FAILURE MODES OF ADJOINT SAMPLING: INSTABILITY AND SLOW MULTIMODAL ADAPTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

The Adjoint Sampling has been presented as a theoretically grounded approach to learning diffusion samplers from unnormalized target densities. Its existing theory characterizes the desired optimal policy as a fixed point of an idealized update, but does not establish that the proposed iteration converges to, or is locally stable around, that fixed point. We rigorously analyze its convergence using Gaussian and Gaussian-mixture targets. Their amenability to closed-form and asymptotic analysis allows us to isolate two fundamental challenges for the iterative sampler: local stability and probability redistribution across separated modes. In the Gaussian case, we prove that, over a nontrivial parameter regime, an arbitrarily small terminal-mean error alternates in sign and grows geometrically across iterations. For well-separated Gaussian mixtures, asymptotic analysis reveals two additional failure modes with exponentially slow convergence: discovery of an unrepresented mode and correction of inaccurate mode weights. These results hold with exact conditional expectations and density evolution, separating intrinsic properties of the iteration from errors due to finite sampling or neural-network approximation. Direct quadrature and Fokker–Planck solutions reproduce the predicted failure modes. Neural-network experiments show that direct velocity damping strongly suppresses the local Gaussian instability, while replay provides weaker moderation; neither intervention reliably resolves missing-mode discovery or probability redistribution across separated modes. Additional numerical experiments revealed that our identified failure modes persist beyond the idealized Gaussian setup and are also observed in higher dimensional state spaces. Together, these results show that fixed-point correctness does not imply convergence of AS method and that stabilizing local update dynamics does not by itself resolve global probability-allocation failures. Finally, we discuss the broader implications of our findings for the design and analysis of future diffusion-based sampling algorithms.

1 INTRODUCTION

Statistical sampling from an unnormalized density function is a fundamental task in Bayesian inference, statistical physics, and molecular simulation Faulkner & Livingstone (2023); Hénin et al. (2022); Wilkinson (2007); Sharma (2017); Jackman (2000); McCartan & Imai (2023). Given an unnormalized density function $\widetilde{\nu}(x)$, the goal is to generate i.i.d. samples from $\nu(x) = \widetilde{\nu}(x)/Z$ without explicitly calculating the normalization constant Z. Markov chain Monte Carlo (MCMC; Metropolis et al. (1953)) and its variants, such as Hamiltonian Monte Carlo (HMC; (Duane et al., 1987; Neal, 2012)), together with Sequential Monte Carlo (SMC; (Del Moral et al., 2006)) provide general-purpose solutions, but their performance can depend strongly on the design of transition kernels and intermediate distributions, particularly for high-dimensional or multimodal targets.

Recent advances in neural generative modeling have motivated learning models that transport samples from a tractable reference distribution to the target. Unlike conventional generative modeling, samples from the target are unavailable; learning is driven instead by evaluations of its unnormalized density or score. Annealed Flow Transport Monte Carlo (AFTMC) combines normalizing-flow transports with annealing and Monte Carlo correction (Arbel et al., 2021). Path Integral Sampler (PIS) and Denoising Diffusion Samplers (DDS) formulate sampling through stochastic optimal control or variational inference over path measures (Zhang & Chen, 2022; Vargas et al., 2023). Schrödinger bridge methods use iterative proportional fitting to learn controlled dynamics connect-

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.

ing prescribed endpoint distributions, with formulations developed both for sampling from unnormalized targets (Bernton et al., 2019) and for sample-based generative modeling (De Bortoli et al., 2021). Liouville Flow Importance Sampler (LFIS) learns a continuous-time deterministic probability flow along a prescribed annealing path and uses accumulated residuals as importance weights that correct imperfect learned transport (Tian et al., 2024). These methods differ in their dynamics and correction mechanisms, but share the aim of replacing or augmenting hand-designed local transitions with some form of learned transport.

Adjoint Sampling (AS; Havens et al., 2025) is a recently proposed diffusion-based sampler built upon Adjoint Matching (AM; Domingo-Enrich et al., 2025), a regression-based method for solving stochastic optimal control problems without explicit importance weighting. At each outer iteration, AS generates terminal samples from the current controlled process and applies a reciprocal projection (Shi et al., 2023), replacing the conditional path law of that process with the corresponding reference bridge while preserving its current terminal marginal. Because the reciprocally projected and target-inducing path measures share the same conditional reference bridge, the Radon–Nikodym derivative of the target-inducing measure with respect to the reciprocally projected measure reduces to the density ratio between their terminal marginals. Instead of explicitly evaluating this density ratio, AS updates the control through a regression driven by the prescribed terminal-cost gradient. This construction removes explicit importance weights and enables repeated, inexpensive generation of regression data—by resampling intermediate states conditioned on stored terminal samples—thereby permitting many neural-network parameter updates for each terminal sample and target-gradient evaluation. **Existing theory identifies the desired optimal control as a fixed point of the idealized update, but does not establish that the complete outer iteration decreases the objective, is locally stable around that fixed point, or converges to it.** The correctness of the resulting iterative sampler therefore depends on an unresolved convergence question.

Contributions. We analyze both the idealized exact update and the practical neural-network implementation of AS and identify three distinct failure modes: instability of the outer iteration, slow discovery of unrepresented modes, and slow redistribution of probability mass between separated modes. Our main contributions are:

- **Exact update dynamics and instability for Gaussian distributions.** We derive the closed-form AS update and its stability criterion. We prove that AS can be intrinsically unstable, that an arbitrarily small terminal-mean error can grow geometrically across iterations.
- **Asymptotic analysis of multimodal adaptation.** For well-separated, two-mode Gaussian mixtures, a rigorous asymptotic analysis reveals that using AS for discovering an initially unrepresented mode and for correcting inaccurate mode weights can be impractically slow.
- **Comprehensive numerical analyses support the theoretical prediction and generalization.** Numerical solution of the Fokker–Planck Equation reproduces all three failure modes, excluding errors from Monte Carlo estimation, neural-network approximation, and incomplete optimization. Through complementary neural-network ablations, we systematically distinguish the effects of finite network optimization, replay, and explicit damping, neither of which resolve the missing-mode discovery or redistribution of probability among separated modes. The same qualitative behaviors persist in higher dimensions (2 & 128) and for non-Gaussian targets, substantiating the generality of our theoretical characterization.

Together, these results demonstrate that fixed-point correctness and convergence are distinct requirements in the design and analysis of learned diffusion samplers.

Related work. Our study is closely related to three lines of work. First, practical implementations of AS use replay and incomplete inner-loop optimization, or velocity field damping (Blessing et al., 2026), which were reported to smooth optimization but depart from the idealized update analyzed theoretically (Havens et al., 2025). Recent work interprets AM as a continuous-time method of successive approximations and employs explicit damping in its numerical implementation (Domingo-Enrich & Han, 2026). Classical deterministic and stochastic-control studies have shown that unregularized successive-approximation updates may diverge and have introduced penalized variants with convergence guarantees under suitable assumptions (Li et al., 2018; Kerimkulov et al., 2021). These results motivate moderating the AM update, but do not apply directly to the reciprocal-projection iteration used by AS. Related trust-region methods for stochastic optimal control constrain the path-space divergence between successive controlled processes, replacing an unrestricted update with

a sequence of smaller transport steps (Blessing et al., 2025). Although these approaches apply to diffusion-based sampling, they modify the underlying control problem and do not analyze the stability of the reciprocal-projection iteration used by AS. Our analysis provides this missing characterization and directly exposes the instability and slow convergence as an intrinsic mathematical defect of the method, instead of a numerical one. Moreover, the difficulty of recovering separated modes and their relative probabilities from local score information has been identified as the blindness of score-based methods to isolated components and mixing proportions (Wenliang & Kanagawa, 2021). Our results establish how this information limitation manifests specifically in AS: for well-separated Gaussian mixtures, we derive the reciprocal update and quantify the exponentially weak signals governing both missing-mode discovery and correction of inaccurate mode weights. Finally, reciprocal and Markovian projections also appear in iterative Schrödinger-bridge methods, for which convergence of the idealized alternating-projection sequence has been established under appropriate assumptions (Shi et al., 2023; Kholkin et al., 2024). AS uses a different outer iteration, combining reciprocal projection with an AM control update. Consequently, convergence of the projection procedure itself does not imply convergence of AS. Our analyses and counterexamples demonstrate that retaining the desired solution as a fixed point is insufficient: the resulting outer iteration may be unstable or exponentially slow.

2 Background: Adjoint Sampling via Adjoint Matching

For completeness and to set up notations, we provide a brief review of AS, leaving detailed derivations to the original references. Let $p^{\text{tar}}(x) = \widetilde{\nu}(x)/Z$ denote the target density. AS formulates sampling as a stochastic optimal-control problem on a controlled diffusion

$$\mathrm{d}X_t = \sigma(t)\, u(X_t, t)\, \mathrm{d}t + \sigma(t)\, \mathrm{d}B_t, \qquad X_0 = 0. \tag{1}$$

Throughout this work, $\mathbf{X} := \{X_t\}_{t\in[0,1]}$ denotes an entire stochastic trajectory, whereas X_t denotes its state at time t and X_1 denotes its terminal state. We use lowercase letters, such as x_t and x_1, for their realizations. Accordingly, $p(\mathbf{X})$ denotes a path density, while $p_t(\cdot)$, $p_{t|s}(\cdot,\cdot)$ and $p_{t,s}(\cdot,\cdot)$ denote the time-t marginal density, two-time (t,s) conditional and joint densities, respectively. The objective is to find a control u whose terminal density agrees with p^{tar}. Given a control u, we denote its induced path density by $p^u(\mathbf{X})$ and its t-marginal by $p_t^u(x)$. We denote the path measure of the uncontrolled ($u = 0$) reference process, termed the *base process* in Havens et al. (2025), by $p^{\text{ref}}(\mathbf{X})$ and its terminal marginal by $p_1^{\text{ref}}(x)$. Define the terminal cost and its gradient as $g(x) := \log p_1^{\text{ref}}(x) - \log p^{\text{tar}}(x)$ and $\nabla_x g(x) = \nabla_x \log p_1^{\text{ref}}(x) - \nabla_x \log \widetilde{\nu}(x)$, respectively. The sampling problem is equivalent, up to an additive constant, to minimizing

$$\mathcal{J}(u) = \mathbb{E}_{p^u}\left[\frac{1}{2}\int_0^1 |u(X_t,t)|^2\, \mathrm{d}t + \nabla_x g(X_1)\right]. \tag{2}$$

The optimal path density retains the conditional bridge of the reference process while replacing its terminal marginal with the target, $p^*(\mathbf{X}) = p^{\text{ref}}(\mathbf{X}|X_1)\, p^{\text{tar}}(X_1)$. AM converts Eq. 2 into a regression problem (Domingo-Enrich et al., 2025). At outer iteration i, let u_i denote the current control and define its terminal density as $p^{\text{cur}}(x) := p_1^{u_i}(x)$. AS first constructs the reciprocal projection (Shi et al., 2023), $p^{\text{rec}}(\mathbf{X}) = p^{\text{ref}}(\mathbf{X}|X_1)\, p^{\text{cur}}(X_1)$. This projection preserves the current terminal density while replacing the conditional path law with the corresponding reference-process bridge. Since p^* and p^{rec} share the same conditional bridge, their density ratio reduces to $p^*(\mathbf{X})/p^{\text{rec}}(\mathbf{X}) = p^{\text{tar}}(X_1)/p^{\text{cur}}(X_1)$. AS does not evaluate this density ratio explicitly. Instead, AM back-propagates the terminal-cost gradient through an adjoint equation and uses the resulting adjoint variable as the regression target for updating the control:

$$\mathcal{L}_{\text{RAM}}(u; p^{\text{cur}}) = \mathbb{E}_{t\sim\lambda,\, X_1\sim p^{\text{cur}},\, X_t\sim p^{\text{ref}}_{t|1}(\cdot|X_1)}\left[\frac{1}{2}\left|u(X_t,t) + \sigma(t)\,\nabla_x g(X_1)\right|^2\right], \tag{3}$$

where λ is the distribution used to sample the regression time. With unrestricted function approximation and exact minimization, the updated control is[1]

$$u_{\text{new}}(x_t,t) = -\sigma(t)\, \mathbb{E}_{p^{\text{rec}}_{1|t}(X_1|X_t=x_t)}\left[\nabla_x g(X_1)\right]. \tag{4}$$

[1] Although RAM samples the reciprocal distribution $p^{\text{rec}}_{t,1}$ by first drawing $X_1 \sim p^{\text{cur}}$ and then $X_t \sim p^{\text{ref}}_{t|1}(\cdot \mid X_1)$, the exact least-squares minimizer is the conditional mean of the regression target X_1 given its input X_t.

In practice, terminal samples from p^{cur} are stored and reused ("replay buffer"). Multiple regression pairs are generated for each terminal sample X_1, by resampling $t \sim \lambda$ and $X_t|X_1 \sim p^{\text{ref}}_{t|1}(\cdot|X_1)$.

The existing theory shows that reciprocal projection does not increase $\mathcal{J}$ and interprets Eq. 4 as a unit functional-gradient step from the projected control. The optimal control is therefore a fixed point of the idealized update. These results, however, do not establish that the complete outer iteration decreases $\mathcal{J}$, converges to the desired fixed point, or is locally stable around it. Our analysis directly addresses these questions by studying the dynamics of Eq. 4.

3 THEORETICAL ANALYSIS

Here, we outline the explicit calculations and results, leaving the detailed derivations to Appendices. For all analyses in this section, we choose $\sigma(t) = 1$ without loss of generality, which is equivalent to a rescaling of time in Eqs. 1 and 4; our conclusion remains valid for general $\sigma(t)$.

We set out to consider Gaussian p^{tar}, Gaussian p^{cur} as the basic block of our analysis. We consider a one-dimensional state space $x \in \mathbb{R}^1$, $p^{\text{tar}}(x) = \mathcal{N}(x; \mu, \sigma^2)$ and $p^{\text{cur}}(x) = \mathcal{N}(x; m, s^2)$, where $(\mu, \sigma^2) \in \mathbb{R} \times \mathbb{R}_+$ and $(m, s^2) \in \mathbb{R} \times \mathbb{R}_+$ are the mean and variance parameters of the target and current distributions, respectively; the score terminal-density ratio is $\nabla_x g(x) = \left(1 - \sigma^2\right) x/\sigma^2 - \mu/\sigma^2$. This particular setup allows us **to study the stability of AS algorithm** in its purest form. We present the key results Prop. 3.1, Thm. 3.2, and Thm. 3.3 below and provide proofs in Appendix A.

Proposition 3.1. *For the one-dimensional Gaussian case defined above, the reciprocal projection leads to a Gaussian conditional distribution:*

$$p^{rec}_{1|t}(X_1|X_t) = \mathcal{N}\left(X_1; \frac{s^2 X_t + (1-t)m}{1 - t\left(1 - s^2\right)}, \frac{s^2(1-t)}{1 - t\left(1 - s^2\right)}\right). \tag{5}$$

Theorem 3.2. *The optimal control $u_{G,G}(x_t, t)$ for the one-dimensional Gaussian case is given by*

$$u_{G,G}(x_t, t) = -\frac{\left(1 - \sigma^2\right)\left[s^2 x_t + (1-t)m\right]}{\sigma^2\left[1 - t(1 - s^2)\right]} + \frac{\mu}{\sigma^2}. \tag{6}$$

Multi-dimensional Gaussian case can be treated similarly; see Corollary A.1 in Appendix A.2.

Theorem 3.3. *Given the control in Eq. 6, the resulting controlled process $dX_t = u_{G,G}(X_t, t)\ dt + dW_t$ will have the terminal distribution $p^{next}(x_1) = \mathcal{N}(x_1; m_+, s^2_+)$ where*

$$m_+ = \begin{cases} \mu + \left(1 + \frac{1}{1-s^2}\log s^2\right)(m - \mu), & \gamma = 1, \\ \mu + A_m\,(m - \mu), & \gamma \neq 1, \end{cases} \tag{7a}$$

$$s^2_+ = \begin{cases} \frac{s^2}{1-s^2}\log\frac{1}{s^2} & \gamma = \frac{1}{2}, \\ \frac{s^2 - s^{4\gamma}}{(1-s^2)(2\gamma - 1)} & \gamma \neq \frac{1}{2}, \end{cases} \tag{7b}$$

with

$$\gamma := \frac{s^2\left(1 - \sigma^2\right)}{\sigma^2\left(1 - s^2\right)}, \quad \textit{and} \quad A_m := \frac{(\gamma - 1)\left(1 - s^{2\gamma}\right) - \gamma\left(1 - s^{2(\gamma-1)}\right)}{\left(1 - s^2\right)(\gamma - 1)}. \tag{8}$$

The joint update of (m, s^2) has a triangular structure. The variance update is independent of m, whereas the mean error is slaved to the variance dynamics through equation 7a. As illustrated in Fig. 1, s^2 may undergo an initial transient before converging to the target variance σ^2. Consequently, the mean amplification factor approaches $A_m(\sigma^2; \sigma^2) = 1 + (1 - \sigma^2)^{-1}\log\sigma^2$. If the target variance lies in the unstable regime, $A_m(\sigma^2; \sigma^2) < -1$, then the magnitude of the mean amplification factor eventually remains greater than one, manifesting in linear instability of the updates. Unless an iteration happens to produce $m = \mu$ exactly, the mean error therefore alternates in sign and eventually grows geometrically. Moreover, even this exceptional case is unstable: an arbitrarily small perturbation away from $m = \mu$ is amplified by subsequent iterations. Thus, convergence of the variance does not imply stability of the joint mean and variance update.

We next consider mixture target and single-mode current distributions **to study if AS can efficiently sample a missing-mode**. This setting represents a practical difficulty in statistical and soft-matter

physics: local relaxation followed by sampling within a basin may characterize one metastable state, while transitions to other states remain rare (Frenkel & Smit, 2002; Hénin et al., 2022). By initializing the current distribution with the correct location and variance of one target component, we isolate mode discovery from errors in modeling the already captured mode.

Specifically, we consider

$$p^{\text{tar}}(x) = \frac{1}{2}\mathcal{N}(x; -\nu, \sigma^2) + \frac{1}{2}\mathcal{N}(x; \nu, \sigma^2), \tag{9a}$$

$$p^{\text{cur}}(x) = \mathcal{N}(x; \nu, \sigma^2), \tag{9b}$$

and define the scale-separation parameter $\varepsilon := \sigma/\nu \ll 1$. The current distribution has the correct location and variance for the right mode but assigns only exponentially small probability to the basin $x < 0$ of the left mode. We provide the key results, Prop. 3.4, Thm. 3.5, and Cor. 3.6 below, and leave the proofs in Appendix B.

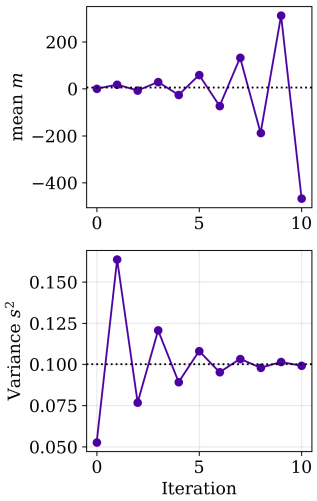


Figure 1: Joint update dynamics of m and s^2. The target parameters are $(\mu, \sigma^2) = (5.0, 0.1)$, and the initial parameters are $(m, s^2) = (0.2, 0.05263)$. The variance evolves independently of the mean and approaches the target variance after an initial transient. As the variance enters the unstable region, the mean error alternates in sign and grows in magnitude.

Proposition 3.4 (Localization of the missing-mode information)**.** *The target–reference score difference is* $-g_{GMM}(x) = x - x/\sigma^2 + (\nu/\sigma^2)\tanh(\nu x/\sigma^2)$. *For the corresponding single-Gaussian target* $\mathcal{N}(x; \nu, \sigma^2)$, $-g_G(x) = x - x/\sigma^2 + \nu/\sigma^2$. *The difference of the two scores,* $(\nu/\sigma^2)\left[\tanh(\nu x/\sigma^2) - 1\right]$, *is appreciated only when* $x \lesssim \mathcal{O}(\nu\varepsilon^2)$. *For any fixed* $\delta > 0$ *and* $x \geq \delta\nu$, $|g_{GMM}(x) - g_G(x)| \leq (2/\nu\varepsilon^2)\exp(-2\delta\varepsilon^{-2})$.

The missing mode can therefore influence the AM update only if either the reciprocal conditional reaches the score-transition layer or the forward process leaves the region concentrated around the captured mode.

Theorem 3.5 (Exponentially weak missing-mode correction)**.** *Fix* $\delta \in (0, 1)$ *independently of* ε, *with* $\varepsilon^2 \ll \delta$. *There exists a time-dependent safe region* $\mathcal{S}_t$ *carrying all but exponentially small probability under the captured-mode dynamics such that* $|u_{G,GMM}(x_t, t) - u_{G,G}(x_t, t)| = \mathcal{O}\left(\text{poly}(\varepsilon^{-1})e^{-\mathcal{C}\varepsilon^{-2}}\right)$, $x_t \in \mathcal{S}_t$ *and* $\mathbb{P}_{G,GMM}\{X_t \notin \mathcal{S}_t\} = \mathcal{O}\left(\text{poly}(\varepsilon^{-1})e^{-\mathcal{C}\varepsilon^{-2}}\right)$, *for some* $\mathcal{C} > 0$ *independent of* ε *and* t. *Here,* $u_{G,GMM}$ *and* $u_{G,G}$ *denote the optimal controls for the respective current-target pairs.*

Corollary 3.6 (Exponentially slow mode discovery)**.** *For the setting in Eq. 9, the AM dynamics associated with the bimodal target are exponentially close to those of the already captured right mode on trajectories carrying asymptotically dominant probability. Consequently, the probability available for discovery of the missing basin is exponentially small in* $1/\varepsilon^2$.

The proof of Thm. 3.5 combines two complementary large-deviation estimates. For typical states of the forward process, reaching the score-transition layer requires sampling far into the tail of the reciprocal conditional distribution. Conversely, bringing that transition layer into a region of appreciable reciprocal probability requires tail events of the forward process. Both events are exponentially unlikely in $1/\varepsilon^2$. The target mixture therefore produces no $\mathcal{O}(\varepsilon^0)$ correction to the single-mode control in the well-separated limit.

To quantify this scale, the single-Gaussian control u_{G} results in $\mathcal{N}(\nu, \sigma^2)$. To the leading order, the probability assigned to the left basin after an additional AM iteration is $\sim (\varepsilon/\sqrt{2\pi})\exp(-1/2\varepsilon^2)$. Interpreting this quantity as the characteristic inter-basin mass available per outer iteration gives the estimate for the mode-discovery timescale $\tau_{\text{disc}} \sim (\sqrt{2\pi}/\varepsilon)\exp(1/2\varepsilon^2)$. This estimate is analogous to the exponentially long transition times of metastable diffusion (Assaf & Meerson, 2010) or discrete-state stochastic processes (Doering et al., 2005). Reciprocal projection does not provide an order-one mechanism for discovering an unrepresented, well-separated mode.

We finally consider the case in which both the target and current distributions contain two modes, **to study if AS can efficiently adjust the relative weights of discovered modes**. This is a related but

distinct task: initialization and local relaxation may populate several basins, but their relative populations depend on the initialization and need not reflect the free-energy differences that determine their equilibrium populations. By choosing a current mixture with correct component locations and variances but incorrect mode weights, we isolate the recovery of these equilibrium populations from mode discovery and local fitting errors. Consider the target and current distributions

$$p^{\text{tar}}(x) = \varpi\mathcal{N}(x; -\nu, \sigma^2) + (1-\varpi)\mathcal{N}(x; \nu, \sigma^2), \tag{10a}$$

$$p^{\text{cur}}(x) = w\mathcal{N}(x; -m, s^2) + (1-w)\mathcal{N}(x; m, s^2), \tag{10b}$$

where $\varpi, w \in (0,1)$. As before, we consider a well-separated target with $\varepsilon := \sigma/\nu \ll 1$. The target–reference score difference is $-\nabla_x g_{\text{GMM}}(x) = x - x/\sigma^2 + \nu/\sigma^2 \tanh\left((x - x_c^{\text{tar}})/\nu\varepsilon^2\right)$, where $x_c^{\text{tar}} := \sigma^2/(2\nu)\log(\varpi/(1-\varpi))$. Thus, the target weight ϖ affects the score primarily through the location of a narrow transition layer of width $\mathcal{O}(\nu\varepsilon^2)$ between the two modes.

Proposition 3.7 (Exact reciprocal decomposition)**.** *Define $\psi_s(t) := 1 - t(1-s^2)$ and $\theta(x_t, t) := mx_t/\psi_s(t)$. The reciprocal conditional distribution admits the exact decomposition*

$$p^{rec}_{1|t}(x_1|x_t) = r_L(x_t,t)p^{rec}_{1|t,L}(x_1|x_t, L) + r_R(x_t,t)p^{rec}_{1|t,R}(x_1|x_t, R), \tag{11}$$

where

$$r_L(x_t,t) = \frac{w}{w + (1-w)e^{2\theta(x_t,t)}}, \quad r_R(x_t,t) = \frac{(1-w)e^{2\theta(x_t,t)}}{w + (1-w)e^{2\theta(x_t,t)}}. \tag{12a}$$

The exact AM control consequently decomposes as

$$u(x_t,t) = r_L(x_t,t)u_L(x_t,t) + r_R(x_t,t)u_R(x_t,t), \tag{13}$$

where $u_{\{L,R\}}(x_t,t) := \mathbb{E}_{p^{rec}_{1|t,\{L,R\}}(X_1|X_t=x_t,\{L,R\})}[-g_{GMM}(X_1)]$. Equivalently,

$$u(x_t,t) = \frac{u_L(x_t,t) + u_R(x_t,t)}{2} + \frac{u_R(x_t,t) - u_L(x_t,t)}{2}\tanh\left[\frac{m}{\psi_s(t)}\left(x_t - x_c^{cur}(t)\right)\right], \tag{14}$$

with $x_c^{cur}(t) := (\psi_s(t)/(2m))\log(w/(1-w))$.

Proposition 3.7 is exact and does not require mode separation. Importantly, the component responsibilities r_L and r_R depend on the current weight w, whereas the target weight ϖ enters only through the component controls u_L and u_R via the target score.

Theorem 3.8 (Exponentially weak correction of inaccurate mode weights)**.** *Suppose that the current mixture has the correct component locations and variances, $m = \nu$ and $s^2 = \sigma^2$, but an inaccurate weight $w \neq \varpi$. Assume that w and ϖ remain bounded away from zero and one as $\varepsilon \downarrow 0$. Then, on regions carrying all but exponentially small probability, $u_{\{L,R\}}(x_t,t) = u^G_{\{L,R\}}(x_t,t) + \mathcal{O}\left(\text{poly}(\varepsilon^{-1})e^{-c\varepsilon^{-2}}\right)$, for some $c > 0$, where u^G_L and u^G_R are the controls induced by the corresponding left and right single-Gaussian targets. Consequently,*

$$u(x_t,t) = r_L(x_t,t)u^G_L(x_t,t) + r_R(x_t,t)u^G_R(x_t,t) + \mathcal{O}\left(\text{poly}(\varepsilon^{-1})e^{-c\varepsilon^{-2}}\right). \tag{15}$$

The leading-order control depends on the current weight w through r_L and r_R, but is independent of the target weight ϖ.

Corollary 3.9 (Exponentially slow weight equilibration)**.** *In the infinitely separated limit, every mixture weight $w \in (0,1)$ is a leading-order fixed point of the AM iteration when the component locations and variances are correct. For finite but small ε, the signal that distinguishes w from the target weight ϖ is exponentially small in $1/\varepsilon^2$. Correction of inaccurate mode weights therefore occurs on an exponentially long characteristic iteration scale.*

Proofs of Prop. 3.7, Thm. 3.8, and Cor. 3.9 are provided in Appendix C.

This failure mode is distinct from missing-mode discovery. Here, both modes are already represented, but the local target score is asymptotically insensitive to their relative normalization. The global target weight is encoded only in the exponentially unlikely overlap region between the modes.

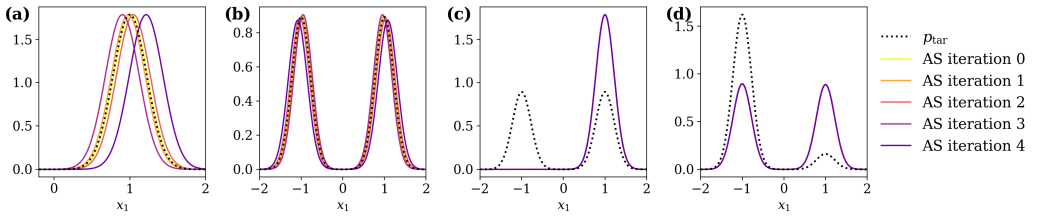


Figure 2: Direct quadrature and FPE evolution of successive AS terminal distributions. Curves show probability density functions, colors indicate AS iterations, and dotted lines are the target density functions. (a) An initially small mean error produces alternating, growing displacements for an effectively single-Gaussian target. (b) Symmetric perturbations produce alternating inward and outward motion of two modes, exposing instability in a second-moment deformation. (c) A correctly captured right mode shows little visible discovery of the missing left mode. (d) Correctly located modes retain nearly equal weights despite a 10:1 target component-weight ratio. Nearly overlapping curves in (c–d) indicate little change over the displayed iterations.

4 Numerical Analysis

We investigate the predicted failure modes through two complementary numerical approaches. Direct quadrature and Fokker–Planck solutions test the idealized AS iteration without learning or finite-sample errors. Neural-network implementations examine how practical approximation, replay, and damping modify these dynamics.

4.1 Direct solve of the Fokker–Planck Equation

We first investigate the iterative dynamics of AS by directly computing quadrature and evolving probability densities by the Fokker–Planck Equation (FPE). This provides a controlled, best-case setting for examining the update itself: the target score is evaluated analytically, the reciprocal conditional expectation is computed by deterministic quadrature, and the resulting controlled process is propagated by solving the FPE. With this setup, we avoid approximations from neural-network learning, stochastic optimization and finite-sample estimation, but introduce controllable and verifiable domain truncation, numerical quadrature, spatial and temporal discretization, and round-off error of the terminal density. All experiments in this subsection (Sec. 4.1) are in one dimensional setting ($D = 1$). A detailed documentation of the numerical methods, including the discretization schemes, quadrature rules, and experimental setup, is provided in Appendix D.3.

Figure 2 shows the evolution of the AS terminal distributions for different initial conditions and target configurations, under various experimental setups. Specifically, we adopt the most general mixture-to-mixture setup Eq. 10, fixing $\nu = 1$ and $\sigma^2 = s^2 = 0.05$ but with different initial perturbations and weight configurations corresponding to panels (a–d) in the figure:

- **(a)** By setting $\varpi = w = 10^{-20}$ and $m = 1.01$, the target and current distributions are effectively single Gaussians with a small initial mean error. The result exhibits alternating, growing displacements as predicted by Thm. 3.3.
- **(b)** By setting $\varpi = w = 1/2$ and $m = 1.01$, we aim to simulate a symmetric bimodal current distribution with slightly excessive mode separation from the symmetric bimodal target. The result exhibits alternating but symmetric inward and outward motion. This shows more complex update dynamics incurred by the linear instability prescribed in Thm. 3.3.
- **(c)** With $\varpi = 1/2$, $w = 10^{-20}$, and $m = 1$, we study the scenario where the current distribution has a correctly captured right mode but missing the left mode of a symmetric bimodal distribution. The result shows little visible discovery of the missing left mode, illustrating slow inter-mode probability redistribution, predicted by Thm. 3.5 and Cor. 3.6.
- **(d)** With $\varpi = 10/11$, $w = 1/2$, and $m = 1$, we study the scenario where the current distribution has correctly located modes but with a target component-weight ratio of $10 : 1$. The result shows that the modes retain nearly equal weights despite the target imbalance, illustrating the difficulty of adjusting component weights in the AS dynamics, as predicted by Cor. 3.9.

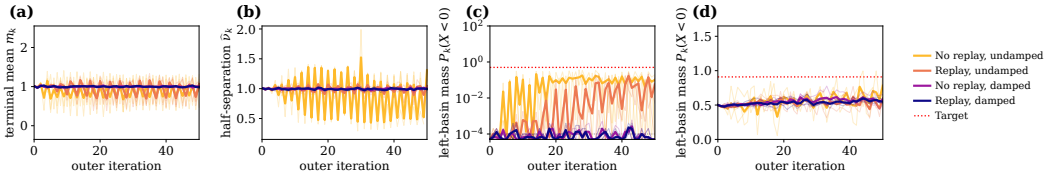


Figure 3: Neural-network Adjoint Sampling for the four one-dimensional Gaussian failure-mode tests. Thin curves show individual seeds and thick curves show the three-seed mean; dashed red lines denote target values. (a) Single-Gaussian terminal mean. (b) Half-separation $\widehat{\nu}_k = (m_{R,k} - m_{L,k})/2$ for the symmetric two-mode problem, where $m_{L,k}$ and $m_{R,k}$ are the empirical conditional means in the left and right basins. (c) Left-basin probability for missing-mode discovery on a logarithmic scale. The dotted line at 10^{-4} marks one count in the 10^4-sample diagnostic population; empirical zeros are displayed below this line for visibility. (d) Left-basin probability $P_k(X < 0)$ for the incorrect-weight problem, whose target is $10/11$.

Generalizability. We provide additional numerical experiments to demonstrate the robustness of the theoretical predictions. Specifically, parallel experiments were conducted to the anharmonic and Ginzburg–Landau potentials, whose log-density functions $\log \widetilde{\nu}(x)$ are $\alpha_2 x^2 + \alpha_4 x^4$ and $\kappa \left(x^2 - \nu^2\right)^2 + \xi x$ up to an additive constant, respectively. Detailed experimental setup can be found in Appendix F.1. All failed modes were observed with these non-Gaussian potentials, showing that they are not restricted to Gaussian family; see Figs. 11 and 12. Particularly, Fig. 12(d) shows a complex update dynamics due to the interplay between the instability and the multimodality.

4.2 Approximations via Neural Networks

We next test the same failure modes in the practical neural-network (NN) implementation of AS, combining replay (Havens et al., 2025)) on/off with direct velocity damping (Blessing et al., 2026) on/off. For each problem and seed, all four variants begin from the same learned initial control. Full initialization, optimization, sampling, and diagnostic details are given in Appendix E.

Figure 3 compares NN-based AS across the four one-dimensional Gaussian experiments in Sec. 4.1, with replay and direct velocity damping independently enabled or disabled. Individual trajectories and three-seed means reveal both systematic effects and seed dependence. Terminal-density and Wasserstein analyses in Appendix E.2 further distinguish stabilization of individual moments from recovery of the full distribution.

(a) Single-Gaussian instability. The NN implementation replaces the geometric divergence of the exact iteration with bounded oscillations. This suggests that imperfect learning—through finite-sample regression, limited function approximation, and incomplete inner-loop optimization—introduces numerical friction that moderates the instability. This moderation occurs without replay or explicit damping, but does not by itself restore convergence. Direct velocity damping substantially improves both the mean and the full terminal distribution without replay. Replay alone provides weaker moderation.

(b) Symmetric bimodal oscillation. The mode separation exhibits inward–outward oscillations resembling the FPE results. Damping keeps the separation close to its target, but the mean final-20 Wasserstein errors remain significant with or without replay, indicating that AS does not recover the target mixture.

(c) Missing-mode discovery. A representative undamped trajectory without replay approaches an approximately period-2 alternation between two distinct terminal distributions, neither matching the target. One places appreciable mass in the missing basin, creating the appearance of successful discovery. Replay produces less phase-locked recurrence; clustering terminal distributions by their empirical quantile functions identifies two recurrent states in Wasserstein distance (Fig. 4). Damping suppresses these excursions but leaves the missing-mode probability near 10^{-4}, the diagnostic resolution of 10^4 samples, compared with the target $1/2$. These observations support an instability-driven interpretation of the apparent recovery, while showing that suppressing the oscillations does not resolve mode discovery.

(d) Weight correction. None of the four variants reliably recovers the target allocation. The final-20 mean left-basin probabilities range from approximately 0.53 to 0.61, far from the target $10/11$. Individual-seed trajectories reveal recurrent, nonconvergent changes that are obscured by seed averaging. Replay and damping therefore do not reliably correct the relative weights, even when both modes are initially represented.

Generalizability. We further test the Gaussian failure modes beyond one dimension. In two dimensions, the local stability behavior remains broadly consistent with the Gaussian theory, while multimodal recovery and weight redistribution remain difficult and can exhibit recurrent dynamics; see Figs. 5, 6,7, 8 and Appendix E.3. In 128 dimensions, the NN implementation does not exhibit the instability predicted for the idealized Gaussian update, but still fails to recover the missing mode or correct the mixture weights over the tested iterations (Fig. 10; Appendix E.4). Finally, curvature-matched logit–Beta targets reproduce the qualitative distinction observed for Gaussian targets: damping improves local stability, while missing-mode discovery and weight correction remain difficult (Fig. 13; Appendix F.2). These experiments show that the observed failures extend beyond 1D Gaussian models, although their manifestation depends on the practical implementation.

5 Discussion

Our analysis identifies intrinsic limitations of the idealized AS iteration: local instability can occur even with exact regression and density evolution, while discovery and weight correction of separated modes can be extremely slow. Analytically tractable targets isolate these mechanisms from learning errors, and numerical experiments demonstrate related failures beyond the Gaussian setting. These findings distinguish correctness of the desired fixed point from stability and convergence of the iteration used to reach it.

Our numerical results clarify both the benefits and limitations of existing stabilization strategies. Direct velocity damping, already used in related work (Blessing et al., 2026; Domingo-Enrich & Han, 2026), substantially improves the single-Gaussian terminal distribution, whereas replay provides weaker moderation. Our Gaussian analysis gives a theoretical basis for choosing the relaxation parameter, complementing its use as a practical heuristic. A natural next step is to develop adaptive relaxation or path-space trust-region updates with stability guarantees for the complete AS iteration, including reciprocal projection and control fitting. These remain directions for further investigation rather than remedies established by the present work.

Stabilization alone, however, does not resolve multimodal sampling. In our experiments, damping suppresses unstable excursions without reliably recovering missing modes or correcting their relative weights. These tasks require both exploration of poorly represented basins and information about their relative probability mass. Combining AS regression with importance weighting, sequential correction, or annealed exploration may provide this information while retaining some of its computational advantages. Importance weighting itself requires adequate proposal coverage and controlled weight variance; it is therefore a component of a possible remedy, rather than a complete solution.

AS offers an efficient regression-based construction, but computational scalability alone does not ensure reliable sampling. The challenge is to retain this efficiency while ensuring stable outer updates and effective probability redistribution across separated modes. Our analysis provides tractable test cases for developing and assessing such modifications. More broadly, it shows that fixed-point correctness, iterative stability, and global probability allocation must be addressed together in the future design of learned diffusion samplers.

AI USE STATEMENT

In this work, we used large language models (LLMs), including ChatGPT and Claude, primarily to polish author-written text and improve its clarity. We also used these tools to reexamine mathematical derivations, sketch initial versions of selected numerical implementations—including PDE quadrature and Fokker–Planck solvers, as well as neural-network implementations of the Adjoint Sampling method—and supplement and proofread an author-conducted literature survey.

All LLM-assisted mathematical derivations were independently checked by the authors. All LLM-generated code snippets were reviewed, revised, tested, and incorporated into author-developed implementations before final numerical experiments. The authors executed all numerical simulations and analyzed the resulting data. The literature survey was initially conducted and drafted by the authors, who subsequently verified any references or claims suggested by the LLM tools against the original sources.

We did not use LLMs to generate the research ideas or scientific conclusions of this work. We have reviewed all LLM-assisted work and take full responsibility for the final content of this manuscript, including all text, mathematical claims, code, and other artifacts produced with the assistance of LLMs.

REPRODUCIBILITY STATEMENT

Detailed derivations and proofs are provided in the appendices so that all theoretical claims can be independently verified. We provide the numerical implementations, including the PDE quadrature and Fokker–Planck solvers and the neural-network-based implementation of the Adjoint Sampling, in an anonymous repository. The repository contains all code, data, configuration files, and instructions necessary to reproduce the numerical results presented in this work.

REFERENCES

Michael Arbel, Alex Matthews, and Arnaud Doucet. Annealed flow transport Monte Carlo. In *International Conference on Machine Learning*, pp. 318–330. PMLR, 2021.

Michael Assaf and Baruch Meerson. Extinction of metastable stochastic populations. *Phys. Rev. E*, 81:021116, Feb 2010. doi: 10.1103/PhysRevE.81.021116. URL https://link.aps.org/doi/10.1103/PhysRevE.81.021116.

Espen Bernton, Jeremy Heng, Arnaud Doucet, and Pierre E Jacob. Schr\" odinger bridge samplers. *arXiv preprint arXiv:1912.13170*, 2019.

Denis Blessing, Julius Berner, Lorenz Richter, Carles Domingo-Enrich, Yuanqi Du, Arash Vahdat, and Gerhard Neumann. Trust region constrained measure transport in path space for stochastic optimal control and inference. In *Advances in Neural Information Processing Systems*, volume 38. Curran Associates, Inc., 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/f1ccd13e51f610ff4a7a08db9a8b458b-Abstract-Conference.html.

Denis Blessing, Lorenz Richter, Julius Berner, Egor Malitskiy, and Gerhard Neumann. Bridge matching sampler: Scalable sampling via generalized fixed-point diffusion matching. *arXiv preprint arXiv:2603.00530*, 2026.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 17695–17709. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/940392f5f32a7ade1cc201767cf83e31-Paper.pdf.

Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(3):411–436, 2006.

Charles R. Doering, Khachik V. Sargsyan, and Leonard M. Sander. Extinction times for birth-death processes: Exact results, continuum asymptotics, and the failure of the fokker–planck approximation. *Multiscale Modeling & Simulation*, 3(2):283–299, 2005. doi: 10.1137/030602800. URL https://doi.org/10.1137/030602800.

Carles Domingo-Enrich and Jiequn Han. Adjoint matching through the lens of the stochastic maximum principle in optimal control. *arXiv preprint arXiv:2604.08580*, 2026. doi: 10.48550/arXiv.2604.08580. URL https://arxiv.org/abs/2604.08580.

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.

Simon Duane, A.D. Kennedy, Brian J. Pendleton, and Duncan Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, September 1987. ISSN 0370-2693. doi: 10.1016/0370-2693(87)91197-X.

Michael F. Faulkner and Samuel Livingstone. Sampling algorithms in statistical physics: A guide for statistics and machine learning, June 2023.

Daan Frenkel and Berend Smit. Chapter 16 - rare events. In *Understanding Molecular Simulation (Second Edition)*, pp. 431–464. Academic Press, San Diego, second edition edition, 2002. ISBN 978-0-12-267351-1. doi: https://doi.org/10.1016/B978-012267351-1/50018-3. URL https://www.sciencedirect.com/science/article/pii/B9780122673511500183.

Aaron J Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Daniel S. Levine, Brandon M Wood, Bin Hu, Brandon Amos, Brian Karrer, Xiang Fu, Guan-Horng Liu, and Ricky T. Q. Chen. Adjoint sampling: Highly scalable diffusion samplers via adjoint matching. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 22204–22237. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/havens25a.html.

Jérôme Hénin, Tony Lelièvre, Michael R. Shirts, Omar Valsson, and Lucie Delemotte. Enhanced sampling methods for molecular dynamics simulations [Article v1.0]. *Living Journal of Computational Molecular Science*, 4(1):1583, December 2022. doi: 10.33011/livecoms.4.1.1583.

Simon Jackman. Estimation and Inference via Bayesian Simulation: An Introduction to Markov Chain Monte Carlo. *American Journal of Political Science*, 44(2):375, April 2000. ISSN 00925853. doi: 10.2307/2669318.

Bekzhan Kerimkulov, David Šiška, and Łukasz Szpruch. A modified MSA for stochastic control problems. *Applied Mathematics & Optimization*, 84:3417–3436, 2021. doi: 10.1007/s00245-021-09750-2. URL https://doi.org/10.1007/s00245-021-09750-2.

Sergei Kholkin, Grigoriy Ksenofontov, David Li, Nikita Kornilov, Nikita Gushchin, Alexandra Suvorikova, Alexey Kroshnin, Evgeny Burnaev, and Alexander Korotin. Diffusion and adversarial Schrödinger bridges via iterative proportional markovian fitting. *arXiv preprint arXiv:2410.02601*, 2024. doi: 10.48550/arXiv.2410.02601. URL https://arxiv.org/abs/2410.02601.

Qianxiao Li, Long Chen, Cheng Tai, and Weinan E. Maximum principle based algorithms for deep learning. *Journal of Machine Learning Research*, 18(165):1–29, 2018. URL https://www.jmlr.org/papers/v18/17-653.html.

Cory McCartan and Kosuke Imai. Sequential Monte Carlo for sampling balanced and compact redistricting plans. *The Annals of Applied Statistics*, 17(4), December 2023. ISSN 1932-6157. doi: 10.1214/23-AOAS1763.

Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. ISSN 0021-9606. doi: 10.1063/1.1699114.

Radford M. Neal. MCMC using Hamiltonian dynamics. *arXiv:1206.1901 [physics, stat]*, June 2012. doi: 10.1201/b10905.

Sanjib Sharma. Markov chain Monte Carlo methods for Bayesian data analysis in astronomy. *Annual Review of Astronomy and Astrophysics*, 55(1):213–259, 2017. doi: 10.1146/annurev-astro-082214-122339.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in neural information processing systems*, 36:62183–62223, 2023.

Simo Särkkä and Arno Solin. *Useful Theorems and Formulas for SDEs*, pp. 98–125. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2019.

Yifeng Tian, Nishant Panda, and Yen Ting Lin. Liouville flow importance sampler. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 48186–48210. PMLR, 21–27 Jul 2024.

Francisco Vargas, Will Sussman Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=8pvnfTAbu1f.

Li K. Wenliang and Heishiro Kanagawa. Blindness of score-based methods to isolated components and mixing proportions. *arXiv preprint arXiv:2008.10087*, 2021. doi: 10.48550/arXiv.2008.10087. URL https://arxiv.org/abs/2008.10087.

Darren J. Wilkinson. Bayesian methods in bioinformatics and computational systems biology. *Briefings in Bioinformatics*, 8(2):109–116, April 2007. ISSN 1467-5463. doi: 10.1093/bib/bbm007.

Qinsheng Zhang and Yongxin Chen. Path integral sampler: A stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=_uCb2ynRu7Y.

A Gaussian-to-Gaussian Adjoint Matching Update

We provide the detailed derivation of the Gaussian-to-Gaussian update presented in Section 3. After the time transformation discussed in Section 3, the noise schedule is homogenized to one, and the controlled process becomes

$$\mathrm{d}X_t = u(X_t, t)\,\mathrm{d}t + \mathrm{d}W_t, \qquad X_0 = 0, \qquad t \in [0, 1]. \tag{16}$$

The reference process is therefore a standard Wiener process. In this appendix, σ^2 denotes the variance of the target distribution, not the noise schedule. We consider

$$p^{\text{tar}}(x) = \mathcal{N}(x; \mu, \sigma^2), \qquad p^{\text{cur}}(x) = \mathcal{N}(x; m, s^2). \tag{17}$$

The terminal marginal of the reference process is $p^{\text{ref}}(x_1) = \mathcal{N}(x_1; 0, 1)$. The target–reference score difference is therefore

$$-\nabla_x g(x_1) = \nabla_{x_1} \log \frac{p^{\text{tar}}(x_1)}{p_1^{\text{ref}}(x_1)} = -\frac{x_1 - \mu}{\sigma^2} + x_1 = -\frac{1 - \sigma^2}{\sigma^2} x_1 + \frac{\mu}{\sigma^2}. \tag{18}$$

The exact Adjoint Matching update is the conditional expectation

$$u_{\text{G,G}}(x_t, t) = \mathbb{E}_{p^{\text{rec}}_{1|t}(X_1 | X_t = x_t)} \left[-\nabla_x g(X_1) \right]. \tag{19}$$

The reciprocal projection replaces the conditional path law of the current controlled process with that of the reference process while preserving the current terminal marginal. Its joint density at times t and 1 is

$$p^{\text{rec}}_{t,1}(x_t, x_1) = p^{\text{ref}}_{t,1}(x_t, x_1) \frac{p^{\text{cur}}(x_1)}{p_1^{\text{ref}}(x_1)}. \tag{20}$$

Conditioning on $X_t = x_t$ gives

$$p^{\text{rec}}_{1|t}(x_1 | x_t) \propto p^{\text{ref}}(x_1 | x_t) \frac{p^{\text{cur}}(x_1)}{p^{\text{ref}}(x_1)}. \tag{21}$$

For the Wiener reference process,

$$p^{\text{ref}}_{1|t}(x_1|x_t) = \mathcal{N}(x_1; x_t, 1-t). \tag{22}$$

Substituting the three Gaussian densities into Eq. 21 yields

$$p^{\text{rec}}_{1|t}(x_1|x_t) \propto \exp\left[-\frac{(x_1-x_t)^2}{2(1-t)} - \frac{(x_1-m)^2}{2s^2} + \frac{x_1^2}{2}\right]. \tag{23}$$

Collecting the quadratic and linear terms in x_1, we obtain

$$\text{Var}^{\text{rec}}_{1|t}(X_1|X_t = x_t) = \left(\frac{1}{1-t} + \frac{1}{s^2} - 1\right)^{-1} = \frac{s^2(1-t)}{1-t(1-s^2)}, \tag{24a}$$

$$\mathbb{E}^{\text{rec}}_{1|t}[X_1|X_t = x_t] = \frac{s^2 x_t + (1-t)m}{1-t(1-s^2)}. \tag{24b}$$

Therefore,

$$p^{\text{rec}}_{1|t}(X_1|X_t) = \mathcal{N}\left(X_1; \frac{s^2 X_t + (1-t)m}{1-t(1-s^2)}, \frac{s^2(1-t)}{1-t(1-s^2)}\right). \tag{25}$$

Substituting the conditional mean into Eq. 19 gives the exact Adjoint Sampling update

$$u_{\text{G,G}}(x_t, t) = -\frac{(1-\sigma^2)\left[s^2 x_t + (1-t)m\right]}{\sigma^2\left[1-t(1-s^2)\right]} + \frac{\mu}{\sigma^2}. \tag{26}$$

In particular, the updated control is affine in x_t:

$$u_{\text{G,G}}(x_t, t) = \alpha(t)x_t + \beta(t), \tag{27}$$

where

$$\alpha(t) = -\frac{s^2(1-\sigma^2)}{\sigma^2\left[1-t(1-s^2)\right]}, \qquad \beta(t) = -\frac{m(1-\sigma^2)(1-t)}{\sigma^2\left[1-t(1-s^2)\right]} + \frac{\mu}{\sigma^2}. \tag{28}$$

Rolling out this updated control gives the affine Itô process

$$\mathrm{d}X_t = \left[\alpha(t)X_t + \beta(t)\right]\mathrm{d}t + \mathrm{d}W_t. \tag{29}$$

The terminal distribution generated by Eq. 29 is derived in the following section (Appendix A.1).

A.1 Moment Equations for the Affine Controlled Process

We derive the terminal distribution generated by the affine controlled process

$$\mathrm{d}X_t = \left[\alpha(t)X_t + \beta(t)\right]\mathrm{d}t + \mathrm{d}W_t, \qquad X_0 = 0, \tag{30}$$

where $\alpha(t)$ and $\beta(t)$ are defined in Eq. 28. Because Eq. 30 is linear with additive Gaussian noise, X_t remains Gaussian for all t. Let

$$M(t) := \mathbb{E}[X_t], \qquad V(t) := \text{Var}(X_t). \tag{31}$$

The mean and variance satisfy Särkkä & Solin (2019):

$$\dot{M}(t) = \alpha(t)M(t) + \beta(t), \qquad M(0) = 0, \tag{32a}$$

$$\dot{V}(t) = 2\alpha(t)V(t) + 1, \qquad V(0) = 0. \tag{32b}$$

Define

$$\psi(t) := 1 - t(1-s^2), \qquad \gamma := \frac{s^2(1-\sigma^2)}{\sigma^2(1-s^2)}. \tag{33}$$

We assume $s^2 \neq 1$ in the derivation below; the result for $s^2 = 1$ follows by continuity. With this notation,

$$\alpha(t) = -\gamma\frac{1-s^2}{\psi(t)}, \qquad \exp\left(\int_0^t \alpha(\tau)\,\mathrm{d}\tau\right) = \psi^{\gamma}(t). \tag{34}$$

The solutions of Eqs. 32 are therefore

$$M(t) = \psi^{\gamma}(t) \int_0^t \psi^{-\gamma}(\tau)\beta(\tau)\,\mathrm{d}\tau, \tag{35a}$$

$$V(t) = \psi^{2\gamma}(t) \int_0^t \psi^{-2\gamma}(\tau)\,\mathrm{d}\tau. \tag{35b}$$

For later use, define

$$I_r := \int_0^1 \psi^{-r}(\tau)\,\mathrm{d}\tau. \tag{36}$$

Direct integration gives

$$I_r = \begin{cases} \dfrac{1}{1-s^2}\log\dfrac{1}{s^2}, & r = 1, \\ \dfrac{1-(s^2)^{1-r}}{(1-s^2)(1-r)}, & r \neq 1. \end{cases} \tag{37}$$

We also use

$$\int_0^1 \frac{1-\tau}{\psi^{\gamma+1}(\tau)}\,\mathrm{d}\tau = \frac{I_\gamma - s^2 I_{\gamma+1}}{1-s^2}, \tag{38}$$

which follows from $1-\tau = [\psi(\tau) - s^2]/(1-s^2)$.

A.1.1 Terminal Mean

From Eqs. 28 and 35, the terminal mean is

$$M(1) = (s^2)^{\gamma}\left[-\frac{m(1-\sigma^2)}{\sigma^2}\frac{I_\gamma - s^2 I_{\gamma+1}}{1-s^2} + \frac{\mu}{\sigma^2} I_\gamma\right]. \tag{39}$$

The definition of γ implies

$$\frac{1-\sigma^2}{\sigma^2} = \gamma\frac{1-s^2}{s^2}, \qquad \frac{1}{\sigma^2} = \frac{s^2 + \gamma(1-s^2)}{s^2}. \tag{40}$$

When $\gamma \neq 1$, substituting Eq. 37 into Eq. 39 and simplifying gives

$$M(1) - \mu = A_m(m - \mu), \tag{41}$$

where

$$A_m := \frac{(\gamma-1)(1-s^{2\gamma}) - \gamma(1-s^{2(\gamma-1)})}{(1-s^2)(\gamma-1)}. \tag{42}$$

Equivalently, the coefficients of m and μ in $M(1)$ are A_m and $1 - A_m$, respectively.

When $\gamma = 1$, Eq. 37 instead gives

$$M(1) - \mu = \left(1 + \frac{\log s^2}{1-s^2}\right)(m - \mu). \tag{43}$$

For $s^2 \neq 1$, the condition $\gamma = 1$ is equivalent to $s^2 = \sigma^2$. Thus, when the current and target variances agree, the mean error is multiplied by

$$A_m(s^2) = 1 + \frac{\log s^2}{1-s^2}. \tag{44}$$

If $A_m(s^2) < -1$, every nonzero mean error alternates in sign and grows geometrically under repeated exact Adjoint Matching updates.

A.1.2 TERMINAL VARIANCE

From Eq. 35,

$$V(1) = (s^2)^{2\gamma} I_{2\gamma}. \tag{45}$$

When $\gamma \neq 1/2$, Eq. 37 yields

$$V(1) = \frac{s^2 - s^{4\gamma}}{(1-s^2)(2\gamma-1)}. \tag{46}$$

When $\gamma = 1/2$, the logarithmic branch of Eq. 37 gives

$$V(1) = \frac{s^2}{1-s^2} \log \frac{1}{s^2}. \tag{47}$$

Identifying the updated terminal parameters as

$$m_+ := M(1), \qquad s_+^2 := V(1), \tag{48}$$

Eqs. 41–47 give the update map stated in Theorem 3.3. In particular, s_+^2 depends only on s^2 and σ^2, whereas the mean error satisfies

$$m_+ - \mu = A_m(s^2; \sigma^2)(m - \mu). \tag{49}$$

The variance dynamics are therefore autonomous, while the mean dynamics are slaved to the evolving variance.

A.2 MULTIVARIATE GAUSSIAN ADJOINT MATCHING UPDATE

Corollary A.1. *Let $p^{tar} = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $p^{cur} = \mathcal{N}(\mathbf{m}, \mathbf{S})$, where $\boldsymbol{\Sigma}$ and $\mathbf{S}$ are positive-definite covariance matrices. Define*

$$\mathbf{C}_t = \left[\frac{1}{1-t}\mathbf{I} + \mathbf{S}^{-1} - \mathbf{I}\right]^{-1}, \qquad \mathbf{a}_t(\mathbf{x}_t) = \mathbf{C}_t \left[\frac{\mathbf{x}_t}{1-t} + \mathbf{S}^{-1}\mathbf{m}\right]. \tag{50}$$

Then the reciprocal conditional distribution and exact Adjoint Matching update are

$$p_{1|t}^{rec}(\mathbf{X}_1|\mathbf{X}_t) = \mathcal{N}\left(\mathbf{X}_1; \mathbf{a}_t(\mathbf{X}_t), \mathbf{C}_t\right), \tag{51}$$

and

$$\mathbf{u}_{G,G}(\mathbf{X}_t, t) = \left(\mathbf{I} - \boldsymbol{\Sigma}^{-1}\right) \mathbf{a}_t(\mathbf{X}_t) + \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}. \tag{52}$$

Proof of Corollary A.1. For the reference Wiener process, the conditional terminal distribution is

$$p_{1|t}^{\text{ref}}(\mathbf{x}_1|\mathbf{x}_t) = \mathcal{N}\left(\mathbf{x}_1; \mathbf{x}_t, (1-t)\mathbf{I}\right), \tag{53}$$

and its terminal marginal is $p^{\text{ref}}(\mathbf{x}_1) = \mathcal{N}(\mathbf{x}_1; \mathbf{0}, \mathbf{I})$. By the definition of reciprocal projection,

$$p_{1|t}^{\text{rec}}(\mathbf{x}_1|\mathbf{x}_t) \propto p_{1|t}^{\text{ref}}(\mathbf{x}_1|\mathbf{x}_t) \frac{p^{\text{cur}}(\mathbf{x}_1)}{p_1^{\text{ref}}(\mathbf{x}_1)}. \tag{54}$$

Substituting the three Gaussian densities and retaining only terms that depend on $\mathbf{x}_1$ gives

$$\log p_{1|t}^{\text{rec}}(\mathbf{x}_1|\mathbf{x}_t) = -\frac{1}{2}\mathbf{x}_1^\top \left[\frac{1}{1-t}\mathbf{I} + \mathbf{S}^{-1} - \mathbf{I}\right] \mathbf{x}_1 + \mathbf{x}_1^\top \left[\frac{\mathbf{x}_t}{1-t} + \mathbf{S}^{-1}\mathbf{m}\right] + \text{const.} \tag{55}$$

Completing the square therefore yields

$$p_{1|t}^{\text{rec}}(\mathbf{X}_1|\mathbf{X}_t = \mathbf{x}_t) = \mathcal{N}\left(\mathbf{X}_1; \mathbf{a}_t(\mathbf{x}_t), \mathbf{C}_t\right), \tag{56}$$

where

$$\mathbf{C}_t^{-1} = \frac{1}{1-t}\mathbf{I} + \mathbf{S}^{-1} - \mathbf{I}, \qquad \mathbf{a}_t(\mathbf{x}_t) = \mathbf{C}_t \left[\frac{\mathbf{x}_t}{1-t} + \mathbf{S}^{-1}\mathbf{m}\right]. \tag{57}$$

Notice that $\mathbf{C}_t^{-1} = \mathbf{S}^{-1} + t(1-t)^{-1}\mathbf{I}$ is positive definite for $0 \leq t < 1$.

For the Gaussian target and reference distributions, the terminal-cost gradient is

$$\boldsymbol{\nabla}_{\mathbf{x}_1} \mathbf{g}(\mathbf{x}_1) = \nabla_{\mathbf{x}_1} \log \frac{p_1^{\text{ref}}(\mathbf{x}_1)}{p^{\text{tar}}(\mathbf{x}_1)} = \left(\boldsymbol{\Sigma}^{-1} - \mathbf{I}\right) \mathbf{x}_1 - \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}. \tag{58}$$

Using the exact Adjoint Matching update and the conditional mean derived above,

$$\mathbf{u}_{\text{G,G}}(\mathbf{x}_t, t) = -\mathbb{E}_{p_{1|t}^{\text{rec}}(\mathbf{X}_1|\mathbf{X}_t = \mathbf{x}_t)}\left[\mathbf{g}(\mathbf{X}_1)\right] = \left(\mathbf{I} - \boldsymbol{\Sigma}^{-1}\right) \mathbf{a}_t(\mathbf{x}_t) + \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}, \tag{59}$$

which proves the result. □

B Asymptotic Analysis of Missing-Mode Discovery

We provide the derivation of Theorem 3.5 and the asymptotic estimate in Corollary 3.6. We specialize to the setting of Eq. 9,

$$p^{\text{tar}}(x) = \frac{1}{2}\mathcal{N}(x; -\nu, \sigma^2) + \frac{1}{2}\mathcal{N}(x; \nu, \sigma^2), \qquad p^{\text{cur}}(x) = \mathcal{N}(x; \nu, \sigma^2), \tag{60}$$

with $\varepsilon = \sigma/\nu \ll 1$.

From the Gaussian reciprocal projection derived in Appendix A, define

$$\psi(t) := 1 - t(1-\sigma^2), \qquad M_t(x_t) := \frac{\sigma^2 x_t + (1-t)\nu}{\psi(t)}, \qquad V_t := \frac{\sigma^2(1-t)}{\psi(t)}. \tag{61}$$

The reciprocal conditional distribution is

$$p^{\text{rec}}(X_1 | X_t = x_t) = \mathcal{N}\left(X_1; M_t(x_t), V_t\right). \tag{62}$$

B.1 Exponential Closeness to the Single-Mode Control

Let u_{GMM} denote the exact AM control induced by the Gaussian-mixture target, and let u_{G} denote the control induced by the corresponding right-mode Gaussian target

$$p_{\text{G}}(x) = \mathcal{N}(x; \nu, \sigma^2). \tag{63}$$

These controls are

$$u_{\text{GMM}}(x_t, t) = \mathbb{E}_{p^{\text{rec}}_{1|t}(X_1|X_t=x_t)}[-g_{\text{GMM}}(X_1)], \qquad u_{\text{G}}(x_t, t) = \mathbb{E}_{p^{\text{rec}}_{1|t}(X_1|X_t=x_t)}[-g_{\text{G}}(X_1)]. \tag{64}$$

Define

$$\Delta\nabla_x g(x) := g_{\text{GMM}}(x) - g_{\text{G}}(x) = \frac{\nu}{\sigma^2}\left[1 - \tanh\left(\frac{x}{\nu\varepsilon^2}\right)\right]. \tag{65}$$

The difference between the controls satisfies

$$|u_{\text{GMM}}(x_t, t) - u_{\text{G}}(x_t, t)| = \left|\mathbb{E}_{p^{\text{rec}}_{1|t}(X_1|X_t=x_t)}[\Delta\nabla_x g(X_1)]\right|. \tag{66}$$

For $x \geq 0$, the inequality $1 - \tanh y \leq 2e^{-2y}$ gives the local bound

$$|\Delta\nabla_x g(x)| \leq \frac{2}{\nu\varepsilon^2}\exp\left(-\frac{2x}{\nu\varepsilon^2}\right). \tag{67}$$

For all $x \in \mathbb{R}$, we also have the global bound

$$|\Delta\nabla_x g(x)| \leq \frac{2}{\nu\varepsilon^2}. \tag{68}$$

Choose a constant δ satisfying

$$\varepsilon^2 \ll \delta < 1 \tag{69}$$

and divide the terminal domain at

$$x_d := \nu\delta. \tag{70}$$

Using this decomposition,

$$\begin{aligned} |u_{\text{GMM}}(x_t, t) - u_{\text{G}}(x_t, t)| \leq{}& \left|\mathbb{E}^{\text{rec}}_{1|t}\left[\Delta\nabla_x g(X_1)\mathbf{1}_{\{X_1 < x_d\}} \mid X_t = x_t\right]\right| \\ &+ \left|\mathbb{E}^{\text{rec}}_{1|t}\left[\Delta\nabla_x g(X_1)\mathbf{1}_{\{X_1 \geq x_d\}} \mid X_t = x_t\right]\right|. \end{aligned} \tag{71}$$

On $X_1 \geq x_d$, Eq. 67 gives

$$\left|\mathbb{E}^{\text{rec}}_{1|t}\left[\Delta\nabla_x g(X_1)\mathbf{1}_{\{X_1 \geq x_d\}} \mid X_t = x_t\right]\right| \leq \frac{2}{\nu\varepsilon^2}\exp\left(-\frac{2\delta}{\varepsilon^2}\right). \tag{72}$$

To control the complementary domain, let $z > 0$ and define the safe region

$$\mathcal{S}_t(z, \delta) := \left\{ x_t : M_t(x_t) - \nu\delta \geq \frac{z}{\varepsilon}\sqrt{V_t} \right\}. \tag{73}$$

For $x_t \in \mathcal{S}_t(z, \delta)$, the cutoff x_d lies at least z/ε standard deviations below the mean of the reciprocal conditional. Consequently,

$$\mathbb{P}^{\text{rec}}_{1|t}\{X_1 < x_d \mid X_t = x_t\} \leq \frac{1}{2}\operatorname{erfc}\left(\frac{z}{\sqrt{2}\varepsilon}\right). \tag{74}$$

Combining this estimate with Eq. 68 gives

$$\left|\mathbb{E}^{\text{rec}}_{1|t}\left[\Delta\nabla_x g(X_1)\mathbf{1}_{\{X_1 < x_d\}} \mid X_t = x_t\right]\right| \leq \frac{1}{\nu\varepsilon^2}\operatorname{erfc}\left(\frac{z}{\sqrt{2}\varepsilon}\right). \tag{75}$$

Using the large-argument expansion of the complementary error function,

$$\operatorname{erfc}\left(\frac{z}{\sqrt{2}\varepsilon}\right) \sim \frac{\sqrt{2}\varepsilon}{z\sqrt{\pi}}\exp\left(-\frac{z^2}{2\varepsilon^2}\right), \tag{76}$$

we obtain

$$\left|\mathbb{E}^{\text{rec}}_{1|t}\left[\Delta\nabla_x g(X_1)\mathbf{1}_{\{X_1 < x_d\}} \mid X_t = x_t\right]\right| \lesssim \frac{1}{\nu\varepsilon}\exp\left(-\frac{z^2}{2\varepsilon^2}\right). \tag{77}$$

Therefore, throughout the safe region,

$$|u_{\text{GMM}}(x_t, t) - u_{\text{G}}(x_t, t)| \lesssim \frac{1}{\nu\varepsilon}\exp\left(-\frac{z^2}{2\varepsilon^2}\right) + \frac{2}{\nu\varepsilon^2}\exp\left(-\frac{2\delta}{\varepsilon^2}\right). \tag{78}$$

B.2 Self-Consistency of the Safe Region

It remains to show that the forward process driven by u_{G} remains in $\mathcal{S}_t(z, \delta)$ with exponentially high probability. Because the current distribution already agrees with the right Gaussian target, the single-mode AM control is

$$u_{\text{G}}(x_t, t) = -\frac{1 - \sigma^2}{\psi(t)}x_t + \frac{\nu}{\psi(t)}. \tag{79}$$

The corresponding affine process has marginal distribution

$$X_t \sim \mathcal{N}\left(t\nu, t\psi(t)\right). \tag{80}$$

Let $z_c > 0$ and define the lower forward-tail boundary

$$b_t := t\nu - \frac{z_c}{\varepsilon}\sqrt{t\psi(t)}. \tag{81}$$

The probability of falling below this boundary is

$$\mathbb{P}_{\text{G}}\{X_t < b_t\} = \frac{1}{2}\operatorname{erfc}\left(\frac{z_c}{\sqrt{2}\varepsilon}\right) \lesssim \frac{\varepsilon}{z_c}\exp\left(-\frac{z_c^2}{2\varepsilon^2}\right). \tag{82}$$

Because $M_t(x_t)$ is increasing in x_t, the interval $[b_t, \infty)$ is contained in the safe region whenever

$$M_t(b_t) - \nu\delta \geq \frac{z}{\varepsilon}\sqrt{V_t}. \tag{83}$$

Substitution of Eqs. 61 and 81 reduces this condition to

$$(1 - \delta)\sqrt{\psi(t)} \geq z_c\sigma\sqrt{t} + z\sqrt{1 - t}. \tag{84}$$

By the Cauchy–Schwarz inequality,

$$z_c\sigma\sqrt{t} + z\sqrt{1 - t} \leq \sqrt{z_c^2 + z^2}\sqrt{t\sigma^2 + 1 - t} = \sqrt{z_c^2 + z^2}\sqrt{\psi(t)}. \tag{85}$$

A sufficient condition, uniform in time, is therefore

$$z^2 + z_c^2 \leq (1 - \delta)^2. \tag{86}$$

A convenient time-uniform allocation is

$$z = z_c = \frac{1-\delta}{\sqrt{2}}. \tag{87}$$

Equations 78 and 82 then give

$$|u_{\mathrm{GMM}}(x_t,t) - u_{\mathrm{G}}(x_t,t)| \lesssim \frac{1}{\nu\varepsilon}\exp\left[-\frac{(1-\delta)^2}{4\varepsilon^2}\right] + \frac{2}{\nu\varepsilon^2}\exp\left(-\frac{2\delta}{\varepsilon^2}\right) \tag{88}$$

for all $x_t \in \mathcal{S}_t$, while

$$\mathbb{P}_{\mathrm{G}}\{X_t \notin \mathcal{S}_t\} \lesssim \frac{\varepsilon}{1-\delta}\exp\left[-\frac{(1-\delta)^2}{4\varepsilon^2}\right]. \tag{89}$$

In particular,

$$|u_{\mathrm{GMM}}(x_t,t) - u_{\mathrm{G}}(x_t,t)| = \mathcal{O}\left(\mathrm{poly}(\varepsilon^{-1})e^{-\mathcal{C}\varepsilon^{-2}}\right), \tag{90}$$

where

$$\mathcal{C} := \min\left\{\frac{(1-\delta)^2}{4}, 2\delta\right\} > 0. \tag{91}$$

This proves a time-uniform, self-consistent regime in which the bimodal-target control is exponentially close to the single-mode affine control on a region carrying all but exponentially small forward probability.

For completeness, a time-dependent allocation that saturates Eq. 84 is

$$z(t) = (1-\delta)\frac{\sqrt{1-t}}{\sqrt{\psi(t)}}, \qquad z_c(t) = (1-\delta)\frac{\sigma\sqrt{t}}{\sqrt{\psi(t)}}. \tag{92}$$

These parameters satisfy

$$z^2(t) + z_c^2(t) = (1-\delta)^2. \tag{93}$$

This allocation makes the tradeoff explicit: reciprocal-tail concentration dominates at early time, whereas forward-process concentration dominates near terminal time. The apparent deterioration of the reciprocal-tail bound as $t \uparrow 1$ reflects the collapse $V_t \to 0$ of the reciprocal conditional; near terminal time, the relevant estimate is instead the forward-process bound controlled by $z_c(t)$.

B.3 Exponentially Small Missing-Basin Probability

The preceding analysis shows that, on trajectories carrying all but exponentially small probability, the control induced by the bimodal target is exponentially close to the single-mode control u_{G}. At leading single-mode order, the terminal distribution therefore remains

$$X_1^{\mathrm{G}} \sim \mathcal{N}(\nu, \sigma^2). \tag{94}$$

Its probability of reaching the missing left basin is

$$\mathbb{P}_{\mathrm{G}}\{X_1^{\mathrm{G}} < 0\} = \frac{1}{2}\,\mathrm{erfc}\left(\frac{\nu}{\sqrt{2}\sigma}\right) = \frac{1}{2}\,\mathrm{erfc}\left(\frac{1}{\sqrt{2}\varepsilon}\right). \tag{95}$$

As $\varepsilon \downarrow 0$,

$$\mathbb{P}_{\mathrm{G}}\{X_1^{\mathrm{G}} < 0\} \sim \frac{\varepsilon}{\sqrt{2\pi}}\exp\left(-\frac{1}{2\varepsilon^2}\right). \tag{96}$$

Thus, the inter-basin mass present under the leading single-mode dynamics is exponentially small in the squared mode separation.

If Eq. 96 is interpreted as the characteristic probability mass made available to the missing basin per outer iteration, its inverse gives

$$\tau_{\mathrm{disc}} \sim \frac{\sqrt{2\pi}}{\varepsilon}\exp\left(\frac{1}{2\varepsilon^2}\right). \tag{97}$$

This inverse-flux interpretation is analogous to WKB estimates in metastable stochastic systems, where exponentially small probability flux produces exponentially long escape times (Doering et al., 2005; Assaf & Meerson, 2010). We emphasize that Eq. 97 is a characteristic asymptotic estimate, not a global mixing-time theorem for repeated nonlinear AM iterations. Once the second mode acquires appreciable probability mass, the two-mixture dynamics analyzed in Section C becomes relevant.

C Reciprocal Projection and Weight Equilibration for Gaussian Mixtures

We derive the exact reciprocal-mixture decomposition in Proposition 3.7 and then examine its well-separated limit. Consider

$$p^{\text{tar}}(x) = \varpi\mathcal{N}(x; -\nu, \sigma^2) + (1-\varpi)\mathcal{N}(x; \nu, \sigma^2), \tag{98a}$$

$$p^{\text{cur}}(x) = w\mathcal{N}(x; -m, s^2) + (1-w)\mathcal{N}(x; m, s^2). \tag{98b}$$

C.1 Target Score and Its Transition Layer

Expanding the target density gives

$$p^{\text{tar}}(x) \propto \varpi \exp\left[-\frac{(x+\nu)^2}{2\sigma^2}\right] + (1-\varpi)\exp\left[-\frac{(x-\nu)^2}{2\sigma^2}\right]. \tag{99}$$

Direct differentiation yields

$$\nabla_x \log p^{\text{tar}}(x) = -\frac{x}{\sigma^2} - \frac{\nu}{\sigma^2}\tanh\left(-\frac{x\nu}{\sigma^2} + \frac{1}{2}\log\frac{\varpi}{1-\varpi}\right) \tag{100}$$

$$= -\frac{x}{\sigma^2} + \frac{\nu}{\sigma^2}\tanh\left(\frac{x - x_c^{\text{tar}}}{\nu\varepsilon^2}\right), \tag{101}$$

where

$$x_c^{\text{tar}} = \frac{\sigma^2}{2\nu}\log\frac{\varpi}{1-\varpi}. \tag{102}$$

Including the reference score gives

$$-g_{\text{GMM}}(x) = x - \frac{x}{\sigma^2} + \frac{\nu}{\sigma^2}\tanh\left(\frac{x - x_c^{\text{tar}}}{\nu\varepsilon^2}\right). \tag{103}$$

If ϖ is bounded away from zero and one, then

$$\frac{x_c^{\text{tar}}}{\nu} = \frac{\varepsilon^2}{2}\log\frac{\varpi}{1-\varpi} = \mathcal{O}(\varepsilon^2). \tag{104}$$

Thus, in the well-separated limit,

$$-g_{\text{GMM}}(x) \sim \begin{cases} x - \dfrac{x+\nu}{\sigma^2}, & x < x_c^{\text{tar}}, \\ x - \dfrac{x-\nu}{\sigma^2}, & x > x_c^{\text{tar}}. \end{cases} \tag{105}$$

The target weight affects the score through the location and shape of the narrow transition layer between these two componentwise Gaussian scores.

C.2 Exact Reciprocal-Mixture Decomposition

Introduce the latent component variable $K \in \{L, R\}$, with

$$\mathbb{P}\{K = L\} = w, \qquad \mathbb{P}\{K = R\} = 1 - w, \tag{106}$$

and component densities

$$p_L(x) = \mathcal{N}(x; -m, s^2), \qquad p_R(x) = \mathcal{N}(x; m, s^2). \tag{107}$$

The reciprocal conditional satisfies

$$p_{1|t}^{\text{rec}}(x_1, K|x_t) \propto \mathbb{P}\{K\}p_K(x_1)\frac{p_{1|t}^{\text{ref}}(x_1|x_t)}{p_1^{\text{ref}}(x_1)}. \tag{108}$$

Define

$$\mathcal{Z}_K(x_t, t) := \int p_K(x_1)\frac{p^{\text{ref}}(x_1|x_t)}{p^{\text{ref}}(x_1)}\,\mathrm{d}x_1. \tag{109}$$

Conditioned on K, the reciprocal distribution is

$$p^{\text{rec}}_{1|t}(x_1|x_t, K) = \frac{1}{\mathcal{Z}_K(x_t,t)} p_K(x_1) \frac{p^{\text{ref}}_{1|t}(x_1|x_t)}{p^{\text{ref}}_1(x_1)}. \tag{110}$$

Each conditional has the same form as the Gaussian reciprocal distribution derived in Appendix A. Writing

$$\psi_s(t) := 1 - t(1 - s^2), \tag{111}$$

we obtain

$$p^{\text{rec}}_{1|t}(X_1|X_t = x_t, L) = \mathcal{N}\left(X_1; \frac{s^2 x_t - (1-t)m}{\psi_s(t)}, \frac{s^2(1-t)}{\psi_s(t)}\right), \tag{112a}$$

$$p^{\text{rec}}_{1|t}(X_1|X_t = x_t, R) = \mathcal{N}\left(X_1; \frac{s^2 x_t + (1-t)m}{\psi_s(t)}, \frac{s^2(1-t)}{\psi_s(t)}\right). \tag{112b}$$

Completing the square in the two normalization integrals gives

$$\frac{\mathcal{Z}_R(x_t,t)}{\mathcal{Z}_L(x_t,t)} = \exp\left(\frac{2mx_t}{\psi_s(t)}\right). \tag{113}$$

Marginalizing over the latent component therefore yields

$$p^{\text{rec}}_{1|1}(x_1|x_t) = r_L(x_t,t) p^{\text{rec}}_{1|t}(x_1|x_t, L) + r_R(x_t,t) p^{\text{rec}}_{1|t}(x_1|x_t, R), \tag{114}$$

where

$$r_L(x_t,t) = \frac{w\mathcal{Z}_L}{w\mathcal{Z}_L + (1-w)\mathcal{Z}_R} = \frac{w}{w + (1-w)e^{2mx_t/\psi_s(t)}}, \tag{115a}$$

$$r_R(x_t,t) = \frac{(1-w)\mathcal{Z}_R}{w\mathcal{Z}_L + (1-w)\mathcal{Z}_R} = \frac{(1-w)e^{2mx_t/\psi_s(t)}}{w + (1-w)e^{2mx_t/\psi_s(t)}}. \tag{115b}$$

Define the componentwise controls

$$u_L(x_t,t) := \mathbb{E}_{\text{rec}}[-g_{\text{GMM}}(X_1) \mid X_t = x_t, L], \tag{116}$$

and

$$u_R(x_t,t) := \mathbb{E}_{\text{rec}}[-g_{\text{GMM}}(X_1) \mid X_t = x_t, R]. \tag{117}$$

Linearity of conditional expectation gives the exact decomposition

$$u(x_t,t) = r_L(x_t,t) u_L(x_t,t) + r_R(x_t,t) u_R(x_t,t). \tag{118}$$

Let

$$\theta(x_t,t) := \frac{mx_t}{\psi_s(t)}. \tag{119}$$

Using $r_L + r_R = 1$, Eq. 118 can be written as

$$u(x_t,t) = \frac{u_L + u_R}{2} + \frac{u_R - u_L}{2} \tanh\left(\theta + \frac{1}{2}\log\frac{1-w}{w}\right). \tag{120}$$

Defining

$$x_c^{\text{cur}}(t) := \frac{\psi_s(t)}{2m} \log\frac{w}{1-w}, \tag{121}$$

we recover

$$u(x_t,t) = \frac{u_L(x_t,t) + u_R(x_t,t)}{2} + \frac{u_R(x_t,t) - u_L(x_t,t)}{2} \tanh\left[\frac{m}{\psi_s(t)}\left(x_t - x_c^{\text{cur}}(t)\right)\right]. \tag{122}$$

This completes the proof of Proposition 3.7. Notice that no well-separated approximation has been used.

C.3 Asymptotic Insensitivity to the Target Weight

We now specialize to

$$m = \nu, \qquad s^2 = \sigma^2, \tag{123}$$

so that the current mixture has the correct component locations and variances but may have an incorrect weight $w \neq \varpi$.

Define the left- and right-component Gaussian score differences

$$-g_L^{\mathrm{G}}(x) = x - \frac{x+\nu}{\sigma^2}, \tag{124a}$$

$$-g_R^{\mathrm{G}}(x) = x - \frac{x-\nu}{\sigma^2}. \tag{124b}$$

Equation 105 shows that $-g_{\mathrm{GMM}}$ agrees with $-g_L^{\mathrm{G}}$ throughout the left mode and with $-g_R^{\mathrm{G}}$ throughout the right mode, except near the inter-mode transition layer.

More explicitly,

$$[-g_{\mathrm{GMM}}(x)] - [-g_R^{\mathrm{G}}(x)] = \frac{\nu}{\sigma^2}\left[\tanh\left(\frac{x - x_c^{\mathrm{tar}}}{\nu\varepsilon^2}\right) - 1\right], \tag{125a}$$

$$[-g_{\mathrm{GMM}}(x)] - [-g_L^{\mathrm{G}}(x)] = \frac{\nu}{\sigma^2}\left[\tanh\left(\frac{x - x_c^{\mathrm{tar}}}{\nu\varepsilon^2}\right) + 1\right]. \tag{125b}$$

For the right component, the score correction is appreciable only near or to the left of x_c^{tar}; for the left component, it is appreciable only near or to the right of x_c^{tar}. Because $x_c^{\mathrm{tar}}/\nu = \mathcal{O}(\varepsilon^2)$, either region lies $\mathcal{O}(\varepsilon^{-1})$ standard deviations from the corresponding component center.

Applying the domain-decomposition and Gaussian-tail argument of Appendix B separately to the left and right reciprocal conditionals gives

$$u_L(x_t, t) = u_L^{\mathrm{G}}(x_t, t) + \mathcal{O}\left(\mathrm{poly}(\varepsilon^{-1})e^{-c\varepsilon^{-2}}\right), \tag{126a}$$

$$u_R(x_t, t) = u_R^{\mathrm{G}}(x_t, t) + \mathcal{O}\left(\mathrm{poly}(\varepsilon^{-1})e^{-c\varepsilon^{-2}}\right), \tag{126b}$$

on regions carrying all but exponentially small probability, for some $c > 0$. The constant c can be chosen uniformly when w and ϖ remain in a compact subset of $(0, 1)$.

Substituting Eq. 126 into the exact decomposition gives

$$u(x_t, t) = r_L(x_t, t)u_L^{\mathrm{G}}(x_t, t) + r_R(x_t, t)u_R^{\mathrm{G}}(x_t, t) + \mathcal{O}\left(\mathrm{poly}(\varepsilon^{-1})e^{-c\varepsilon^{-2}}\right). \tag{127}$$

The leading term depends on w through the exact reciprocal responsibilities r_L and r_R, but contains no dependence on the target weight ϖ. The target-weight information is confined to the exponentially small remainder generated by the overlap region.

At leading order, u_L^{G} and u_R^{G} preserve their corresponding Gaussian components, while r_L and r_R recombine them according to the current mixture weight w. Hence,

$$p_{\mathrm{next}}(x) = w\mathcal{N}(x; -\nu, \sigma^2) + (1-w)\mathcal{N}(x; \nu, \sigma^2) + \mathcal{O}\left(\mathrm{poly}(\varepsilon^{-1})e^{-c\varepsilon^{-2}}\right), \tag{128}$$

in the separated-mode asymptotic sense. Equivalently,

$$w_+ - w = \mathcal{O}\left(\mathrm{poly}(\varepsilon^{-1})e^{-c\varepsilon^{-2}}\right). \tag{129}$$

Therefore, every $w \in (0, 1)$ is a leading-order fixed point as $\varepsilon \downarrow 0$, and the correction toward the target weight ϖ is exponentially slow.

D Solving the Adjoint Matching Equations by Computing the Quadrature and Solving the Fokker–Planck Equation

D.1 Density-level iteration

Let $p^{(k)}$ denote the terminal density obtained at Adjoint Matching iteration k, with $p^{(0)} = p^{\mathrm{cur}}$. For the standard Wiener reference process, define

$$K(z,t) = \frac{1}{\sqrt{2\pi(1-t)}} \exp\left(-\frac{z^2}{2\,(1-t)}\right), \qquad h^{(k)}(y) = \frac{p^{(k)}(y)}{p^{\mathrm{ref}}(y)}, \tag{130}$$

where $p^{\mathrm{ref}}\,(y) = e^{-y^2/2}/\sqrt{2\pi}$. The reciprocal conditional expectation in Eq. 4 can then be evaluated as

$$u^{(k)}(x,t) = -\frac{\displaystyle\int_{-\infty}^{\infty} K(y-x,t)\, h^{(k)}(y)\, \nabla_x g(y)\, \mathrm{d}y}{\displaystyle\int_{-\infty}^{\infty} K(y-x,t)\, h^{(k)}(y)\, \mathrm{d}y}, \qquad 0 \le t < 1. \tag{131}$$

Here, $\nabla_x g = \partial_x \log(p^{\mathrm{tar}}/p^{\mathrm{ref}})$ is the full target-to-reference score difference, evaluated using the exact expression in Eq. 103.

After the velocity is computed, we solve

$$\partial_t \rho^{(k)}(x,t) = -\partial_x\left(u^{(k)}(x,t)\rho^{(k)}(x,t)\right) + \frac{1}{2}\partial_{xx}\rho^{(k)}(x,t), \qquad \rho^{(k)}(x,0) = \delta_0(x), \tag{132}$$

and set

$$p^{(k+1)}(x) = \rho^{(k)}(x,1). \tag{133}$$

Thus, each Adjoint Matching iteration starts a fresh controlled process from $X_0 = 0$; the preceding terminal density determines its velocity through the reciprocal expectation in the next Adjoint Matching step.

Only the initial terminal-density guess is prescribed as a Gaussian mixture. Subsequent terminal densities are retained as general grid-resolved functions and used directly to construct the next velocity. There is no Gaussian-mixture fitting, moment closure, or projection back onto a parametric family. The calculation therefore tests whether the analytical mechanisms persist under the full density-level iteration.

D.2 Numerical discretization

The spatial domain is truncated to $[-5, 5]$, with $N_x = 10001$ uniformly spaced points, giving $\Delta x = 10^{-3}$. The interval $t \in [0,1]$ contains $N_t = 1001$ time levels, giving $\Delta t = 10^{-3}$. Zero-flux boundary conditions are imposed at both spatial endpoints.

The numerator and denominator in Eq. 131 are evaluated by discrete Gaussian convolutions using FFT-accelerated linear convolution. To control roundoff in regions where the denominator is small, entries below 10^{-12} times its maximum are recomputed by direct Gaussian summation. The reciprocal weights are evaluated after subtracting a common maximum from their logarithms; this rescaling cancels in the ratio.

For the forward equation, the implementation uses a nearest-neighbor discretization with right and left transition rates

$$r_i = \frac{1}{2\Delta x^2} + \frac{u_i}{2\Delta x}, \qquad \ell_i = \frac{1}{2\Delta x^2} - \frac{u_i}{2\Delta x}. \tag{134}$$

Writing $P_i(t)$ for the probability mass associated with grid point x_i, the interior equations are

$$\dot{P}_i = r_{i-1}P_{i-1} - (r_i + \ell_i)P_i + \ell_{i+1}P_{i+1}. \tag{135}$$

Outward transitions at the domain boundaries are suppressed, yielding a conservative tridiagonal operator $A[u]$.

Table 1: Settings for the direct quadrature and evolution of FPE-based experiments. Panel labels correspond to Fig. 2.

Panel	Experiment	ϖ	w	m
(a)	Single-mode instability	10^{-20}	10^{-20}	1.01
(b)	Symmetric two-mode instability	$1/2$	$1/2$	1.01
(c)	Missing-mode discovery	$1/2$	10^{-20}	1
(d)	Mixture-weight equilibration	$10/11$	$1/2$	1

Time integration uses Crank–Nicolson with the velocity evaluated at each time-step midpoint:

$$\left(I - \frac{\Delta t}{2} A[u^{(k)}_{n+1/2}]\right) P_{n+1} = \left(I + \frac{\Delta t}{2} A[u^{(k)}_{n+1/2}]\right) P_n. \tag{136}$$

The discrete point-mass initial condition is treated with Rannacher smoothing: the first full time step is replaced by two backward-Euler half steps, with velocities evaluated at their respective midpoints. All linear systems are solved as tridiagonal banded systems. Before constructing the next reciprocal weight, terminal grid masses are floored at 10^{-20} to permit logarithmic evaluation.

D.3 Numerical experiments

The target and initial terminal-density guess are

$$p^{\mathrm{tar}} = \varpi\,\mathcal{N}(-\nu, \sigma^2) + (1-\varpi)\,\mathcal{N}(\nu, \sigma^2), \tag{137}$$

$$p^{(0)} = w\,\mathcal{N}(-m, s^2) + (1-w)\,\mathcal{N}(m, s^2). \tag{138}$$

All four experiments use $\nu = 1$ and $s^2 = \sigma^2 = 0.05$, corresponding to $\varepsilon = \sigma/\nu \approx 0.224$. The weights ϖ and w refer to the left component. A weight of 10^{-20} approximates the single-mode limit while keeping the mixture-score expression finite. Table 1 summarizes the four configurations. Each calculation performs six Adjoint Sampling updates; Fig. 2 displays the initial guess and the first four updates.

E Neural-Network Experiments

This section gives the numerical specification and additional results for the neural-network experiments summarized in Sec. 4.2. We first describe the common training protocol and then organize the Gaussian experiments by dimension: one, two, and 128 dimensions. Non-Gaussian tests are collected separately in Appendix F.

E.1 Common neural protocol

All neural experiments use the controlled process in Eq. 1. Direct velocity damping (Blessing et al., 2026) replaces the ordinary Adjoint Matching regression target Y_k by

$$Y_k^{(\eta)} = (1-\eta)\,u_{k-1}(X_t, t) + \eta\, Y_k, \qquad \eta = 0.5, \tag{139}$$

where the previous network is frozen during the parameter update. We combine velocity damping on/off with replay on/off, giving four algorithmic variants.

Controlled initialization. Because the controlled diffusion always starts from $X_0 = 0$, the prescribed terminal law $p^{(0)}$ must itself be represented by a control before the target-informed iteration begins. For each problem and random seed, we draw 20,000 terminal samples analytically from $p^{(0)}$, temporarily set the RAM terminal target equal to $p^{(0)}$, and perform 20,000 RAM regression updates. The resulting network defines the initial control u_0.

An independent rollout of u_0 using 10^4 terminal samples is retained only as an initialization diagnostic. These generated endpoints are not used for the first target-informed update, and the optimizer and replay state from the initialization stage are discarded.

For each of the four replay/damping variants associated with a fixed problem and seed, we make an independent copy of the same learned control u_0, construct a fresh optimizer, and create a fresh experimental replay buffer. The RAM terminal target is then changed from $p^{(0)}$ to p^{tar}, and a fresh analytic sample of 512 endpoints from $p^{(0)}$ is used for the first target-informed update. Hence velocity damping is active on this first measured update, with $u_{k-1} = u_0$. The preprocessing rollout is never inserted into the experimental replay buffer. In the reported trajectories, $k = 0$ denotes the prescribed initial terminal law and $k = 1$ the first target-informed update.

Numerical specification. Unless stated otherwise, the numerical and optimization settings are

Quantity	Value
Target-informed outer iterations	50
RAM updates per outer iteration	100
Current terminal samples per update	512
Minibatch size	64
Euler–Maruyama steps per rollout	500
Diagnostic terminal samples	10^4
Initialization terminal samples	20,000
Initialization RAM updates	20,000
Random seeds	$0, 1, 2$
Hidden layers	3
Hidden width	64
Activation	SiLU
Optimizer	Adam
Learning rate	10^{-4}
Replay capacity	1000 terminal samples
Terminal-gradient clipping threshold	150
Velocity-damping coefficient	$\eta = 0.5$

Replay stores terminal endpoints together with their corresponding terminal gradients. For a fixed problem and seed, all four variants begin from the same learned u_0 and use matched random-number streams whenever their subsequent algorithms permit it.

E.2 One-dimensional Gaussian experiments

The four target and prescribed initial terminal distributions are

Experiment	p^{tar}	$p^{(0)}$
Single Gaussian	$\mathcal{N}(1, 0.05)$	$\mathcal{N}(1.01, 0.05)$
Symmetric two-mode	$\frac{1}{2}\mathcal{N}(-1, 0.05) + \frac{1}{2}\mathcal{N}(1, 0.05)$	$\frac{1}{2}\mathcal{N}(-1.01, 0.05) + \frac{1}{2}\mathcal{N}(1.01, 0.05)$
Incorrect weights	$\frac{10}{11}\mathcal{N}(-1, 0.05) + \frac{1}{11}\mathcal{N}(1, 0.05)$	$\frac{1}{2}\mathcal{N}(-1, 0.05) + \frac{1}{2}\mathcal{N}(1, 0.05)$
Missing mode	$\frac{1}{2}\mathcal{N}(-1, 0.05) + \frac{1}{2}\mathcal{N}(1, 0.05)$	$\mathcal{N}(1, 0.05)$

The neural single-mode initializations are exact. The 10^{-20} component weights used in Sec. 4.1 are only a numerical regularization for the density-level implementation. For prescribed 50:50 initial mixtures, component counts are stratified to their nominal proportions.

Diagnostics. For multimodal diagnostics, we split the state space at zero and define

$$P_{L,k} = \mathbb{P}_k(X_1 < 0), \qquad m_{L,k} = \mathbb{E}_k[X_1 \mid X_1 < 0], \qquad m_{R,k} = \mathbb{E}_k[X_1 \mid X_1 \geq 0]. \tag{140}$$

When both basins are populated, the empirical half-separation is

$$\widehat{\nu}_k = \frac{m_{R,k} - m_{L,k}}{2}. \tag{141}$$

We additionally record terminal mean, variance, second moment, conditional basin moments, and the empirical one-dimensional Wasserstein distance W_2 to the target. With 10^4 diagnostic samples, a probability of 10^{-4} corresponds to one observed sample; an empirical zero therefore means that no diagnostic sample entered the basin, rather than that the underlying probability is mathematically zero.

Quantitative results. To summarize late-time behavior, for each seed we average the relevant diagnostic over the final 20 outer iterations and then aggregate across the three seeds. For the single-Gaussian and symmetric two-mode problems, the reported errors below are the mean absolute deviations of these seedwise averages from the corresponding target value:

Regime	Single-G. error	Separation error	Missing P_L	Weight P_L
No replay, undamped	0.123	—	0.105	0.608
Replay, undamped	0.070	0.010	0.048	0.532
No replay, damped	0.002	0.005	0.0001	0.582
Replay, damped	0.001	0.007	0.0001	0.551

The target values are 1 for the single-Gaussian mean and mode half-separation, $1/2$ for the missing-mode left-basin probability, and $10/11$ for the incorrect-weight left-basin probability. The separation error is undefined for the no-replay undamped case because one seed loses an entire basin during the late-time window.

For the single-Gaussian experiment, the first column of the table measures the error in the mean. We complement this first-moment diagnostic with the target variance, whose value is 0.05. For the damped runs, the final variances are 0.054 without replay and 0.053 with replay. We additionally report the one-dimensional quadratic Wasserstein distance. For law μ and target law $\mu_\star$, with quantile functions F^{-1} and $F_\star^{-1}$, respectively, we recall that the Wasserstein error between μ and $\mu_\star$ can be written

$$W_2(\mu, \mu_\star) = \left(\int_0^1 \left| F^{-1}(q) - F_\star^{-1}(q) \right|^2 \, \mathrm{d}q \right)^{1/2}. \tag{142}$$

The damped single-Gaussian runs have mean final-20 Wasserstein errors of approximately 0.023 without replay and 0.022 with replay. The agreement of the mean, variance, and W_2 diagnostics indicates that direct velocity damping stabilizes the full terminal law in this local Gaussian experiment.

In the symmetric two-mode problem, damping keeps the half-separation close to its target, although the corresponding final-20 Wasserstein errors remain approximately 0.25 and 0.32. Stabilization of the dominant separation coordinate therefore does not by itself imply convergence of the complete terminal distribution.

The two probability-redistribution experiments remain substantially more difficult. In the missing-mode problem, the damped variants remain near the diagnostic resolution limit; undamped variants generate larger, seed-dependent inter-basin excursions without reproducibly recovering the target. In the incorrect-weight problem, all four variants remain far from the target left weight. These observations retain the distinction predicted by the theory between local stability and global probability redistribution.

Late-time undamped missing-mode dynamics. The seed averages in Fig. 3(c) conceal substantial structure in the individual undamped trajectories. Figure 4 examines a representative seed over iterations $k = 25, \ldots, 50$, using terminal populations retained at every outer iteration. Without replay, the late trajectory exhibits an approximately period-two alternation between two distinct terminal-law branches. This separation persists at the level of the full terminal density and is therefore not merely an oscillation of the scalar basin-probability diagnostic.

Replay produces a less phase-locked recurrence, so iteration parity does not give a useful branch assignment. We instead represent each late terminal population by its empirical quantile function and apply two-means clustering to these vectors. For one-dimensional empirical measures with equal sample counts, squared Euclidean distance between the sorted sample vectors is proportional to empirical W_2^2. The resulting two clusters therefore provide a two-state description in empirical Wasserstein geometry. We refer to each recurrent terminal-law cluster as a W_2 *state*. The labels are oriented so that state A has the smaller pooled value of $P(X < 0)$ and state B the larger value.

E.3 Two-dimensional Gaussian experiments

We next test whether the same local-instability and probability-redistribution mechanisms persist beyond one dimension. The two-dimensional experiments use the common controlled-initialization,

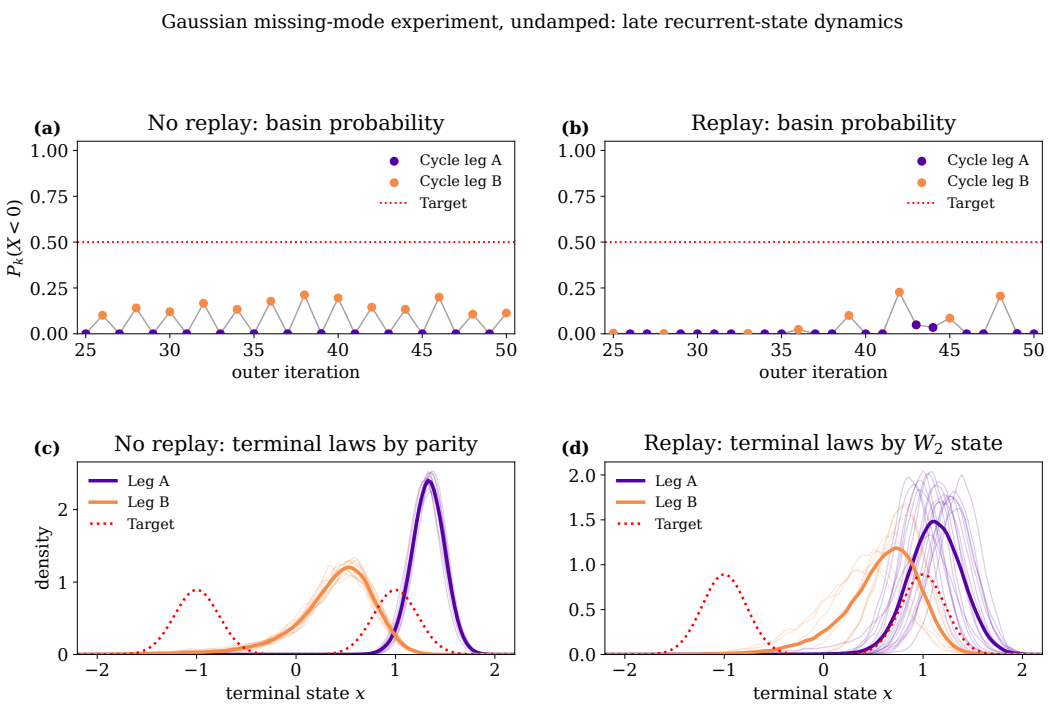


Figure 4: Late-time dynamics of a representative seed in the Gaussian missing-mode experiment with damping removed. (a) No replay: basin probability $P_k(X < 0)$ over iterations $k = 25, \ldots, 50$, with the two approximately alternating branches shown separately. (b) Replay: the same diagnostic, with late iterates colored according to their inferred W_2 state. (c) No replay: terminal densities from every displayed late iteration; faint curves are individual iterates and thick curves are obtained by pooling samples within each parity-defined branch. (d) Replay: terminal densities from every displayed late iteration, colored by W_2 state; thick curves are the corresponding pooled densities. Red dotted curves or lines denote the balanced bimodal target. The no-replay trajectory is consistent with an approximate period-two orbit, while replay perturbs this into a less phase-locked recurrent switching pattern. We conjecture that both reflect the mode-oscillation instability predicted by the reduced Gaussian dynamics, although the neural iterates are not constrained to that reduced family.

optimization, replay, damping, and diagnostic protocol of Appendix E.1. We first examine two direct extensions of the one-dimensional oscillation experiments and then consider mode recovery and weight redistribution on a 3×3 Gaussian grid.

E.3.1 Coordinatewise Gaussian stability

Write the two-dimensional terminal state as $(X, Y) \in \mathbb{R}^2$. Consider the anisotropic target and prescribed initial terminal law

$$p^{\mathrm{tar}} = \mathcal{N}\left(\begin{bmatrix}1\\1\end{bmatrix}, \begin{bmatrix}0.05 & 0\\0 & 0.5\end{bmatrix}\right), \qquad p^{(0)} = \mathcal{N}\left(\begin{bmatrix}1.01\\1.01\end{bmatrix}, \begin{bmatrix}0.05 & 0\\0 & 0.5\end{bmatrix}\right). \tag{143}$$

For diagonal covariance and matching current and target covariances, the exact multivariate Gaussian update of Appendix A.2 decouples across the coordinate axes. Each coordinate therefore obeys the one-dimensional mean-error multiplier

$$A(s^2) = 1 + \frac{\log s^2}{1 - s^2}. \tag{144}$$

For the two axes in Eq. equation 143,

$$A(0.05) \approx -2.153, \qquad A(0.5) \approx -0.386. \tag{145}$$

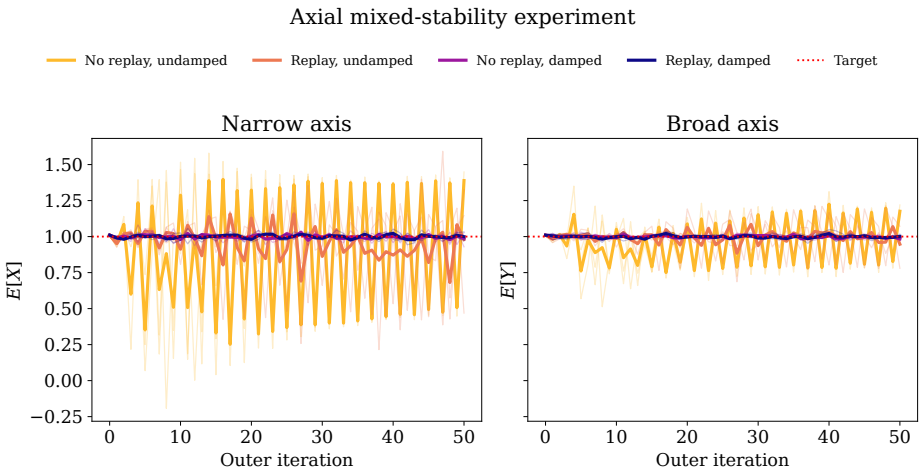


Figure 5: Two-dimensional anisotropic single-Gaussian experiment. The panels show the terminal coordinate means $\mathbb{E}_k[X]$ and $\mathbb{E}_k[Y]$ for the narrow ($\sigma_X^2 = 0.05$) and broad ($\sigma_Y^2 = 0.5$) axes, respectively. Thin curves show individual seeds and thick curves show three-seed means; red dotted lines denote the target coordinate mean 1. The exact Gaussian update decouples across these axes and predicts multipliers $A(0.05) \approx -2.153$ and $A(0.5) \approx -0.386$.

Thus the exact theory predicts an alternating, growing mean error along the narrow coordinate and an alternating, contracting mean error along the broad coordinate.

Figure 5 preserves the ordering predicted by the exact calculation: the broad coordinate is substantially more stable than the narrow coordinate. The learned dynamics, however, no longer exhibit exact coordinatewise decoupling. Approximate representation of the initialized control, partial inner optimization, finite terminal samples, stochastic regression, and the shared neural-network parameterization all modify the exact Gaussian map. We conjecture that these effects contribute to the residual oscillatory motion visible in the broad coordinate, allowing the instability of the narrow direction to leak into the jointly learned control. Its smaller amplitude remains consistent with the greater coordinatewise stability predicted by Eq. equation 145.

E.3.2 Diagonal two-mode oscillation

We next consider an isotropic symmetric two-mode mixture whose component centers lie along

$$e_{\parallel} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \tag{146}$$

The target has component means $\pm e_{\parallel}$, covariance $0.05 I_2$, and equal weights; the prescribed initial law has means $\pm 1.01 e_{\parallel}$ with the same covariance and weights. Define the projected terminal coordinate

$$S = e_{\parallel}^{\mathsf{T}} \begin{bmatrix} X \\ Y \end{bmatrix} = \frac{X+Y}{\sqrt{2}}. \tag{147}$$

When both projected basins are represented, let

$$m_{-,k} = \mathbb{E}_k[S \mid S < 0], \qquad m_{+,k} = \mathbb{E}_k[S \mid S > 0], \tag{148}$$

and define the projected half-separation

$$\widehat{\nu}_k^{\parallel} = \frac{m_{+,k} - m_{-,k}}{2}. \tag{149}$$

Its target value is 1.

The complete trajectory in Fig. 6 exposes an oscillatory mode-location instability directly analogous to the one-dimensional symmetric two-mode experiment. The undefined values in the other no-replay undamped seeds indicate a more severe form of the same practical failure: at the resolution

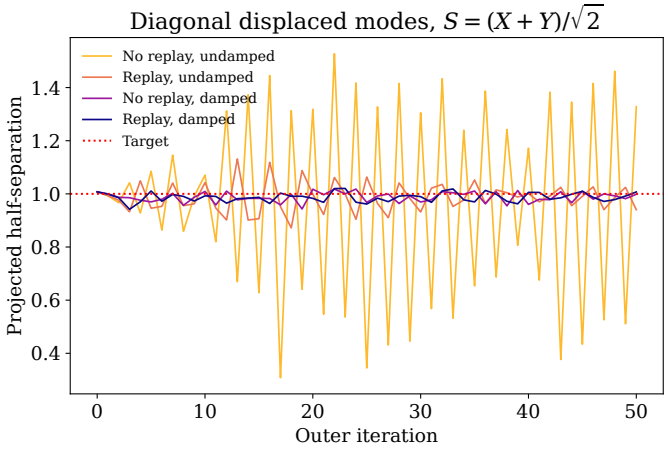


Figure 6: Projected half-separation $\widehat{\nu}_k^{\parallel}$ for the diagonal two-mode experiment. The four curves show the replay/damping variants for seed 0; the red dotted line is the target value 1. Seed 0 is selected mechanically as the unique common seed for which both projected basins are represented at every iteration under all four variants. In the no-replay undamped experiment, seeds 1 and 2 contain respectively 12 and 7 iterations at which one projected basin contains no observations among the 10^4 diagnostic samples, so Eq. equation 149 is undefined at those iterations.

of the diagnostic population, one projected basin can disappear entirely rather than merely oscillating in location. We refer to these events as empirical basin loss; they do not imply mathematically zero probability in the missing basin.

E.3.3 Nine-mode recovery and weight redistribution

We next consider a 3×3 grid of isotropic Gaussian components. Fix

$$(c_1, \dots, c_9) = \big((-2,-2), (-2,0), (-2,2), (0,-2), (0,0), (0,2), (2,-2), (2,0), (2,2)\big), \tag{150}$$

and let

$$\phi_j(x) = \mathcal{N}(x; c_j, 0.05 I_2), \qquad j = 1, \dots, 9. \tag{151}$$

Thus neighboring component centers are separated by 2, and each coordinate has the same component variance 0.05 used in the preceding one-dimensional multimodal experiments.

We consider two groups of problems. In the *mode-recovery* experiments, the target is the uniform nine-component mixture

$$p_{\mathrm{tar}} = \frac{1}{9} \sum_{j=1}^{9} \phi_j, \tag{152}$$

while the prescribed initial law contains only

$$\mathcal{S}_{\mathrm{center}} = \{(0,0)\}, \tag{153}$$
$$\mathcal{S}_{\mathrm{cross}} = \{(0,0), (-2,0), (2,0), (0,-2), (0,2)\}, \tag{154}$$
$$\mathcal{S}_{\mathrm{corners}} = \{(-2,-2), (-2,2), (2,-2), (2,2)\}. \tag{155}$$

The represented components have equal initial weights. These cases require the learned dynamics to create respectively eight, four, and five initially unrepresented modes.

In the *weight-redistribution* experiments, all nine components are present initially with equal weights $1/9$, but the target weights are changed. In the ordering of Eq. equation 150, the three target weight

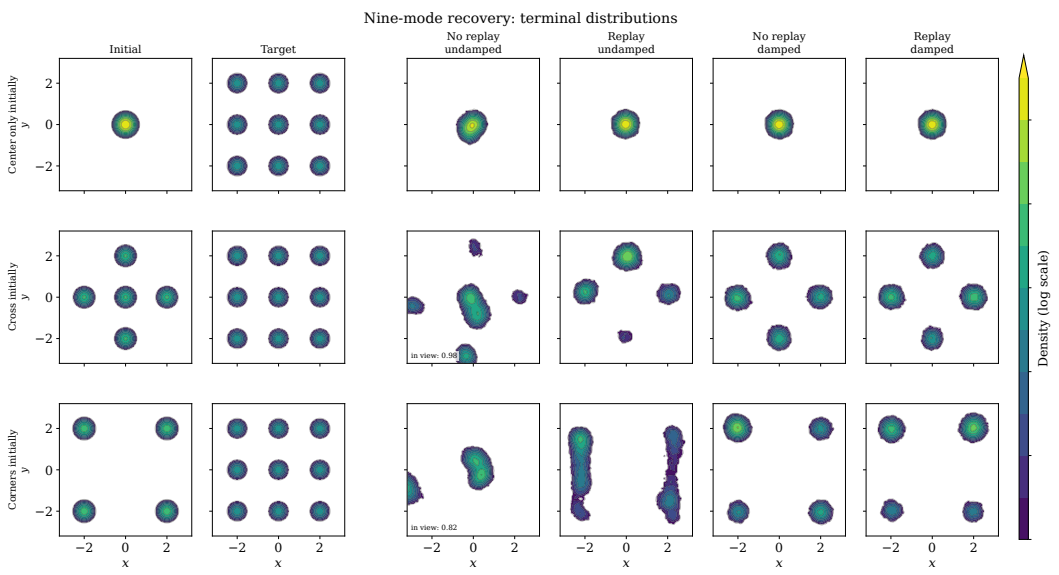


Figure 7: Final terminal densities for the three nine-mode recovery experiments. Rows correspond to center-only, cross, and corner-only initialization. Within each row, the first two panels show the prescribed initial law and target, followed by the four replay/damping variants. Learned densities are estimated from the final diagnostic populations pooled across the three seeds, and a common logarithmic density scale is used throughout the figure.

vectors are

$$w^{\mathrm{center}} = \frac{1}{18}(1,1,1,1,10,1,1,1,1), \tag{156}$$

$$w^{\mathrm{corner}} = \frac{1}{18}(10,1,1,1,1,1,1,1,1), \tag{157}$$

$$w^{\mathrm{graded}} = \frac{1}{42}(10,10,10,3,3,3,1,1,1). \tag{158}$$

The first target favors the center, the second favors the corner $(-2,-2)$, and the third imposes a graded allocation across the x-coordinate, with relative column weights 10:3:1.

Final terminal densities. Figures 7 and 8 show the final terminal distributions. For each learned variant, the displayed density is estimated from the final diagnostic samples pooled across the three seeds, giving 3×10^4 terminal samples per panel. Pooling is used here only to display the typical spatial structure at the final iteration; seed-to-seed temporal variation is retained explicitly in the dynamical diagnostic below.

The recovery contours show that initially absent components are not reliably reconstructed, even when several neighboring modes are already represented. The weight experiments isolate a distinct failure: all target locations are present from the outset, yet substantial allocation errors remain. These observations extend the one-dimensional distinction between local representation of individual modes and global redistribution of probability among separated basins.

Localized allocation discrepancy. Nearest-center classification alone can obscure a second source of error: probability lying diffusely between the target modes is necessarily assigned to some component. We therefore use a localized diagnostic. Let P_k denote the terminal law at outer iteration k. With localization radius

$$r = 0.5, \tag{159}$$

define

$$\ell_{j,k} = P_k(B_r(c_j)), \qquad j = 1,\dots,9, \tag{160}$$

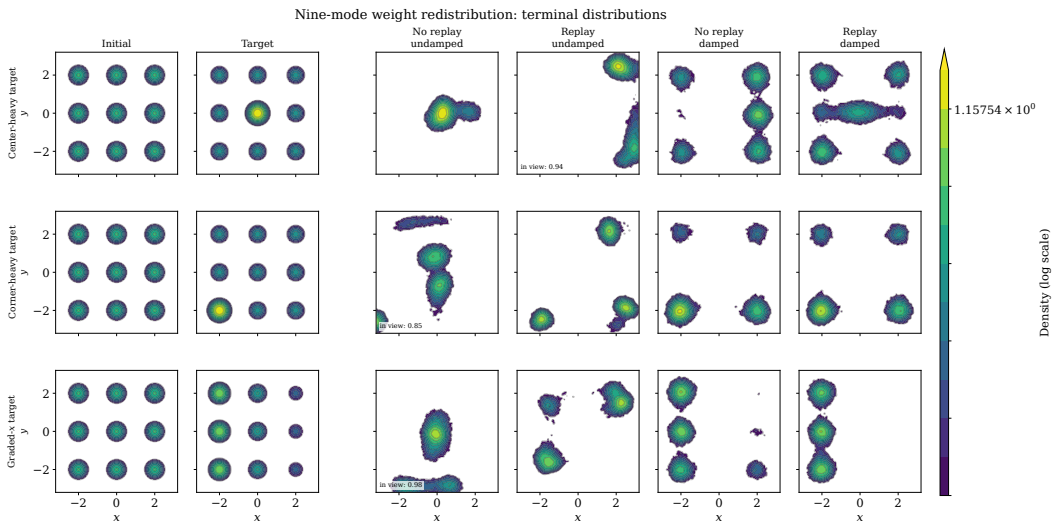


Figure 8: Final terminal densities for the nine-mode weight-redistribution experiments. Rows correspond to the center-heavy, corner-heavy, and graded-x targets in Eqs. equation 156–equation 158. Within each row, the first two panels show the uniform prescribed initial law and target, followed by the four replay/damping variants. Learned densities are estimated from final diagnostic populations pooled across the three seeds, with a common logarithmic density scale.

and let

$$\ell_{\mathrm{out},k} = 1 - \sum_{j=1}^{9} \ell_{j,k} \tag{161}$$

be the probability outside all nine localization balls. Define $\ell_{j,\star}$ and $\ell_{\mathrm{out},\star}$ analogously under the target terminal law $P_\star$.

The *localized allocation total variation* is the ordinary discrete total-variation distance between these ten-category probability vectors:

$$D_{\mathrm{loc}}(k) = \frac{1}{2}\left[\sum_{j=1}^{9} |\ell_{j,k} - \ell_{j,\star}| + |\ell_{\mathrm{out},k} - \ell_{\mathrm{out},\star}|\right]. \tag{162}$$

Thus $D_{\mathrm{loc}}(k) = 0$ means that the learned terminal law has the correct probability in every target-mode neighborhood and the correct off-mode probability. The diagnostic responds both to incorrect mode weights and to excessive diffuse or interstitial mass.

Figure 9 shows that substantial global allocation error persists across all six problems. More importantly, the time-resolved view reveals dynamics that are hidden by the final-density contours: several undamped trajectories exhibit large recurrent or oscillatory changes in allocation rather than monotone relaxation toward the target. We conjecture that the outer-iteration instability identified in the Gaussian analysis contributes to the nonmonotone character of these nine-mode trajectories, superimposed on the weak mechanism for redistributing probability among separated modes. The final contour plots alone do not establish this interpretation, but the corresponding time-resolved allocation dynamics are consistent with it.

E.4 High-dimensional Gaussian experiments

We test the same four failure modes in the practical neural-network implementation of AS in a high dimensional setting, specifically for p^{tar} we embed the 1D single Gaussian $\mathcal{N}(1, 0.05)$ in a 128-dimensional space (128D) resulting in a multivariable Gaussian distribution $\mathcal{N}(\boldsymbol{m}, \Sigma)$ with mean $\boldsymbol{m} \in \mathbb{R}^{128}$, $m_0 = 1$ and $m_i = 0$ for $i \in [1, 127]$; and diagonal covariance $\Sigma = 0.05\boldsymbol{I}$, with $\boldsymbol{I}$ identity matrix in $\mathbb{R}^{128\times128}$. For the symmetric two-mode Gaussian p^{tar} , we use 'left' mean

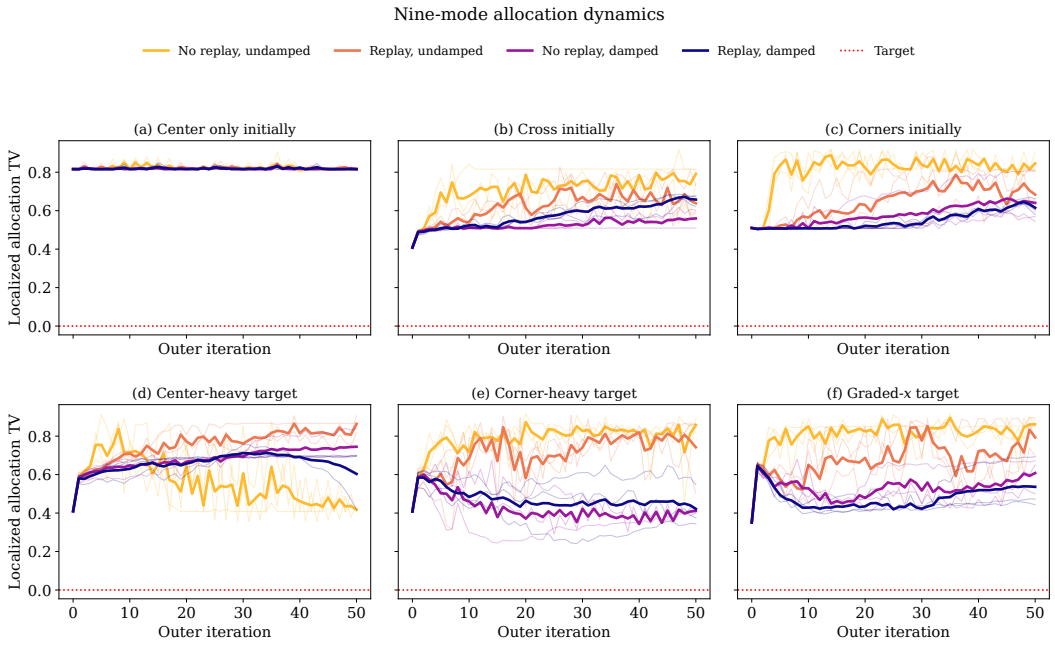


Figure 9: Localized allocation discrepancy $D_{\mathrm{loc}}(k)$ from Eq. equation 162. Top row: center-only, cross, and corner-only mode-recovery experiments. Bottom row: center-heavy, corner-heavy, and graded-x weight-redistribution experiments. Thin curves show individual seeds and thick curves show the three-seed mean; red dotted lines denote the target value $D_{\mathrm{loc}} = 0$.

$\boldsymbol{m}_L \in \mathbb{R}^{128}$, $m_{L0} = -0.8$, $m_{Li} = 0$ for $i \in [1, 127]$ and 'right' mean $\boldsymbol{m}_R \in \mathbb{R}^{128}$, $m_{R0} = 0.8$ and $m_{Ri} = 0$ for $i \in [1, 127]$, and same covariances $\Sigma = 0.05\boldsymbol{I}$. For initialization $p^{(0)}$, we use same covariances $\Sigma = 0.05\boldsymbol{I}$; for single Gaussian, we use $m_0 = 1.05$, while for symmetric two-mode Gaussian, we use 'left' mean $m_{L0} = -0.925$, and 'right' mean $m_{R0} = 0.925$. We follow an analogous configuration of parameters as reported in Havens et al., 2025 for the synthetic energy functions, i.e., we use a neural network with 3 layers and 128 hidden features. We train AS for 1000 outer iterations, generating 512 new samples and energy evaluations per iteration into a buffer. In each iteration we optimize on 500 batches of batch size 512 from the replay buffer. We use a larger buffer of max size 20000 (max size of 10000 was used in the original work) and a slightly smaller learning rate of 10^{-4} (3×10^{-4} was used in the original). For the single Gaussian experiments, we use a geometric noise schedule with $\sigma_{\min} = 10^{-4}$ and $\sigma_{\max} = 2.27 \times 10^{-1}$, while for the two-mode Gaussian experiments we use a geometric noise schedule with $\sigma_{\min} = 10^{-3}$ and the same $\sigma_{\max} = 2.27 \times 10^{-1}$. A summary of results, ran with three different seeds, can be found in Fig. 10. We observe that samplings come close to the means set, with some stiffness after initialization to $p^{(0)}$ (a-b), and fail to find the missing mode (c), or tend to favor one mode (d) corresponding to the light skewness obtained for $p^{(0)}$, that has about 5% more mass in the 'right' than in the 'left' mode.

F GENERALIZATION

We test whether the observed failure modes persist beyond Gaussian targets using two complementary numerical settings: density-level quadrature and Fokker–Planck experiments for non-Gaussian potentials, and neural-network experiments with logit–Beta targets.

F.1 DENSITY-LEVEL NON-GAUSSIAN EXPERIMENTS

We use exact numerical quadrature and solver, detailed in Appendix D.3, to investigate the generalization of the observed failure modes. In particular, we consider the following non-Gaussian models:

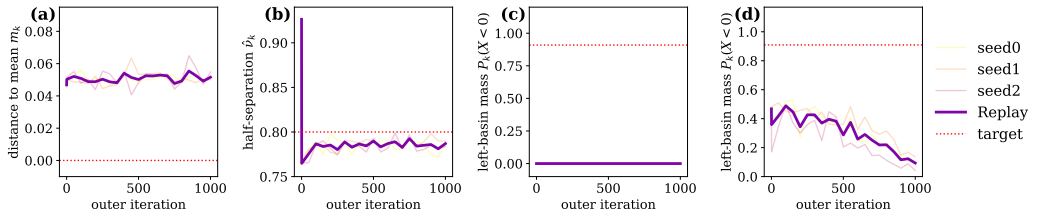


Figure 10: Neural-network Adjoint Sampling for the four failure-mode tests for the 128D case. Thin curves show individual seeds and thick curves show the three-seed mean, dashed red lines denote target values. All the runs use replay buffer. (a) Distance to single-Gaussian terminal mean $\boldsymbol{m}_k$. (b) Half-separation $\widehat{\nu}_k = \|\boldsymbol{m}_{R,k} - \boldsymbol{m}_{L,k}\|/2$ for the symmetric two-mode problem where $\boldsymbol{m}_{L,k}$ and $\boldsymbol{m}_{R,k}$ are the empirical conditional means in the left and right basins. (c) Left-basin probability $P_k(X < 0)$ for missing-mode discovery. (d) Left-basin probability $P_k(X < 0)$ for the incorrect-weight problem, whose target is 10/11.

- For the single-component experiment testing the instability (cf. Figure 2(a)), we replace the Gaussian mixture with anharmonic potential, with the log-density of the form

$$\log p^{\text{tar}}(x) = 0.01x^2 + 3x^4 + \mathcal{C}_{\text{AM}}, \tag{163}$$

  with a shifted initial distribution:

$$\log p^{(0)}(x) = 0.01(x - 0.01)^2 + 3(x - 0.01)^4 + \mathcal{C}_{\text{AM}}, \tag{164}$$

  where $\mathcal{C}_{\text{AM}}$ is from the normalization constant ensuring that p^{tar} and $p^{(0)}$ integrate to one. Fig. 11 shows the result of this experiment.

- For the two-component experiment testing the two-mode oscillation (cf. Fig. 2(b)), missing-mode discovery (cf. Fig. 2(c)), and incorrect-weight allocation (cf. Fig. 2(d)), we replace the Gaussian mixture with a Ginzburg–Landau potential, with the log-density of the form

$$\log p^{\text{tar}}(x; \kappa, \nu, h) = \frac{\kappa}{4}\left(x^2 - \nu^2\right)^2 + \xi x + \mathcal{C}_{\text{GL}}, \tag{165}$$

  where $\mathcal{C}_{\text{GL}}$ is from the normalization constant ensuring that p^{tar} integrates to one. Fig. 12 illustrates the results of the following settings:

  (a) Two-mode oscillation: $(\kappa, \nu, \xi) = (3, 2, 0)$, with an initial distribution taking the same form but with $(\kappa, \nu, h) = (3, 1.95, 0)$.

  (b) Missing-mode discovery: $(\kappa, \nu, \xi) = (3, 2, -0.8)$. Due to the asymmetry introduced by the linear term hx, the initial condition is defined as: $p^{(0)} \propto \mathbf{1}_{\{x \geq 0\}} p^{\text{tar}}(x)$; here, $\mathbf{1}_S$ denotes the indicator function of the set S.

  (c) Incorrect-weight allocation: $(\kappa, \nu, \xi) = (3, 2, -0.8)$. Due to the asymmetry introduced by the linear term hx, the initial condition is defined as: $p^{(0)} \propto 0.5 p_{\text{left}}(x) + 0.5 p_{\text{right}}(x)$, where

$$p_{\text{left}}(x) := \frac{\mathbf{1}_{\{x<0\}} p^{\text{tar}}(x)}{\int_{-\infty}^{0} p^{\text{tar}}(x')\, dx'}, \quad p_{\text{right}}(x) := \frac{\mathbf{1}_{\{x \geq 0\}} p^{\text{tar}}(x)}{\int_{0}^{\infty} p^{\text{tar}}(x')\, dx'}. \tag{166}$$

  (d) Complex update dynamics: target $(\kappa, \nu, \xi) = (3, 2, -0.8)$ and initial distribution with $(\kappa, \nu, \xi) = (3, 2, 0.8)$. The initial density is a left-right mirror image of the asymmetrical bimodal target density. This setup shows a complex update dynamic of oscillation, switching between bi- and single-modal distributions, and over-shooting instability between successive AS updates.

F.2 Neural-network logit–Beta experiment

As a neural-network counterpart to the density-level tests above, we repeat the four failure-mode experiments with a non-Gaussian component family. Let

$$Z \sim \text{Beta}(\alpha, \beta), \qquad X = \text{logit}(Z) = \log \frac{Z}{1 - Z}, \tag{167}$$

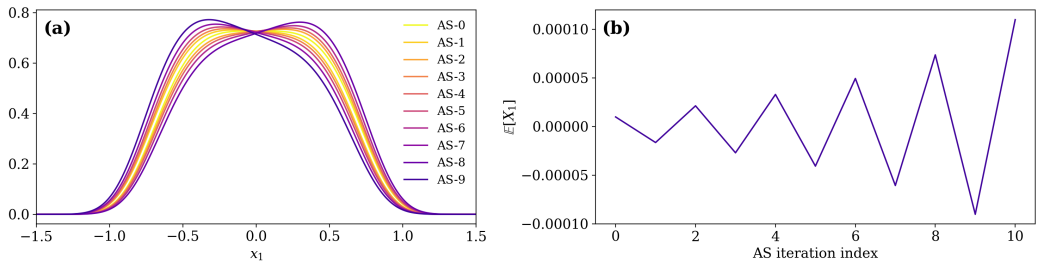


Figure 11: Numerical results for the anharmonic potential illustrating the failure modes of the adjoint matching method. (a) Numerical solution showing the evolution of the distribution oscillating around the target distribution and diverge over time. (b) The mean of the distribution as a function of AS iteration indicates clear instability and divergence over the mean value of the target distribution.

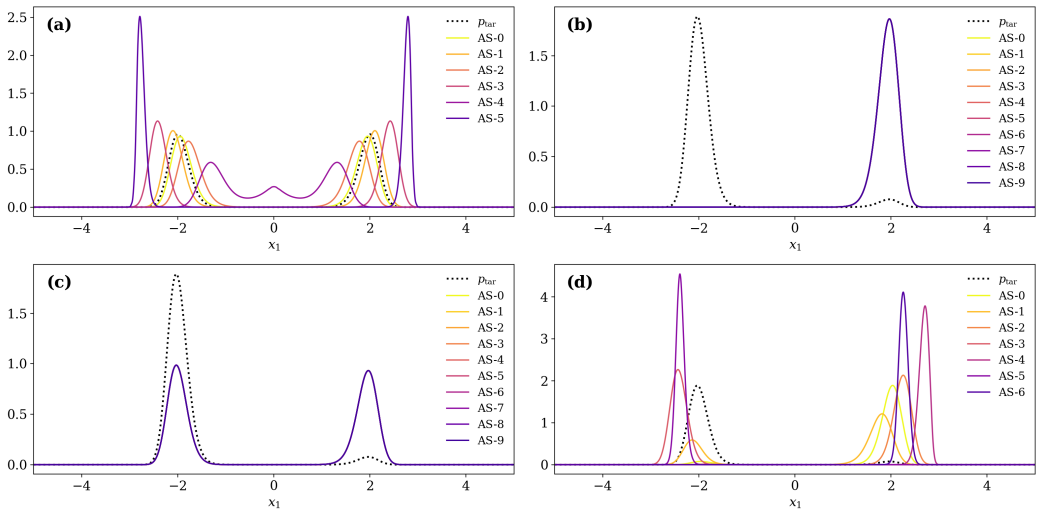


Figure 12: Numerical results for the Ginzburg-Landau potential illustrating the failure modes of the adjoint matching method. Numerical solutions reveal that (a) the evolution of the AS distribution oscillating around the two-mode target distribution and diverge over time, (b) failure of AS discovering missing mode, and (c) failure of AS to correctly allocate weights among the modes.

and write

$$s(x) := \frac{1}{1+e^{-x}} \tag{168}$$

for the logistic map. Adjoint Sampling is performed in the unconstrained coordinate $X \in \mathbb{R}$, using the same Brownian reference process as in the Gaussian experiments. Thus the sampler targets the induced *logit–Beta* density of X; the variable $Z = s(X)$ is used below to report diagnostics in the native Beta coordinate.

Under the transformation in Eq. equation 167, the density of X is

$$q_{\alpha,\beta}(x) = \frac{s(x)^{\alpha}\,(1-s(x))^{\beta}}{B(\alpha,\beta)}. \tag{169}$$

We parameterize the family by a sampling-coordinate mode μ and a concentration parameter κ:

$$m_\mu = s(\mu), \qquad \alpha_\mu = \kappa m_\mu, \qquad \beta_\mu = \kappa(1-m_\mu). \tag{170}$$

The corresponding score is

$$\partial_x \log q_{\mu,\kappa}(x) = \kappa\,(m_\mu - s(x))\,, \tag{171}$$

so $x = \mu$ is the mode and

$$-\partial_x^2 \log q_{\mu,\kappa}(\mu) = \kappa m_\mu (1 - m_\mu). \tag{172}$$

We choose

$$\kappa = \frac{1}{0.05\, s(1)\,(1 - s(1))} \approx 101.7232, \tag{173}$$

so that the components centered at $\mu = \pm 1$ have local log-density curvature 20, matching the Gaussian components of variance 0.05 in the one-dimensional experiments of Appendix E.2. The two families therefore have the same local quadratic stiffness near the relevant modes, while the logit–Beta score is globally nonlinear and saturates in the tails.

Writing $q_\mu := q_{\mu,\kappa}$, the four sampling-coordinate target/initial pairs are

Experiment	p^{tar}	$p^{(0)}$
Single component	q_1	$q_{1.01}$
Displaced modes	$\frac{1}{2}q_{-1} + \frac{1}{2}q_1$	$\frac{1}{2}q_{-1.01} + \frac{1}{2}q_{1.01}$
Missing mode	$\frac{1}{2}q_{-1} + \frac{1}{2}q_1$	q_1
Incorrect weights	$\frac{10}{11}q_{-1} + \frac{1}{11}q_1$	$\frac{1}{2}q_{-1} + \frac{1}{2}q_1$

Equivalently, each component corresponds in the native coordinate to $Z \sim \mathrm{Beta}(\alpha_\mu, \beta_\mu)$.

We use the common initialization and measured-training protocol of Appendix E.1. Analytic samples from $p^{(0)}$ are generated by sampling the corresponding Beta law in Z and mapping them to $X = \mathrm{logit}(Z)$; all other optimization, replay, damping, and diagnostic settings are unchanged.

Figure 13 reports the resulting trajectories after mapping the generated terminal samples back to the native coordinate $Z = s(X)$. For the single-component experiment we plot the terminal mean $\mathbb{E}[Z]$. For the symmetric displaced-mode experiment we plot $\mathrm{Var}(Z)$ as a native-coordinate measure of mode spread. For the missing-mode and incorrect-weight experiments we plot $P(Z < 1/2)$. Because the logistic map is monotone,

$$P(Z < 1/2) = P(X < 0), \tag{174}$$

so the latter two diagnostics retain exactly the basin interpretation of the Gaussian experiments.

The same qualitative division seen for the Gaussian components persists. In the single-component experiment, direct velocity damping keeps the native terminal distribution close to its target: the final-20 native-coordinate Wasserstein errors are approximately 0.0047 without replay and 0.0045 with replay. The symmetric displaced-mode experiment likewise exhibits stable native-coordinate spread under damping, although stabilization of this summary statistic does not imply exact recovery of the full mixture.

The global probability-allocation experiments remain substantially more difficult. For the incorrect-weight problem, the target left-basin probability is approximately 0.909, whereas the largest mean value over the final 20 iterations among the four variants is about 0.649. In the missing-mode problem, both damped variants remain near 5×10^{-5} left-basin probability. Undamped trajectories can produce much larger, strongly seed-dependent excursions; in one no-replay run the left-basin probability approaches one.

The undamped logit–Beta missing-mode traces also display recurrent, nonconvergent behavior qualitatively similar to the Gaussian experiment. Their seed-averaged trajectories should therefore not be interpreted as convergence to an intermediate distribution. This persistence for a non-Gaussian component family is consistent with the conjecture that the recurrence is generated by the iterative mode-oscillation mechanism rather than by the exact Gaussian form of the component density. Thus replacing the Gaussian components by curvature-matched logit–Beta components preserves the observed distinction between local stabilization, recurrent undamped dynamics, and global probability redistribution.

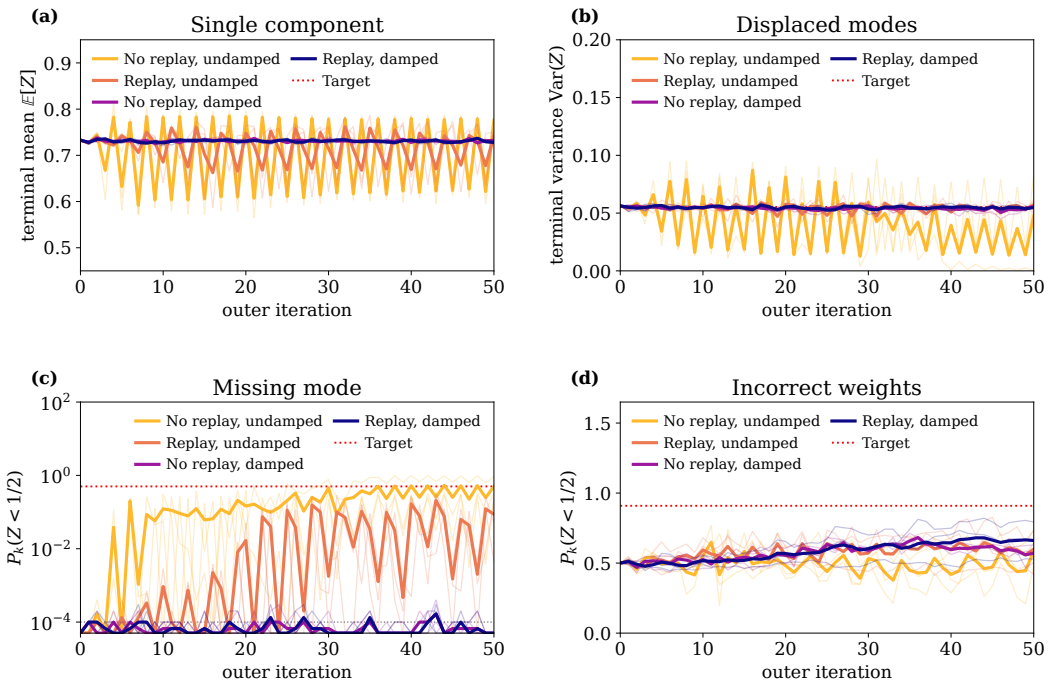


Figure 13: Neural failure-mode experiments for logit–Beta targets, with diagnostics displayed in the native Beta coordinate $Z = s(X)$. Adjoint Sampling is performed in the unconstrained logit coordinate X. Thin curves show individual seeds and thick curves show the three-seed mean; red dotted lines denote target values. (a) Terminal mean $\mathbb{E}[Z]$ for the single-component experiment. (b) Terminal variance $\mathrm{Var}(Z)$ for the symmetric displaced-mode experiment. (c) Missing-mode probability $P(Z < 1/2)$ on a logarithmic scale; the horizontal resolution marker corresponds to one observation among 10^4 diagnostic samples. (d) Left-basin probability $P(Z < 1/2)$ for the incorrect-weight experiment.