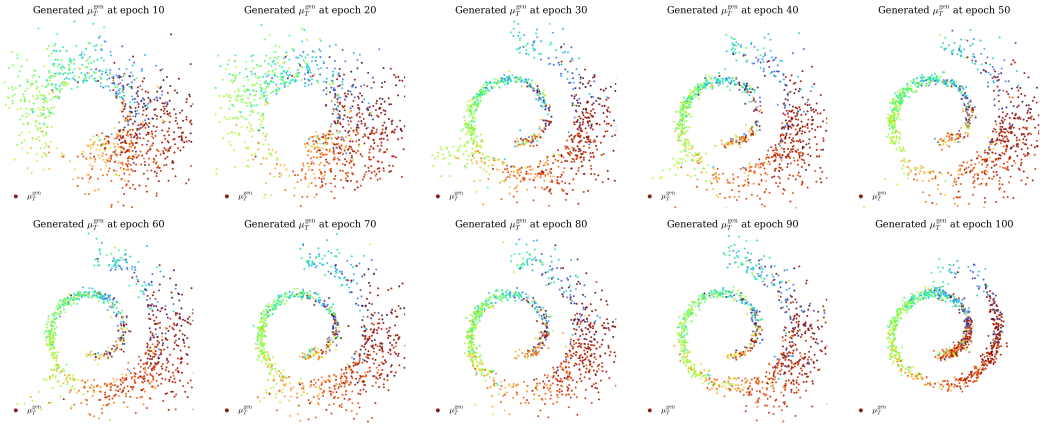

Figure 27: Convergence of $\mathrm{Law}(X_T)$ with training (Moon curves $\to$ Swiss).

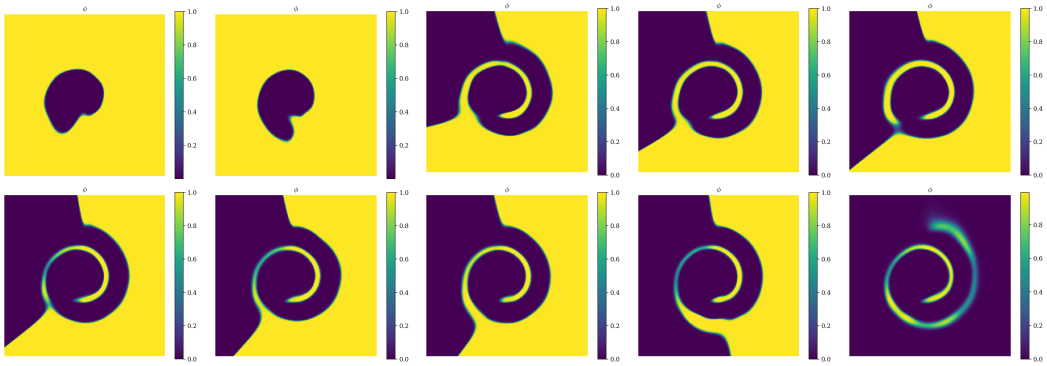

Figure 28: Evolution of potential φ (Moon curves $\to$ Swiss).

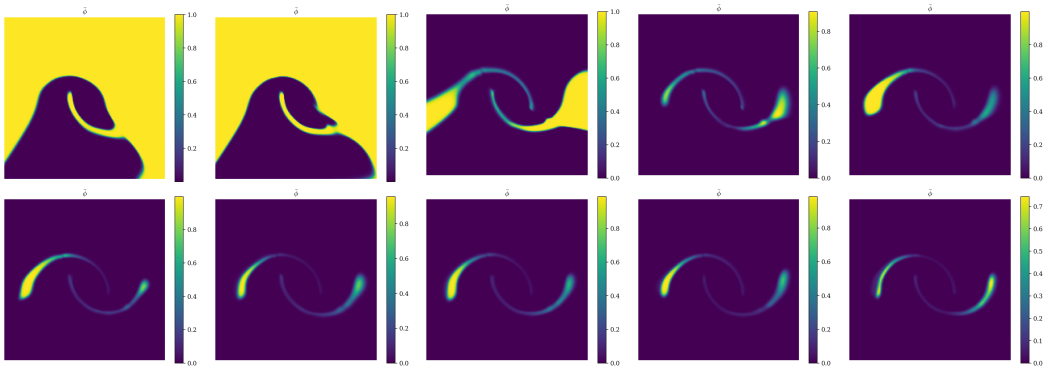

Figure 29: Evolution of potential $\bar{\varphi}$ (Moon curves $\to$ Swiss).

P.2 RELAXATION OF THE HARD CONSTRAINT AND FINITE PENALTY λ

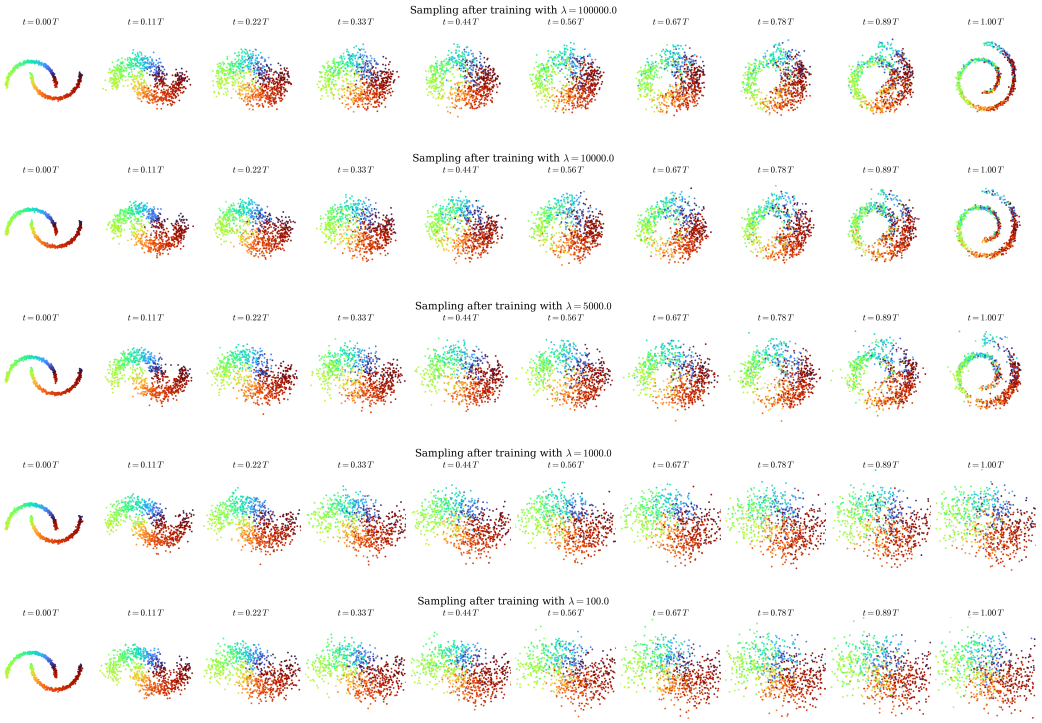

Figure 30: Flow evolution for $\lambda \in \{10^5, 10^4, 5\times10^3, 10^3, 10^2\}$.